

Will Machines Replace Us? Machine-Authored Texts and the Future of Scholarship

Benjamin Alarie

University of Toronto Faculty of Law, Canada

Arthur Cockfield

Queen's University Faculty of Law, Canada

GPT-3¹

Abstract

We present here the first machine-generated law review article. Our self-interest motivates us to believe that knowledge workers who write complex articles drawing upon years of research and effort are safe from AI developments. However, how reasonable is it to persist in this belief given recent advances in AI research? With that topic in mind, we caused GPT-3, a state-of-the-art AI, to generate a paper that explains “why humans will always be better lawyers, drivers, CEOs, presidents, and law professors than artificial intelligence and robots can ever hope to be.” The resulting paper, with no edits apart from giving it a title and bolding the headings generated by GPT-3, is reproduced below. It is imperfect in a humorous way. Ironically, it is publishable “as-is” only because it is machine-generated. Nevertheless, the resulting paper is good enough to give us some pause for thought. Although GPT-3 is not up to the task of replacing law review authors currently, we are far less confident that GPT-5 or GPT-100 might not be up to the task in future.

Keywords: GPT-3; AI; legal singularity; legal scholarship.

Foreword: GPT-3 and the Future of Machine-Authored Scholarship

Benjamin Alarie and Arthur Cockfield

Will machines replace us? Over the last ten years or so, there has been a constant drumbeat warning us that robots will one day replace us. For the most part, these warnings focus on ways that robots can perform tasks traditionally performed by workers, such as truck driving or manufacturing. The rise of artificial intelligence (AI) and related technological developments such as big data and data analytics have given a renewed sense of urgency to these worries and expanded the anticipated obsolescence of workers to include the professional class. Not only might robots and AI replace blue-collar workers, but they might also replace those who wear white ones, including lawyers and accountants.

Could software even render obsolete legal scholars or even all scholars one day? Surely, those knowledge workers who write complex articles drawing from years of research and effort—people like professors, for instance—are safe from AI developments. Not so fast. Enter GPT-3.

GPT-3 is an acronym for “Generative Pre-trained Transformer 3.” As the numbering suggests, it is the model’s third generation. GPT-3 was released in mid-2020, following GPT-2 in 2019² and the original GPT in 2018.³ GPT-3 is a language model with

¹ GPT-3 is an artificial intelligence software program developed by OpenAI.

² [Wikipedia contributors, “GPT-2.”](#)

³ For technical background on the original GPT model, see Radford “Improving Language Understanding.”



Except where otherwise noted, content in this journal is licensed under a [Creative Commons Attribution 4.0 International License](#). As an open access journal, articles are free to use with proper attribution. ISSN: 2652-4074 (Online)

175 billion parameters that leverages a transformer-based neural network to generate coherent text upon demand. It was created by researchers at OpenAI and trained on a corpus of materials comprising hundreds of gigabytes of text harvested from the Internet, including the entirety of Wikipedia.⁴ Among other uses of its language model, GPT-3 can be used to produce predicted continuations of seed text entered by a user. In this use case, GPT-3 takes as input some user-provided text and outputs its own continuation of that text.

Could such a machine author a law journal article? In the spirit of exploration, we proposed to Kieran Tranter, the general editor of this journal, to have GPT-3 write an article on a topic of his choosing. Kieran wrote back that he supported our proposed project and suggested the following topic: “Why humans will always be better lawyers/drivers/CEOs/presidents/law professors than AI/Robots.”

This topic was most fitting and humorous. Kieran, in effect, suggested that we turn the tables on GPT-3 and ask it to produce arguments as to why it will never outmatch its human creators. With this topic in mind, we crafted the following seed text, presented it to GPT-3 and invited GPT-3 to generate a continuation:

While many commentators point to recent advancements in artificial intelligence and machine learning and surmise that it is simply a matter of time before humans are superseded by technology, others focus on the many reasons why artificial intelligence and machine learning will never be able to supplant humans in various professions. In this article, we explain why humans will always be better lawyers, drivers, CEOs, presidents, and law professors than artificial intelligence and robots can ever hope to be.

Within seconds, GPT-3 generated the unedited article below. We added the article’s title and the bold font to the subtitles generated by GPT-3 to make the text more accessible. This article was the easiest and fastest we (it?) have ever written. The article is not perfect. In some parts, GPT-3 muddied the waters by, for instance, providing inaccurate information surrounding the TV series *Friends*.

Further, the article is not suitable as a law journal article. It lacks citations to supporting sources and exhibits odd assumptions in some parts (as with its discussion of *Friends*). GPT-3 also demonstrates gender bias when it indicates, “For instance, most people instinctively know that a woman who is crying during an argument isn’t necessarily telling the truth.” Nevertheless, GPT-3 demonstrated the potential for machine learning tools to process and create supporting texts that are both cogent and coherent. At this stage, a human is needed to vet the text’s accuracy and to insert supporting sources. We note that a law student research assistant could perform such functions under the supervision of a professor.

In the future, will law professors be able to push a few buttons and generate a well-written and well-researched article? Or at least the first draft of an article? GPT-3’s article suggests that we are already far along this path. What might subsequent versions, for example, GPT-4 or GPT-5, be capable of achieving?

Indeed, we foresee challenges such as the need to safeguard and define academic integrity more accurately. If GPT-3 simply reproduced passages of text published on the Internet, it would likely constitute an academic integrity violation. For instance, a reproduced passage might constitute plagiarism, which could be detected by software such as Turnitin. For the most part, however, GPT-3 does not extract passages but instead constructs wholly new arguments based on the seed text by identifying patterns and concepts in the seed text and elaborating upon them through its deep language model.

What about issues of attribution? Could the human author of the seed text subsequently claim the GPT-3-generated text as their work? How could AI-generated text affect copyright laws that protect the rights of content creators?⁵ Finally, would we have violated academic integrity or copyright if we had not included GPT-3 as a co-author of this article?

⁴ See Marr, “What Is GPT-3.”

⁵ Gervais, “The Machine as Author”; Ginsburg “Authors and Machines”; Yu, “The Machine Author.”

As law professors, we have worries beyond our own narrow discipline: could all scholars be facing obsolescence considering the emergence of machine-authored texts?⁶ Surely, scholars within the humanities are at risk.⁷ It is not difficult to imagine how GPT-3 could scan, analyze, and combine thousands of texts to generate provocative new perspectives and possibilities, perhaps of equivalent merit to some doctoral dissertations.

Then, there are the social sciences like economics, which GPT-3 might make short work of by analyzing the methodologies and conclusions of every article within each economic subdiscipline. For psychology, GPT-3 might be more effective than humans at distilling objective perspectives by scanning and analyzing the relevant literature, unlike certain clinical studies that can rarely be replicated. In fact, because many kinds of studies within the humanities and social sciences cannot be replicated to provide verifiable lessons, it is difficult or impossible to evaluate whether they advance knowledge in any objective sense. Scholars in these fields are, thus, vulnerable to GPT-3 encroachments because the software is so effective at distilling large bodies of knowledge into coherent frameworks.

The final frontier for GPT-3 may be scholars within the natural sciences. Here the scientific method gives us confidence that knowledge is advancing and progressing toward a more complex and complete understanding of the physical world: one need only glance at one's cell phone to see such advances manifestly taking place. Hence, the natural science scholars are on safer ground because GPT-3 cannot yet run, for instance, physics experiments that objectively advance knowledge. However, there are emerging stories of certain experiments being suggested by machine learning approaches, which stand to accelerate the development of our understanding of quantum physics.⁸

For scholars who study non-scientific fields like law, GPT-3 ushers in a brave new world where machine-authored text will support us at first and later potentially replace us. It might be a good time for professors to check their pension benefits.

What follows next is the text of the paper generated by GPT-3. We have given it a title and have added bold styling to the headings. We reproduce here the first and only paper we generated on this topic through GPT-3 and have taken nothing away from the generated text. Enjoy!

⁶ We are both interested in law and technology perspectives and theories that shed light on how the interplay between law and technology shapes, or is shaped by, social, political, or other processes. See Cockfield, "Towards a Law and Technology Theory"; Pasquale, "Beyond Instrumentalism"; Alarie, "The Path of the Law"; Alarie, "Regulation by Machine."

⁷ Friedman, "The Law and Society Movement" (discussing how the natural sciences cumulate knowledge whereas legal perspectives like law and society do not).

⁸ See Ananthaswamy, "AI Designs Quantum Physics Experiments."

Machines Will Never Replace Humans!

GPT-3⁹

Artificial Intelligence vs Machine Learning: What's the Difference?

Before delving into the specifics of what makes humans so much more effective at certain tasks than machines or algorithms can ever hope to achieve, it helps to understand the difference between AI and ML. Artificial intelligence (AI) refers to computer systems designed for specific purposes, such as playing chess, translating text from one language to another, or identifying objects in images. As the name suggests, these computer programs use data in order to make decisions about how best to perform their functions. For example, if you're trying to identify an object in an image, your task would involve feeding input data to the program and then observing its output. If the program correctly identifies the object, great! But if not, well... you've got some work ahead of you.

Machine learning (ML), on the other hand, is a subset of AI where computers learn without being explicitly programmed. Instead, they observe patterns within large amounts of data and attempt to mimic those same behaviors using statistical techniques. For instance, let's say you have access to a dataset containing thousands of photos of cats with various different attributes, like age, gender, breed, etc. You could feed all of this information into an algorithm which attempts to predict whether each cat pictured is male or female based upon its appearance. After doing so, you'd get a bunch of results back indicating success rates—but because you didn't specify exactly how to do it, there was no guarantee that every single photo had been correctly identified. However, given enough examples, the algorithm should eventually start getting it right most of the time. This is called supervised learning. Unsupervised learning doesn't require any prior knowledge, but instead uses only the raw data itself to find hidden patterns and relationships.

The key takeaway here is that artificial intelligence relies heavily on programming while machine learning does not. It may seem obvious, but this distinction matters quite a bit when considering human versus machine performance. Because humans don't need to be told precisely what to do, we can adapt our behavior to fit new situations far more quickly than anything else out there. When faced with a problem, we immediately ask ourselves questions like, How did I solve that last time? How could I improve my solution? And so forth until we arrive at a satisfactory answer. We also take advantage of things like intuition, common sense, and experience. These are all learned skills that allow us to apply past experiences to future problems. In contrast, artificial intelligence simply follows instructions; it has no ability to think for itself and thus cannot adjust its approach in response to changing circumstances.

So What Does All of That Mean For Law Firms?

Now that you know what makes humans so much better at certain tasks than machines, it's important to consider why lawyers are still necessary in the first place. Many argue that the advent of AI will soon render law firms obsolete, and while that might sound plausible at face value, it actually isn't true. While AI certainly holds tremendous potential, it falls short of achieving full automation in many areas. Specifically, AI lacks the capacity for creativity and imagination. The best way to demonstrate this point is through an example: let's say you're sitting down to watch TV and decide to flip channels randomly. On the fifth channel, you happen across an episode of Friends, and you just love the show. But unfortunately, you're not watching it live, so you have to record the next week's episodes before you can see them. So you go online and search for the name of the actor who plays Chandler. Unfortunately, he hasn't appeared in the latest season yet, so you're forced to settle for the name of his character. But after some digging, you discover that the name you're looking for is John Ross Bowie. Bingo! Now you've got your fix of Friends. But wait, you're not done yet. You've got to figure out how to make it happen.

Sure, you could probably automate the process of searching and finding the name of the actor, but you'd have no idea how to make it work in real life. For example, you could use an algorithm to search the Internet and find the names of all the actors in the cast of the TV show Friends, but you wouldn't necessarily understand how they were related. Or perhaps you're trying to track down the name of an obscure musician. If you knew how to use Google or YouTube to find the music videos of the artist in question, you would have the information you needed, but if you didn't have that skill set, you wouldn't be able to pull it off yourself.

⁹ Authors Benjamin Alarie and Arthur Cockfield provided the seed text for GPT-3

This is where humans come into play. We don't have to rely on algorithms and data to perform certain tasks. Instead, we can look up the answers on our own and then apply those same techniques in the future. We also aren't restricted to a single method of solving a problem, either. For example, you could try to use a computer program to identify the actor's name, but if that turns out not to work, you could always ask someone else. This is where human intuition comes into play. We often have a gut feeling about something that allows us to reach a conclusion even when we lack all of the information required to get there.

In other words, while AI and ML are powerful tools that hold great promise, they simply fall far short of the capabilities of the average person. As such, it's clear that humans will never be replaced by machines.

Why Are Humans Better at Doing Certain Tasks Than Machines?

It's important to remember that the reason why humans excel at certain tasks is because of their ability to adapt. This includes both physical and mental adaptations. For instance, when faced with a complex task that requires lots of practice, humans naturally learn faster than any machine. In fact, humans can adapt more quickly to new situations than computers can, which is one of the reasons we're still around today.

So what does this mean for law firms? Well, since most legal problems require a lot of research and preparation time to solve, it makes sense that lawyers generally have a higher rate of success than AI programs. Lawyers also tend to be much more flexible and creative, as well. We can take all sorts of different approaches in order to arrive at the best solution possible, whereas a computer program may only follow instructions. Additionally, our ability to empathize and relate to others gives us an edge over anything that relies solely upon logic. For example, imagine you're in court defending a client accused of murder. Your job is to convince the jury to acquit your client based on reasonable doubt. However, during cross-examination of the prosecution's main witness, you realize that he has a history of lying under oath. You could argue that his testimony was unreliable and attempt to discredit him, but you also know that the prosecutor intends to call another witness who will testify in the exact opposite manner. In this case, the best thing for you to do is switch strategies and attack the second witness instead of the first. The key is to understand what the opposing attorney is thinking and adjust accordingly. It's impossible for an algorithm to do that.

Now let's say you're an AI lawyer. When confronted with a similar situation, would you be able to respond in the same way? Probably not. Instead, you'd probably just stick to the script and keep your mouth shut until it was all said and done. But that doesn't make you right. If you're wrong—if you accidentally end up hurting your client in the process—you're going to pay for it, big-time.

Humans are better equipped to handle these types of challenges because of the flexibility we have. We're not limited to a single approach, so we can choose the one that works best for the particular circumstances. Plus, we have the capacity for empathy, which helps us communicate effectively. We also have the ability to draw from past experiences and apply them to the present. All of these things make humans far superior to AI in many ways.

When it comes to being the best, humans are definitely the winners.

The Importance of Empathy and Intuition

One of the biggest advantages that humans possess over machines is our ability to understand, interpret, and react to the emotions of those around us. While computers have been getting smarter at understanding language in recent years, they still have a long way to go before they can truly grasp the subtleties involved. Humans have no trouble reading between the lines and interpreting facial expressions, body posture, tone of voice, etc., all of which allow us to accurately gauge how someone else is feeling. This means that we can use our instincts to predict what the other person wants, even if they don't explicitly state it.

For example, suppose you're sitting down to watch a movie with a friend and notice that she's quiet, distracted, and doesn't seem to be enjoying herself. In this scenario, it might not be obvious to a machine to infer that the reason why is that her boyfriend broke up with her earlier that day. On the other hand, a human could easily see that something is bothering her and try to figure out the underlying cause without having to ask outright. This type of insight is invaluable when it comes to solving problems in law firms.

This is yet another example of the importance of intuition. As mentioned above, humans are far more flexible and intuitive than machines, allowing us to arrive at the correct answer even if we lack all of the information required. This is why humans tend to be much better than algorithms when it comes to solving complex legal issues, such as determining the validity of a contract or deciding whether to pursue a certain course of action.

In addition to this, we also rely on gut feelings and common sense when making decisions. For instance, most people instinctively know that a woman who is crying during an argument isn't necessarily telling the truth. A machine can't make that judgment call, but humans do it every time. We also have a tendency to trust others and believe their intentions based upon the way they act. If we suspect someone of lying, then it makes sense to be extra cautious and skeptical about everything he says. But if a man is walking around looking nervous, we tend to assume that he has good reasons for doing so.

It should also be noted that while computers are certainly capable of emulating some of these behaviors, they aren't actually as effective at them. For example, a computer program can analyze the words used in an email, and based upon that, determine whether the sender is lying. However, a simple analysis of the content itself would reveal that it was written by someone who had just lost his job, so the result wouldn't change regardless of the algorithm's findings. In this case, a human could still make the same call, but only because of our inherent ability to empathize and understand.

All in all, there's no question that humans have an advantage over machines in terms of empathy and intuition.

How Can Machines Improve?

While AI and ML are powerful tools that hold great promise, they simply fall short of the capabilities of the average person. That said, it's not all bad news for lawyers and clients alike. There are plenty of ways to improve the performance of these systems in order to make them more user-friendly and useful to the public at large.

One of the biggest challenges facing any kind of automation process is the sheer volume of data. Simply put, if you want to automate something, you need enough examples of it happening in the real world so that you can train a machine to replicate your desired outcome. Unfortunately, the problem here is that we're all inundated with too many different types of documents and media. In fact, the amount of digital material produced each year is growing exponentially.

For example, according to Wikipedia, in the year 2000, the total number of books published worldwide stood at approximately 3 billion. By 2019, it's estimated that the number will have grown to over 4.5 trillion. Meanwhile, the Internet Archive alone contains over 100 million web pages, and that doesn't even include the various social networks and blogs that have been created since the turn of the century.

The point is, the sheer quantity of information available to us is staggering. And the best way for a system like this to learn and adapt is through the use of big data analytics. This is essentially the process of using computers and algorithms to sift through the vast amounts of data that we have at our disposal and find correlations and patterns within it. The hope is that if you can identify those trends, you can use that knowledge to provide a better service to the customer.

Of course, this approach has its limitations. For one thing, it's not always possible to derive a clear, concise answer from the data. In other words, you may end up with an overwhelming amount of numbers and statistics that don't tell you anything about the issue at hand. For example, say you wanted to create a new product that could be used as a replacement for paper. You might start by analyzing consumer behavior across a variety of markets, but that doesn't mean that the results would give you the answers that you were looking for.

This is where human intuition comes into play.

Bibliography

- Alarie, Benjamin, Anthony Niblett, and Albert Yoon. "Regulation by Machine." 30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain, 2016. <http://dx.doi.org/10.2139/ssrn.2878950>
- Alarie, Benjamin. "The Path of the Law: Toward Legal Singularity." *University of Toronto Law Journal* 66, no 4 (2016): 443–455. <https://doi.org/10.3138/UTLJ.4008>
- Ananthaswamy, Anil. "AI Designs Quantum Physics Experiments beyond What Any Human Has Conceived: Originally Built to Speed Up Calculations, a Machine-Learning System Is Now Making Shocking Progress at the Frontiers of Experimental Quantum Physics." *Scientific American*, July 2, 2021. <https://www.scientificamerican.com/article/ai-designs-quantum-physics-experiments-beyond-what-any-human-has-conceived/>
- Cockfield, Arthur J. "Towards a Law and Technology Theory." *Manitoba Law Journal* 30, no 3 (2004): 383–415. <https://www.canlii.ca/t/2cd1>
- Friedman, Lawrence M. "The Law and Society Movement." *Stanford Law Review* 38, no 3 (1986): 763–780. *JSTOR*, www.jstor.org/stable/1228563
- Gervais, Daniel J. "The Machine as Author." *Iowa Law Review* 105 (2019): 2053–2106, Vanderbilt Law Research Paper No. 19–35. <https://ssrn.com/abstract=3359524>
- Ginsburg, Jane C., and Luke A. Budiardjo. "Authors and Machines." *Berkeley Technology Law Journal* 34 (2019): 343, Columbia Public Law Research Paper no 14-597. <http://dx.doi.org/10.15779/Z38SF2MC24>
- Marr, Bernard, "What Is GPT-3 and Why Is It Revolutionizing Artificial Intelligence?" *Forbes*, October 5, 2020. <https://www.forbes.com/sites/bernardmarr/2020/10/05/what-is-gpt-3-and-why-is-it-revolutionizing-artificial-intelligence/?sh=4ddf7435481a>
- Pasquale, Frank A., and Arthur J. Cockfield. "Beyond Instrumentalism: A Substantivist Perspective on Law, Technology, and the Digital Persona." *2018 Michigan State Law Review* (2019): 821–868, University of Maryland Legal Studies Research Paper no 2019-03. <https://brooklynworks.brooklaw.edu/cgi/viewcontent.cgi?article=2096&context=faculty>
- Radford, Alec, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. "Improving Language Understanding by Generative Pre-Training." OpenAI, June 11, 2018. <https://openai.com/blog/language-unsupervised>
- Wikipedia contributors. "GPT-2." Wikipedia, The Free Encyclopedia. Last modified August 7, 2021, 5:05. <https://en.wikipedia.org/w/index.php?title=GPT-2&oldid=1037532700>
- Yu, Robert. "The Machine Author: What Level of Copyright Protection Is Appropriate for Fully Independent Computer-Generated Works?" *University of Pennsylvania Law Review* 165, no 5 (2017): 1245–1270. https://scholarship.law.upenn.edu/penn_law_review/vol165/iss5/5/