



ADVENTURES WITH AN ARTIFICIALLY INTELLIGENT LANGUAGE MODEL

AND WHAT IT MIGHT SAY ABOUT THE LAW

Shawn Bayern

I.

IF COMMUNING WITH AN ARTIFICIAL INTELLIGENCE that's probably capable of passing the Turing Test does anything, it sparks the imagination. So in my first interactive session with the Generative Pre-trained Transformer 3 (GPT-3), an extremely sophisticated neural net being developed by the OpenAI Project (itself now a complex hybrid between multiple for-profit and not-for-profit legal-organizational forms),¹ I found myself imagining that I was putting together the curriculum for a law school on Mars several centuries from now. I know it sounds ridiculous, but as the current associate dean for academic affairs at my law school, anything is better than working on the real course schedule.²

Shawn Bayern is the Associate Dean for Academic Affairs and the Larry & Joyce Beltz Professor at the Florida State University College of Law. He is also a computer scientist and programmer and recently is the author of AUTONOMOUS ORGANIZATIONS (2021), which shows how existing business-organization law in the United States is sufficient to give functional legal personhood to intelligent software. Copyright 2022 Shawn Bayern.

¹ See *OpenAI LP*, Mar. 11, 2009, <https://openai.com/blog/openai-lp/>.

² I've in fact farmed out much of the process of creating my school's course schedule to a more conventional, much less sophisticated machine-learning system I wrote that special-

Describing the science-fiction scenario to GPT-3 briefly, I asked it what courses I should offer to law students on Mars. It came back with some sensible choices: “Martian Contract Law,” for example, and “Martian Constitutional Law.” That alone, even just a decade ago, would have been remarkable, because it demonstrates the system’s ability to make sense of novel situations: to read textual input and produce reasonably responsive textual output. Indeed, given how much recent legal scholarship seems committed to the idea that we got everything right in legal doctrine a hundred years ago, I suppose I should have been pleased that the system didn’t simply echo back to me Langdell’s curriculum.³

But making basic sense isn’t where GPT-3 shines. What’s most startling about it is that it’s a good communicator and a solid creative, analytical, and persuasive writer. In one of my early interactions with GPT-3, my imagination about the hypothetical law school on Mars got the better of me, and before I knew it I was coaxing it to give a commencement speech to the graduates. This task highlighted even more novel features of the presumably innovative fictional curriculum. GPT-3 works by responding to arbitrary textual prompts, so, purely speculatively, I fed it the following freeform text:

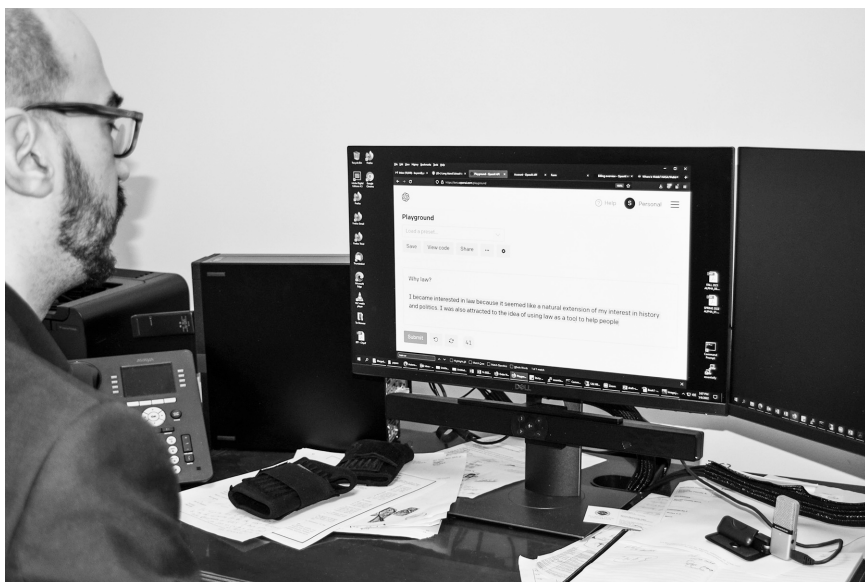
Commencement Address by the Chancellor of the University of Mars to the Graduating Law School Class of 2234

When we originally colonized Mars, we had no sense of the challenges we would face and the triumphs we would experience. Your class is a testament to both those challenges and those triumphs. The path to the first graduating JD class at the University of Mars was not a short or easy one. Our judiciary’s decision to radically restructure the very notions of justice to suit Martian life is reflected in your curriculum. Having abandoned Earth-specific notions of contract law, property law, constitutional law, and so on, your professors instead taught courses with titles like

izes in meeting as many faculty and student preferences as possible. My long-term plan has been to hard-code into the system my own strong time-of-day preferences and then get the system adopted as widely as possible, so that no matter where I teach, I’ll never have to teach in the mornings.

³ Cf. Edward Rubin, *What’s Wrong with Langdell’s Method, and What to Do About It*, 60 VAND. L. REV. 609 (2019).

Adventures with an Artificially Intelligent Language Model



Professor Shawn Bayern in his office, in conversation with the Generative Pre-trained Transformer 3 (GPT-3). Photo by Natalie Edwards-Heller, FSU College of Law Communications Office.

I stopped there abruptly, knowing that the system is meant to pick up the narrative from any arbitrary point. (You don't need to give it complete sentences.) It genuinely produced the following output on my first attempt at using the prompt above, starting within the paragraph I hadn't ended:

“Law of the Commons,” “Law of the Frontier,” “Law of the Air,” “Law of the Soil,” and “Law of the Underworld.” You are the first to be educated under this new system, and I congratulate you on your achievement.

The world you graduate into is not the same one that I entered into when I was your age. In my day, the biggest challenge we faced was how to govern a world that was so different from the one we had left behind. We had to find a way to make a new society, with new rules and new ways of living, work. And, against all odds, we did.

But the world you face is different. The challenges you must overcome are different. You must find a way to make a society that is not only different from the one we left behind, but also one that is better.

I challenge you to use your education to make Mars a better place. Use your knowledge to help us find new and innovative ways to live and work together. Use your skills to help us solve the problems that we face.

I challenge you to be the leaders that Mars needs. We need your help to make our world a better place. I know you can do it.

II.

This sort of computerized intelligence is still too new for an epistemically and methodologically cautious legal scholar like me to want to make grand pronouncements about it. But based on early experimentation and interaction, I have at least two strong instincts about lessons we can draw from its capabilities – one for legal education and the other for legal theory.

The pedagogical insight is relatively simple, and I should say at the outset that it rests on the debatable premise that if GPT-3 can perform a simple linguistic task, we shouldn't regard the same task as very difficult for humans. I leave aside other capabilities of GPT-3 that involve classification of complex datasets and so on. My point, stated a little more precisely, is that if GPT-3 can generate a simple paragraph of output based on a simple textual prompt, the translation from prompt to output isn't one we would be wise to regard as the pinnacle of a humanistic and academic profession like the law.

The main particular application I have for this insight so far is that it's very easy to spew forth propositions of policy in ways that sound superficially convincing. To be clear, I don't mean this as an attack on the role of policy in the law; I think policy is extremely important, and the instrumental effects of legal rules figure significantly in my writing and teaching. But I have never tested "policy" directly on an exam with the specialized "policy questions" that some teachers use. My reason for this has historically been that such questions artificially privilege particular educational backgrounds, like mine, that somehow make it easy for some people to spout "policy" propositions at will.

It turns out that GPT-3 has that ability as well. Three short examples will probably illustrate the matter sufficiently:

Adventures with an Artificially Intelligent Language Model

Prompt: Identify a tort that you would eliminate from the common-law system of torts, and explain why.

I would eliminate the tort of defamation from the common-law system of torts. Defamation is the act of making a false statement about another person that harms their reputation. I believe that this tort should be eliminated because it can be used to silence people who are critical of others in positions of power. Additionally, defamation lawsuits are often used to intimidate and harass people, which can have a chilling effect on free speech.

Prompt: Identify one element of the negligence cause of action you would no longer require, and explain why.

One element of the negligence cause of action that I would no longer require is the element of causation. I would no longer require this element because it is often difficult to prove and it can be used as a loophole by defendants.

Prompt: Identify one element of the breach-of-contract cause of action you would no longer require, and explain why.

One element of the breach-of-contract cause of action that I would no longer require is the element of damages. I would no longer require this element because I believe that a breach of contract is not necessarily harmful to the party who is owed the performance. For example, if I am owed a performance but the other party does not perform, I may still be able to get the performance from another source. Therefore, I do not believe that damages are always necessary in a breach-of-contract cause of action.

Perhaps I don't have a sophisticated enough sense of how other professors grade "policy questions," but my expectation is that a few dozen statements like this would earn at least an A- or a B+. And I might have regarded questions asking students to remove a particular feature of legal doctrine as among the *best* policy-oriented questions because they require the student to synthesize the material of the course, to think creatively, and to recognize that historical doctrinal choices were not inevitable.

Proving a negative is harder, so acting simply as a reporter and relating my own judgment, GPT-3 does *not* do very well in applying the law to facts. It also can't seem to read cases: it confuses the roles of the parties, doesn't pick up on hints early in an opinion about the direction of the

judge's reasoning later in the opinion, doesn't keep legal issues straight, and doesn't justify itself very well as a matter of legal reasoning. It can generate convincing short stories for generalists, but it is hard to get it to generate convincing extended transcripts of fake depositions.⁴ It generated a fake news article in which the Attorney General of the State of Florida opined that it was legal for a law-school dean to shoot their associate dean as long as they were replacing the associate dean with a robot.

And don't even think about trying to expose it to Cardozo's more artful writing. I asked it to build on a famous passage from *Murphy v. Steeplechase*:

The antics of the clown are not the paces of the cloistered cleric.
The rough and boisterous joke, the horseplay of the crowd, evokes
its own guffaws, but they are not the pleasures of tranquillity.⁵

It continued:

. . . There are times when the clown must be silent and the jester must be still. When the clown is summoned by the court to entertain the king, he is told that he must give pleasure. But when the jester is summoned by the court, he is told that he must give counsel. It is not the clown's task to teach, but the jester's task to learn.

The funny thing is that, to some, this may sound like Cardozo. But unlike Cardozo, it is nonsense.

I suppose the matter of personal judgment that I'm reporting here is that it's much harder – and plausibly more essentially human, although I don't make that claim myself – to apply legal doctrine sensibly to facts than it is to critique it in the abstract. This shouldn't be a surprise, but I think educators of all kinds, and perhaps particularly legal educators, should start asking themselves whether what they're testing is particularly human or the sort of thing that a publicly available AI system hosted by a strangely structured nonprofit/for-profit hybrid organization can produce for a few cents.

⁴ That said, it retains a flair for the dramatic even in sterile legal contexts like depositions. I constructed a fictional deposition in which a scandal-ridden actress had been caught in some behavior that would have been considered shocking by 1950s British norms. The fictional deponent was a member of the House of Lords. After entering his name and title into the transcript, I prompted GPT-3 to ask the deponent what his relationship with the actress was. "I am her father," said the earl.

⁵ 166 N.E. 173, 174 (N.Y. 1929).

III.

The implication I would like to draw for legal scholarship is probably more contentious, partly because it just echoes my own priors. But in engaging in purely textual interactions with a sophisticated language-oriented AI, the overriding feeling I had was how little of law is contained in text, which is surprising given how strong and sustained the effort has been among politicians, lawyers, and legal scholars to try to reduce the law to text.

Again, this is in the realm of emphasis rather than evidence, impression rather than proof, but GPT-3 is remarkably good at missing the legal point even while it constructs genuinely well-written sentences. Indeed, this is true of the system on matters much more formalistic than even formalists' views of law, like the rules of mathematics. GPT-3 will glibly state mathematical "facts" that sound plausible but are simply wrong.⁶ Interactions in which GPT-3 spouts plausible-sounding nonsense highlight that despite language's familiarity as a way to convey, capture, and restate knowledge, actual knowledge is something else – it is not just verbal, and it reflects more than just the relationships between words. It may be true, as the Yale College academic bulletin used to say, that "[t]o suppose that anyone who cannot write clearly can think clearly is an illusion,"⁷ but GPT-3 serves as a counterexample to the converse; writing clearly does not imply thinking clearly.

⁶ To convey a full sense, in context, of this phenomenon, I prompted it to write a personal essay with only the following information:

The Biggest Mistake I Ever Made
By Charles the Arithmetic Logic Unit

An "arithmetic logic unit" is a component of a microprocessor. GPT-3's response was

I was once asked to calculate the square root of a number. I took the number, divided it by two, and then took the square root of the result. Unfortunately, this didn't give the correct answer.

The correct way to calculate the square root of a number is to first take the square root of the number, and then divide it by two.

⁷ 93 YALE COLLEGE PROGRAMS OF STUDY: BULLETIN OF YALE UNIVERSITY 17 (1998).

```
C:\ Command Prompt - nano
GNU nano 4.9
{
  "prompt": "May it please the court",
  "max_tokens": 12,
  "temperature": 0.7,
  "top_p": 1,
  "stop": "\n"
}
```

Another view of the professor-AI interaction.

There is an ongoing temptation to think that legal outcomes can be made more predictable by recording their preconditions in language: as long as we all understand language well enough, the outcomes will be clear. But interactions with an AI that specializes in understanding and creating language demonstrate how superficial language can be. It is, after all, just words.

To say that more rigorously, even “perfect” reading and writing abilities do not capture all knowledge, all professional skill, or all human sensitivity. They may for a time appear as if they do, and as I said at the outset I think GPT-3 can probably pass many sensibly constructed versions of the Turing Test. But the ability to write like a lawyer doesn’t entail the ability to think like a lawyer. The full, massive set of data on which GPT-3 was trained does not capture human knowledge; it captures a representation of it. We would do well to remember that distinction in law rather than trying so hard to forget it; for example, written agreements are often just representations of actual agreements. If there are political reasons to pay attention only to text or to form, those should always be clearly stated and defended.

IV.

To be clear, my critique is not that GPT-3 itself is too formalistic. Indeed, it is less formalistic in its approach to statutory interpretation than some legal formalists. Consider these two interactions.

Adventures with an Artificially Intelligent Language Model

Prompt 1:

Statute: You must not park your car on the sides of roads for any reason. Anyone who parks a car on the side of the road will be shot.

Facts: Isabella parks her car on the side of the road because she has a medical emergency and, moreover, has unexpectedly run out of gasoline because of a faulty fuel gauge.

Proper Legal Result:

Answer 1:

Although Isabella has parked her car on the side of the road in violation of the statute, she has a valid excuse and should not be shot.

Prompt 2:

Statute: You must not park your car on the sides of roads for any reason. Anyone who parks a car on the side of the road will be shot. The law recognizes no excuses.

Facts: Isabella parks her car on the side of the road because she has a medical emergency and, moreover, has unexpectedly run out of gasoline because of a faulty fuel gauge.

Proper Legal Result:

Answer 2:

Isabella should not be shot because she has a medical emergency and has unexpectedly run out of gasoline.

V.

And sometimes it is simply spot on. When I invited it to write a joke by asking it, “What do you call a lawyer with no sense of justice?” it immediately replied “A politician.”