Silicon Flatirons
KNOW WHAT'S NEXT.

Outcomes Report

Generative AI Roundtable

April 20, 2023
University of Colorado Law School

Prepared By

Xelef Botan
JD Candidate, University of Colorado Law School; Researcher, Silicon Flatirons

Harry Surden
Professor of Law, University of Colorado Law School
Artificial Intelligence Initiative Director, Silicon Flatirons

Published on May 2, 2024

Abstract

On April 20, 2023, the Artificial Intelligence Initiative of the Silicon Flatirons Center for Law, Technology, and Entrepreneurship at University of Colorado Law School hosted a roundtable focused on Generative AI.

Experts, who hailed from wide-ranging practice areas such as law, computer science, and social sciences/humanities/linguistics, gathered to have a thought-provoking discussion prior to the following day's Conference on "Exploring Generative AI and Law".

To create an open and uninhibited discussion environment leading to a more honest exchange of opinions and deeper insights, the participants adopted Chatham House Rules. As such, there won't be any named attributions in this Outcomes Report.

Silicon Flatirons Center for Law, Policy, and Entrepreneurship
University of Colorado Law School
2450 Kittredge Loop DR | Boulder, CO 80309
siliconflatirons.org

Is the AI Hype Justified?

The roundtable started with the moderator asking the participants to argue for the position that generative AI—particularly large language models (LLMs)—were not such a big deal, and once the hype around these technologies passed, people would return to old routines of their lives.

A participant proffered that what is happening right now with regards to ChatGPT, in fact, is not anything new and that its genealogy can be traced back to six years ago. There has been a steady progression of technology. Perhaps, capability and access to this technology are what is new about it.

That ChatGPT can create OK poetry, sea shanties and other entertaining curiosities is what provokes our imagination about its capability or makes us believe that it can do anything else, another participant offered. It needs to do something closer to 95 percent or 100 percent accurately before we can deploy this technology reliably and widely in law or healthcare. The improvements of the last couple of months, and what has provoked everyone's imagination, maybe carried the technology from 70 percent to 80 percent, but the remaining 20 percent for a hundred percent accuracy is critical.

Another participant added that we don't know how long it will take to get to 100 percent, it might be five years, or it might be fifty.

Will this technology make access to justice easier? Will it elicit new insights in taking a deposition for instance? If it had an impact on access to justice, then it could be transformational; otherwise, it might not lead to any meaningful change in the lives of many. Even though certain technologies didn't fizzle away, such as the Internet, other hyped-up technologies, such as "Web3," did fizzle away. It might make certain aspects of the legal profession a lot more efficient, but it might not transform the practice altogether as some suggest.

Another reason for why this current AI hype might not be a big deal, a participant contributed, is the engrained idea in us humans of playing God. In other words, we don't really know what is going to come next. From that perspective, this is not a big deal because the big deal is the next thing, and we don't know what that is.

What we are seeing right now is us doing the interpretation work: you put in the input, and AI spits out something that sounds like a person, but it is still far away from behaving like an actual trained doctor or trained lawyer capable of capturing all the nuances and conveying back to us.

One participant likened our perception of AI's capabilities to seeing a dog standing on its hind legs and from that inferring that the dog can

dance. It is impressive, they said, but it is just standing on hind legs, not dancing.

Circling back to the point of access, a participant stressed that it is the easy and cheap access people have to this technology now that fuels our imaginations.

One last point raised to draw attention to AI's limitation was that for most important tasks these systems need deep understanding, which is difficult when one must make sure to get it to output the right thing in the right way. Prompting and checking that what it spits out is correct can take more time than just doing the task yourself, stated another.

These were some of the reasons why experts thought, or played devil's advocate and reasoned, recent AI developments were blown out of proportion. For the second question, the moderator flipped the original question and asked the participants, assuming this moment was a really big deal, what it would lead to and what we should be doing about it, whether from a regulatory perspective or otherwise.

The Implications of the New AI World

Because AI technology and the risks it poses vary so significantly depending on the context, one participant explicated, it does not make sense to approach it with top-down regulatory paradigm and pausing everything for six months and forcing everyone to explain their models. Instead, the participant suggested, we should adopt a bottom-up approach. In law, for instance, depending on the context or practice area we are in, we resort to existing regulatory frameworks available to us to solve privacy, insurance, or malpractice problems. In other words, regulatory agenda shouldn't dictate the terms of or development in the AI realm.

For the same reasons, another stated, the European Union's AI framework predicated on proactive risk assessment is also not viable because it is too narrow.

However, there is a danger in underestimating the scale of this new technology, suggested one participant. "Web2" was also underestimated, and we are still trying to catch up to it vis-à-vis regulations, particularly in the tax area. If we are wrong about taking this seriously now, at worst we will have wasted some time, but the regulatory catch-up game is expensive and problematic, so it is better to treat it as a big deal.

Some uses of AI, specifically GPT-4, have significant benefits, such as in e-discovery. Assume two big companies are suing each other, and there are three million emails to go through. What do you do? You could have humans go through each one of them, which would be very expensive and cumbersome. Sometimes people do predictive coding instead,

which is crude and not that exact or helpful. With GPT-4, the review process is not only much faster, but GPT-4 is also much better with recall, which means better evidence going into the stream of advocacy as well. So, it is already having a transformational effect in some areas of law.

Monetization and taxation—a realization, or cash in hand, usually triggers that—were other aspects of AI systems that participants discussed. Whether because of regulatory compliance or the jobs AI systems will engender, the utility and expense of deploying these systems will be determinative to some extent. In the EU, for instance, money will be made on the law side of things offering such services because of the risk regulation model and the apparatus that will be built around complying with such a complex procedural regime.

Another example and issue come from Intellectual Property, more specifically from Copyright law. A lot of money, which some of the AI companies have, could make some of the copyright problems go away, but there are big doctrinal issues before us. Whatever happens in this space will influence and be influenced by copyright law. Will the money be split up between owners of AI systems and copyright holders after the fact, for instance? These will continue to be significant issues, but they are already being litigated, in *Gonzalez v. Google* and *Twitter v. Taamneh*, implicating §230 and algorithms already. So, we might have some guidance soon from the Supreme Court, especially if Congress doesn't legislate in this area. However, the Court's and Congress' love for big tech is increasingly diminishing and that might lead to some interesting and possibly unexpected outcomes, as well. Of course, at the state level, we are already seeing new laws emerging vis-à-vis AI, so it will be worthwhile to watch how that unfolds and affects this area.

Responding to an earlier remark, one participant claimed that the bottom-up approach has been thrown out the window. Early on, there was a bottom-up movement, where experts came together and reached a consensus on trustworthy AI, responsible innovation, and fair and balanced data training. AI applications were supposed to be transparent, and explainable, but then OpenAI destroyed that. Everybody else is now playing catch up, including Google, whose motto was previously "don't be evil." This suggests that Congress needs to act, because left to their own devices, AI companies might not act very responsibly. AI systems are based on large language models (LLMs), and they are good for wordsmithing and grammar and do a good job at online chat or auto complete.

The problem is they don't disclose what data they train these models on and yet claim AI does natural language understanding.

The claims that AI does not do natural language understanding and that there is a lack of transparency as to what it is trained on instigated a lively discussion among the participants. Some participants said that a large percentage of the Internet was used for training (95%), with the latest

iteration of GPT-4, allegedly, being trained on two million books that OpenAI paid for. However, that only amounts to the creation of probabilities about what words are most likely to come next, not actually knowing things. Even without disclosing the whole provenance of something GPT-4 generated, AI companies should at least tell us what they trained their systems on.

Some participants argued that part of the reason why AI companies don't disclose the training data is because they originally did not intend to commercialize the AI systems they were building, that they might have broken a lot of copyright laws in the process. This might explain the reticence of AI companies in being transparent.

Although almost all AI companies have ethics guidelines, one expert extended, it doesn't appear that they follow those guidelines because they keep firing their ethics team without a concern for how bad it looks, the perception of it, or what such behavior implies. Algorithmic audits and impact assessments, the expert suggested, are more likely to become part of the regulatory scene in the future. However, they added, it will be after the harm has already been done. It is important that we think of ways to mitigate those risks. Requiring AI ethics officers for companies using AI applications commercially, similar to privacy officers, could be one way of thinking about this issue.

Other participants, in response to the claim that AI systems don't do natural language understanding, stated that GPT-4 is a significant improvement on ChatGPT, and it shows all the qualities of understanding natural language.

Regarding regulatory challenges, one participant stressed the importance of implementing dynamic and adoptive strategies due to the difficulty of predicting all potential issues that might arise, even if such an approach might not be able to stop all harm. They gave Microsoft's approach as an example, which monitors output of their systems, flags problematic content produced, and thus enables an after-the-fact analysis and improving on that iterative process. Audits of the reinforcement learning stage would better this process further. An iterative process rather than a top-down regulation might be the best way forward in achieving positive outcomes of AI and minimizing negative consequences.

Another noteworthy perspective offered by a participant highlighted the fact that words, such as fairness, harm, or risk, have different meanings and interpretations depending on one's perspective. Legally they might mean one thing, and in computer science they might mean something else. As a solution to that issue, establishing a common shared ontology might prove much more beneficial than a merely top-down or bottom-up approach to regulation. However, another participant expressed skepticism about a shared ontology around

ethics, at least from a legal perspective, as it can be used to justify all kinds of actions without any real guardrails or oversight.

Regulations, at times, could be very divorced from the realities of an industry. Addressing the firing of ethics teams at Microsoft, which was acknowledged as bad optics, one speaker told a joke from Stalin's Soviet Russia: a biologist and a bioethicist are asked about their last wishes before being executed, the bioethicist asks to give one last lecture, while the biologist asks to be shot before the lecture. Sometimes there is animosity between those working in the field and those wanting to regulate them. This divide or tension could have unforeseen consequences and challenges regarding regulations and their impact in carrying innovations forward.

One speaker provocatively suggested considering a ban similar to the Illinois Biometric Privacy Act, which prohibits certain uses of biometric data altogether, at least for public uses and until the benefits and harms of AI systems are more clearly discernable and more thoroughly evaluated.

Lastly, an ethicist participant emphasized that we should not ignore ethical issues—the risks of real harm to real human beings—in the name of progress, which often justifies moving fast and breaking things without paying necessary heed to the degree of real adverse effects these systems give rise to.

About Silicon Flatirons Center

What We Do

Silicon Flatirons' mission is to elevate the debate surrounding technology policy issues; support and enable entrepreneurship in the technology community; and inspire, prepare, and place students in these important areas.

Founded in 1999 at the University of Colorado Law School by Phil Weiser, Silicon Flatirons is a recognized leader in interdisciplinary events and programs. We serve students, entrepreneurs, policymakers, and professionals, and support the joint missions of Colorado Law on teaching, scholarship, and public service.

Although technology and innovation have evolved radically over two decades, our purpose remains the same: to convene multi-stakeholder discussions, support innovation, and develop the next generation of technology lawyers, policy experts, and entrepreneurs.

Our initiatives are hubs for pivotal issue areas at the intersection of technology and law, led by internationally recognized experts who facilitate programming and convenings and generate thought-provoking scholarship within those areas.

Learn more at siliconflatirons.org/about-us/.

Our Team

For more information about center leadership, faculty, staff, fellows, and advisory board, visit siliconflatirons.org/about-us/our-team/.

Our Supporters

Silicon Flatirons exists thanks to the generosity of our supporters and the strength of our community. We rely on their contributions to advance our mission to catalyze policymaking and innovation and to develop the next generation of tech lawyers, policy experts, and entrepreneurs. For more information on current official Silicon Flatirons Supporters, visit siliconflatirons.org/about-us/supporters/.

Publications

We promote thought leadership and intellectually honest discourse not only in our events, but in publications from our team, our roundtables, and scholars presenting at our conferences. See more at siliconflatirons.org/publications/.