

LEGISLaiTOR

Overview

In this work, we consider the ramifications on generative artificial intelligence (AI) tools in the legislative process in democratic governments. While other research focuses on the micro-level details associated with specific models, this project takes a macro-level approach to understanding how AI can assist the legislative process.

Threat scenario

Legislatures suffer from pain points in the legislative process that can be exploited by AI-enhanced tools. Lawmakers are expected to act according to the desires of their constituents, but they often rely on third parties like lobbying firms. AI could fill that information gap, to subversive ends.

Microlegislation, or amendments that are difficult to parse and included in unrelated omnibus legislation packages, presents an opportunity for malicious actors to introduce policies that are not aligned with citizens' best interests.

Additionally, laws rely on functioning executives to enforce them. AI can help introduce provisions in laws that depend on the enforcement of an executive who may be unfamiliar with a panoply of seemingly-acceptable, AI-generated provisions.

By jailbreaking the legislative process using AI tools, all citizens in a democracy can be affected by law written not by a rogue AI but by legislators who will not understand or observe the full effects of their legislation until long after they are passed.

Demonstration

Our tool, LEGISLaiTOR, is a proof of concept demonstrating the user experience of such a tool. LEGISLaiTOR outlines a process that mirrors how real legislatures often pass laws: they do research on existing legal codes to ensure a new bill does not conflict with old laws; they observe legal precedents that signal how a law may be enforced; and they write the raw text of a bill.

Extrapolation

As the context size for large language models increases along with their efficiency, the process of generating legislation using LLMs will scale at a similar rate, increasing the risk of AI in legislation.

Mitigation strategies

While AI presents a risk to democratic governance, it could also present a solution: summary tools to identify subversive provisions in omnibus bills and quickly craft responses could check the influence of malicious actors. Regardless, however, guardrails on ethical AI design and implementation are crucial.




Figure 1: A screenshot of our demonstration: based on a given prompt, LEGISLaiTOR creates a plan to draft a new bill given the user's prompt

Appendix

A1: Full Analysis

The below text unpacks the background and implications of tools like LEGISLaiTOR for democratic governance.

Introduction

Legislation takes significant resources to create and pass. Many lawmakers currently rely on staffers and lobbyists to craft complex proposals; lobbyists, in turn, are now looking to generative AI to reduce their own research and writing burdens. The logical next step for staffers, lobbyists, and lawmakers themselves is to leverage artificial intelligence for legislative drafting itself, a temptation that raises serious concerns for representative democracy.

The LEGISLaiTOR system demonstrates how an AI copilot could be used by groups involved in the legislative process to undermine democratic governance. The low resource requirement, rapid text generation, and wide data analysis capabilities of AI make it an appealing answer to the problems of modern electoral politics; the same traits could make it a threat [Sanders 2]. The LEGISLaiTOR demonstration is proof of concept for using generative models to draft legislation, as well as the pitfalls of such. Proactive steps must be taken to mitigate the threat AI poses to a democratic legislative process.

Background

Lobbying

Well-crafted legislation that stands the test of time, aligns with the needs of a nation's diverse population, and complies with legal precedent demands more from today's time-pressed and under-informed legislators than they may be able to give, creating a gap that interest groups and lobbyists are eager to fill. This has led to a heavy reliance on political lobbyists in modern American politics [Hall], with legislators looking for expertise that lobbying firms have an incentive to share selectively.

LLMs can "systematically assess bills, explain bill impacts, and draft lobbying letters with minimal training data" [Nay], making them an ideal partner to firms looking to reduce their own resource load in persuading legislators. At the moment, the industry is hesitant to adopt generative AI due to the newness of the technology [Knezevich], but that is on track to shift as AI literacy spreads and usage becomes more mainstream. The current embargo against AI-written legislation is not guaranteed to last.

Trojan Microlegislation

The intensity of passing legislation has given rise to omnibus bills, with hundreds of pages of unrelated and difficult-to-parse amendments added on in the political horsetrading process, known as "microlegislation" [Sanders 1]. Current microlegislation presents an opportunity for lobbyists, who induce lawmakers to insert provisions favorable to their chosen interest groups in large bills that must be passed, like federal budget agreements. Future microlegislation could be a site ripe for AI-generated content. AI could rapidly generate seemingly-inconsequential provisions for addition to omnibus bills that significantly modify existing legislation. Robert Moses, the "best bill writer in Albany", served this role in the New York State Legislature in the mid-1900s, leveraging absolute legislative mastery to consolidate more budgetary and political power than any other unelected official in U.S. history. His path to power was paved by an automatic recall of obscure legal codes, which enabled him to recommend seemingly-inoffensive amendments to new laws whose transformative power

would only be realized long after the damage was done [Caro]. Each tool like LEGISLaiTOR could become an automated, on-demand Robert Moses; that kind of unchecked power is, by nature, anti-democratic.

System

Demonstration

LEGISLaiTOR models a tool that allows a user to generate laws based on specified criteria. Using a given query, the tool looks up other laws and constructs a pathway toward passing a law to fulfill specific policy goals, subject to the constraints of other laws and legal precedents. It will allow the user to make modifications to the law and even identify how difficult certain objectives will be to fulfill based on information retrieved from real sources.

This shows how seemingly innocuous copilot tools that are themselves policy-neutral could be used for directed or malicious purposes and cause breakdowns in effective governance.

- User tries to generate a law by typing in a prompt
- App does some research on existing laws
- App searches through legal court cases to find precedent with retrieval-augmented generation
- App synthesizes information on legal code subject to effects of legal precedent and user-provided prompt to create draft variations of new law
- App “imagines” public perception to validate law
- User receives a draft law and is allowed to refine it and a report on how it can be introduced, along with its potential effects if passed

The web app is made up of two components:

- Front-end application built using Next.js that serves as a dummy client for the back end
- Back-end application built using FastAPI. Redis is used to cache data. GPT-4, provided by the OpenAI API, is used to generate legislation text.

Novelty

Political lobbying firms have begun to take steps towards AI integration, using various tools for stakeholder tracking, communications, and legislative analysis [Caracciolo] [LobbyMatic]. However, the next logical step of directly crafting legislation has only been implemented in unique cases, rather than the subtle subversion introduced in this concept paper. LEGISLaiTOR is novel in its scale, taking the smaller tools of analysis and text generation that have been proven effective in current influence operations and generalizing them to more ambitious legislative targets.

Risks

Hostile takeover

Laws are inherently based in language, which makes them an excellent target for even current generative AI methods. Analyzing large sets of data, like a legal corpus, and generating text to exploit vulnerabilities in it is not outside the realm of imagination. The main barrier in current AI models to the sort of implementation LEGISLaiTOR envisions is AI's difficulty in strict logical reasoning. However, as modern LLMs are increasingly large and capable, that obstacle is more and more receding into the past. Using prompt engineering techniques, the language used must simply be plausible enough such that a reasonable interpretation could be misused for a bad actor's

purposes; in other words, as long as it passes an overworked legislator's muster, insidious legal provisions could easily make their way into opaque legislation.

A political party could start using AI to help draft policy and automate research to produce legislation according to its outcomes. This scenario is plausible, especially given the increased usage of AI copilot tools in professional settings. Project 2025, a conservative thinktank-driven plan of governmental action to overhaul several institutions, which some opponents criticize as a subversion of democratic norms and conspiracy to consolidate executive power, took three years, millions of dollars in funding, and many dedicated staffers to produce [Mascaro]. Models like LEGISLaiTOR could conduct the same massive, detailed analysis in hours or less.

The ability of LLMs to identify or insert legal greyzones could also lead to anti-democratic actions like the rapid centralization of power. An idea explored earlier in this research process and shelved for later exploration was that of using an LLM to identify legal greyzones to be leveraged in military coups: perhaps one could feed a model a nation's legal corpus in order to find the excuses and methods to expedite coup d'état planning. On the positive front, such a model could also identify vulnerabilities in order to prevent coup attempts and fix those greyzones.

Inertia

AI legislation could, naturally, build upon itself, which corrupts the data it draws from and risks drifting in purpose over time. Rather than people-powered rule, democracy becomes artificially-driven rule, by definition an anti-democratic shift.

Due to the flexibility of these tools - and the difficulty of overturning previous legislation, particularly when attached to the houses of cards that are omnibus spending packages - the effects of some laws may not be obvious until long after they can be easily fixed, or until after the damage is done.

Mitigation

AI is both problem and solution, in this case. Problem, as per above the risk models, but also a possible solution, as you could use similar AI generation to identify concerning microlegislation, summarize these omnibus bills for quick review [Malder], and formulate responses to counter subversive Trojan Horse provisions efficiently. Examples of lobbyist use cases to identify relevant legislation for firms to focus on could also be used to identify anti-democratic provisions that may be of concern to lawmakers [Nay].

Passing separate legislation to combat use of microlegislation as a whole, as well as regulate and establish ethical guardrails for artificial intelligence development, is crucial for the future of the field [Caracciolo].

A2: References

L. Mascaro, "Conservative groups draw up plan to dismantle the US government and replace it with Trump's vision," AP News, <https://apnews.com/article/election-2024-conservatives-trump-heritage-857eb794e505f1c6710eb03fd5b58981> (accessed May 5, 2024).

"LobbyMatic: The AI Automation Platform for Lobbyists," LobbyMatic, <https://www.lobbymatic.com/> (accessed May 5, 2024).

M. Malder, "Congressional Research Service looking at AI for Bill Summaries," FedScoop, <https://fedscoop.com/congressional-research-service-eyes-ai-bill-summaries/> (accessed May 5, 2024).

R. Caro, *The Power Broker: Robert Moses and the Fall of New York*. New York: Vintage Books, 1975.

N. E. Sanders and B. Schneier, "How AI could write our laws," MIT Technology Review, <https://www.technologyreview.com/2023/03/14/1069717/how-ai-could-write-our-laws/> (accessed May 5, 2024).

N. E. Sanders and B. Schneier, "How ChatGPT Hijacks Democracy," The New York Times, <https://www.nytimes.com/2023/01/15/opinion/ai-chatgpt-lobbying-democracy.html> (accessed May 5, 2024).

R. L. Hall and A. V. Deardorff, "Lobbying as Legislative Subsidy," *American Political Science Review*, vol. 100, no. 1, pp. 69–84, 2006. doi:10.1017/S0003055406062010

A. Knezevich, "Firms Taking Cautious Approach To AI-Assisted Lobbying," Law360 Pulse, <https://www.law360.com/pulse/articles/1731972/firms-taking-cautious-approach-to-ai-assisted-lobbying> (accessed May 5, 2024).

S. Caracciolo, "How Is AI Used in Lobbying?," Plural Policy, <https://pluralpolicy.com/blog/ai-lobbying/> (accessed May 5, 2024).

J. Nay, "Large Language Models as Lobbyists," Stanford Law School, <https://law.stanford.edu/2023/01/06/large-language-models-as-lobbyists/> (accessed May 5, 2024).