DOGE Plan to Push AI Across the US Federal Government is Wildly Dangerous

BEN GREEN / MAR 6, 2025

February 20, 2025—Elon Musk speaking at the 2025 Conservative Political Action Conference (CPAC) at the Gaylord National Resort & Convention Center in National Harbor, Maryland. Gage Skidmore via Wikimedia CC BY-SA 2.0

Last month, Elon Musk's DOGE team outlined a plan to transform the United States federal government through an "AI-first strategy."

Led by Thomas Shedd, a former Tesla engineer now in charge of the General Services Administration's technology team, the plan is to deploy AI widely across the federal government. The overall goal is to cut the agency's budget by fifty percent. Shedd suggested using AI to analyze contracts for redundancies, root out fraud, and facilitate a reduction in the federal workforce by automating much of their work.

It's no surprise that Musk, who co-founded OpenAI in 2015 and now (among his various endeavors) runs the company xAI, is advocating for an AI-driven government transformation. This ambition may seem plausible since AI tools are capable of many impressive feats, from writing essays to answering physics questions. But unleashing AI to revolutionize the federal government is a recipe for disaster.

To predict how Musk's proposed transformation will go, we can look at prior examples of integrating AI into government decision-making and operations. There are already many instances of AI-driven governance gone wrong—particularly in cases where the goal is to root out fraud and cut budgets.

In my home state of Michigan, the Unemployment Insurance Agency (UIA) adopted an algorithm (the Michigan Integrated Data Automated System, aka MiDAS) to streamline its operations. The UIA's goals were to prevent unemployment fraud and to eliminate one-third of the existing staff by automating their work. MiDAS quickly boosted the cases of suspected welfare fraud by a factor of five, which led to a 23-fold increase in the UIA's revenues.

Based on these numbers, MiDAS seems to have accomplished the UIA's goals. The only problem? Almost every accusation of fraud—93% of supposed fraud cases—was incorrect. Then, even once these errors became clear, it took years of litigation for the wrongly accused to receive the money they were owed. For many people, the fraud charge remained on their criminal record for years, barring them from jobs.

Around the United States and the rest of the world, similar algorithms have cut vulnerable people off healthcare and violated human rights.

A central reason for these flawed AI tools is that engineers underestimate the complexity of government processes. In turn, engineers embed their superficial assumptions into software and overestimate their tools' capabilities.

It is a typical story that engineers walk into government departments with swashbuckling confidence, believing that technical expertise can break through existing bottlenecks and solve all the government's problems. These quick-fix proposals are typically laughable to the existing staff in those agencies, who have deep expertise in navigating the complex nuances of government processes. The engineers do not understand the intricate workflows that government agencies follow. And because they do not understand these workflows, the engineers do not know how to interpret government data and forms.

In a recent study, colleagues and I asked computer scientists to develop a software tool that gives automated advice about whether people are eligible for bankruptcy relief. The computer scientists completed the task quickly, with one even noting, "this seems very simple." However, when we analyzed the software tools they built, we found that errors and legal misconceptions were rampant. Despite these flaws, the computer scientists were confident that their tools would be beneficial in

practice. Seventy-nine percent of them stated that they would be comfortable with their tool replacing judges in bankruptcy court.

The DOGE team is displaying this engineering hubris to an extreme degree. They are racing forward with little care for existing laws and protocols. Without spending the time to understand government operations, Shedd has asserted that AI can replace many federal workers.

What will happen if these plans for an AI revolution in government move forward? Already, we are seeing evidence that DOGE is repeating the same mistakes that have plagued other instances of government automation.

Most notably, DOGE will likely make many baseless judgments of waste and fraud based on AI results. This outcome will be partly due to Musk's own wildly over-expansive definitions of waste and fraud. He has stated, without any evidence and contrary to other estimates, that "A trillion dollars can be saved just by addressing waste, fraud, and abuse." Given this viewpoint, DOGE's AI developers will surely be more worried about avoiding false negatives (overlooking an instance of fraud) than false positives (incorrectly labeling a contract as wasteful or fraudulent), particularly in cases where they may have a conflict of interest.

One cause of mistaken accusations will be misinterpreting information in government databases. Government data is notoriously complex. Understanding its meaning requires detailed knowledge of how that data is collected and used. An early instance of this error came when Musk tweeted in outrage about millions of people older than 100 supposedly receiving Social Security checks. His post included a screenshot of a simple analysis of the Social Security database, counting the number of people marked "alive" at different ages. Although his analysis correctly counted the data, Musk misinterpreted what it meant. Millions of individuals over 100 are marked as alive due to technical limitations and gaps in the data. Crucially, almost none of them are receiving checks. In fact, several years ago, the Social Security Administration considered updating its systems for annotating deaths but decided that the cost of the upgrade would not be worth the benefits. Had Musk simply asked

Social Security experts about the data, he could have gained a correct understanding. Instead, he jumped confidently to an incorrect conclusion.

Another likely cause of mistaken fraud accusations is AI's inability to interpret the meaning of contracts and the relationships between them. Governments often divide work into underline{multiple contracts}, each serving different aspects of larger programs. These contracts may look redundant to an AI tool when, in fact, they are distinct and complementary. To be sure, there can be benefits to streamlining government contracts. But, that requires carefully reconstructing contracts to serve the underlying policy goals, not sending existing contracts to the woodchipper.

DOGE has already demonstrated a pattern of misinterpreting contracts, leading it to vastly overstate the savings it has generated. On February 19, DOGE posted a "wall of receipts," boasting $16.5 billion in savings from canceled contracts. Their calculations were underline{riddled with mistakes}. For instance, many "canceled" contracts had ended years ago or were terminated under prior administrations. Since their initial post, DOGE has updated 40% of the contracts with more accurate—and smaller—values. These updates included deleting the five largest contracts for which DOGE had initially claimed credit. That $16.5 billion now underline{looks more like} $8 billion, and additional errors remain uncorrected.

To make matters worse, DOGE is making these decisions in an opaque and unaccountable manner. It is unclear what role AI has played in the actions that DOGE has already taken. Moving forward, an AI's classification of fraud could lead to contracts being canceled immediately without any notice or appeals process. When this action harms a person or organization, they may not be able to get a clear explanation of what happened aside from "the AI said so."

Of course, for Musk and his allies, these problems may be features rather than bugs. The AI will help them rapidly cut the federal budget and assert their dominance over the government. So what if many fraud accusations are incorrect and some essential government operations get waylaid?

None of this is to say that AI cannot be a helpful tool for improving government services and operations. But, achieving that goal requires taking the opposite approach of DOGE.

During my own time working as a data scientist in local government, I quickly learned that my technical know-how was wildly insufficient on its own. To advance socially beneficial projects, I had to recognize the limits of technology and avoid the assumption that the most cutting-edge technical tools would be the most useful in practice. I also had to appreciate the wisdom of existing staff. I relied on their guidance to make sense of government processes and datasets. I needed to view my work as augmenting the work of existing teams, not replacing them.

The most unmistakable evidence that Musk and DOGE are not committed to thoughtfully developing technology that improves government is that they are dismantling the teams that have been pioneers in this work. Since it was founded in 2014, the United States Digital Service (USDS) has underline{generated}

benefits such as improving healthcare access for 18 million veterans and enabling millions of rural households to access affordable high-speed internet. However, DOGE quickly fired a third of USDS staff. A few days later, another group of USDS staff resigned, writing that "DOGE's actions … are not compatible with the mission we joined the United States Digital Service to carry out: to deliver better services to the American people through technology and design." Then, last weekend, DOGE eliminated 18F, a group that helps federal agencies improve their technology. This move came just weeks after Shedd told 18F staff that they represented the "gold standard" of government technology developers.

Given the frantic pace and hubris of Musk and the DOGE team, I fear that this AI incursion will do incredible damage to federal operations and the people who depend on them. But, at the very least, those of us who teach courses about technology and public policy will have a new case study of government AI gone wrong to share with our students.

AUTHORS

BEN GREEN

Ben Green is an assistant professor in the University of Michigan School of Information and an assistant professor (by courtesy) in the Gerald R. Ford School of Public Policy. His first book, The Smart Enough City: Putting Technology in Its Place to Reclaim Our Urban Future, was published in 2019 by…

RELATED

Anatomy of an AI Coup

How Congress Can Delete DOGE

At DOGE, Musk's Billionaire Interests Likely to Trump Public Interest

DOGE Understands Something the US Policy Establishment Does Not: Technology is the Spinal Cord of Government

A National Heist? Evaluating Elon Musk's March Through Washington

TOPICS

Artificial Intelligence   Democracy

A nonprofit media and community venture intended to provoke new ideas, debate and discussion at the intersection of technology and democracy.

About	Fellows	Articles
Donate	Contributors	Podcast
Privacy Policy	Submissions	Research Library