

I hacked the government, and it was easier than you may think

February 7 2020



Max Weiss '20 conducted research on bots and hacked a government website comment section with 1,000 fake comments. Credit: Stephanie Mitchell/Harvard Staff Photographer

Max Weiss never intended to hack the government. His discovery of

how easy it is to do—outlined in a new paper he authored—came of the best of intentions.

Weiss, a government concentrator from Cincinnati, was doing advocacy work for state expansion and defense of Medicaid last summer, a project that combined his interests in public policy and healthcare. While studying the ways in which various advocacy groups can influence pending legislation, he learned how valuable such groups find the federal government's comment period, when members of the public are invited to weigh in on new or pending legislation via online forms. He realized how easy it would be to manipulate the results using bots—computer programs that generate automated responses—to flood the sites with fake responses for or against any proposal.

The 21-year-old detailed his findings in a recent *Technology Science* piece, "Deepfake Bot Submissions to Federal Public Comment Websites Cannot Be Distinguished from Human Submissions."

"We were spending a lot of time and energy getting high-quality comments from constituents," said Weiss. "I wanted to make sure these [federal agencies](#) understood the potential consequences of their policies, and I had the idea that I could use a bot and submit a lot of fake comments."

He paused, recognizing that corrupting the process was fraught: "This would be bad for democracy."

But the Leverett House resident couldn't shake the idea, and he began to research the feasibility of such a scheme. Turns out submission is easy to automate. Federal agencies have some leeway to discount comments that are obviously duplicated or irrelevant. But the typical technological defenses against attack, including CAPTCHAS, anomaly detection, and outside verification—all of which are integrated into online activity from

banking to email log-in—were pretty much absent.

Options for Search-and-Replace Comment-Building Method							
1.	2.	3.	4.	5.	6.	7.	8.
I strongly	demand	the FCC	to undo	Obama's	order	to take over	the Internet.
I want to	ask	you	to repeal	Tom Wheeler's	plan	to control	broadband.
I'd like to	urge	Ajit Pai	to reverse	Barack Obama's	policy	to regulate	the web.

Examples of Search-and-Replace Comment Sentence Combinations
I strongly ask you to undo Tom Wheeler's plan to take over the web.
I want to ask Ajit Pai to undo Barack Obama's plan to regulate the web.
I strongly urge you to undo Barack Obama's policy to take over the Internet.
I'd like to demand the FCC to repeal Obama's order to control broadband.
I'd like to urge the FCC to undo Tom Wheeler's policy to regulate the Internet.

Figure 1. Example of Synonym Replacement Used to Build Sentences in Large FCC Public Comment Campaign. The figure shows five examples of sentences (bottom panel) built from eight sentence components, each with three near-term options (top panel). The near-term options used to build combinations for one sentence were taken directly from just one FCC commenting campaign (comprising 1.3 million comments) discovered and dissected by Jeff Kao [1]. Given only the near-term options shown, $38 = 6,561$ variations of this same sentence could be created.

"Most of those websites really just have a text box for your public comments and then a submit button," he said.

In the course of writing the *Tech Science* paper, Weiss realized that cybersecurity experts have been sounding the alarm on federal website vulnerability for years, but previous transgressions had used relatively unsophisticated substitution methods. "In 2017, there were 22 million comments posted for the FCC proposal to repeal net neutrality," he

recalled. "And it was found that 96 percent of those were part of duplicative campaigns."

Weiss used AI methods to generate a high volume of unique deepfake comments about a proposed Medicaid waiver. He then wrote a program that automated the submission process, and ran it from a laptop in his dorm room over the course of a few days. He submitted more than 1,000 fake comments that comprised 55 percent of the total submissions and that were found by survey respondents to be indistinguishable from human comments. Afterward, he notified the federal Centers for Medicare and Medicaid Services which comments were part of his demonstration to prevent their interference with authentic public comment evaluation.

Among the scarier revelations was Weiss' admission that he was successful without being an expert coder and without special equipment. "I've learned to code in the last four years, just through a series of personal projects and summer jobs, and one class," said Weiss, who has taken some courses in the new program in technology science. "I think one of the very important findings from the study is that someone like me who's a very novice coder was able to Google his way through hacking the government.

"I've always been very interested in [public policy](#)," said Weiss, who also enjoys writing and performing comedy. "Most of my government study has been in health policy or in technology policy or public interest technology, so this was just kind of a synthesis of a lot of different things that I'd learned in the Government Department and just some personal tech projects that I have done in the past."

"Max did groundbreaking work, exactly the kind of real-world-impact work we encourage our students to do" in technology science classes, said Latanya Sweeney, professor of government and technology in

residence and director of the Data Privacy Lab at the Institute for Quantitative Social Science, who serves as editor-in-chief of Technology Science.

"Thanks to Max's work, several groups within the federal government are now actively making changes to combat these kinds of vulnerabilities," she added.

More information: Deepfake Bot Submissions to Federal Public Comment Websites Cannot Be Distinguished from Human Submissions, Technology Science. doi.org/10.7910/DVN/OQCPOT.
techscience.org/a/2019121801/

This story is published courtesy of the [Harvard Gazette](#), Harvard University's official newspaper. For additional university news, visit [Harvard.edu](#).

Provided by Harvard Gazette

Citation: I hacked the government, and it was easier than you may think (2020, February 7) retrieved 27 April 2024 from <https://techxplore.com/news/2020-02-hacked-easier.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.