AI Insight Forum: Elections & Democracy

Statement of Lawrence D. Norden
Senior Director, Elections and Government,
Brennan Center for Justice at NYU School of Law
November 8, 2023

Majority Leader Schumer, Senators Rounds, Heinrich and Young:

Thank you for the opportunity to discuss the critical question of how Artificial Intelligence (AI) will reshape the landscape of American electoral processes and democratic governance. For over a year, the Brennan Center has explored this question, partnering with technical experts and election officials to understand the ways in which AI may impact American democracy, and to identify steps that government, private and nonprofit sectors should take to minimize the dangers and increase the benefits of this powerful new technology. In my statement below, I look at the ways in which AI is likely to intersect with our democracy in five important arenas: the information environment for elections; election administration; election security; voter suppression; and participatory democracy. For each of these topics, I also offer suggestions for how Congress and federal agencies can maximize the benefit and reduce the harm of AI's extraordinary power.

AI and the Information Environment for Elections

New developments in AI are poised to disrupt and damage the information landscape surrounding elections in myriad ways. While machine-manipulated media has existed for years, malefactors can now create AI-generated deepfakes and other manipulated images and audio more cheaply, easily, and swiftly than before. Duplicitous actors have already used generative AI to spread scandalous falsehoods about a Chicago mayoral candidate on the eve of an election and to possibly skew a national election in Slovakia. The use of deceptive AI-generated content by campaigns, PACs, and paid advertisers also poses a significant concern. From fakeries showing former President Trump hugging Anthony Fauci, to false depictions of a dystopic future under President Biden, to fabrications of candidate voices, generative AI has already made appearances in the 2024 U.S. presidential campaign — and its disruptive potential is only expected to grow.

While much of Congress' focus to date has been on visual and audio AI-generated content in elections, large language models have potential to be similarly destructive. Bad faith actors could deploy the technology underlying chatbots to spread falsehoods online or via robocalls faster and on a much larger scale. They could spoof election official websites and produce reams of content to game platform algorithms so that false narratives trend on social media. Poynter has demonstrated how antagonists can use AI to generate whole misinformation-filled fake news websites in a matter of minutes— a special risk for elections given that millions of eligible U.S. voters live in counties with no enduring local newspaper.

Advances in AI also have the potential to reduce trust in elections and authoritative sources of information. Malign foreign influence campaigns such as Russia's have long sought to muddy the information landscape so that U.S. voters find it more difficult to discern between truth and falsity. A related risk is the phenomenon of the "liar's dividend," whereby politicians and other actors can deny true events by claiming that documentary evidence is AI-generated; this is an outcome playing out prominently in the current conflict in Israel and Gaza, and one that we should expect to similarly manifest in future elections.

Among the actions Congress can take to mitigate the risk of deception and distrust in elections it should:

- mandate that campaigns, PACs, and other paid political advertisers label AI-generated images, video, and audio depicting candidates and events in substantially misleading ways;
- require campaigns, PACs, and paid advertisers to label a subset of content produced by large language models as generated by AI — for instance, if campaigns deploy interactive chatbots to engage with voters;
- modernize transparency rules for online political ads generally and require information about the role AI played in creating an ad to be disclosed where appropriate, such as in the public file requirement of the bipartisan Honest Ads Act, S.486;
- explore mandatory watermarking of AI-generated content by AI developers, and digital signatures for verified content that could then be labeled by social media, messaging and other digital communications platforms;
- fund the development of new technologies, tools and standards for establishing digital content provenance tracking systems and authentication of verified content.

AI and Election Administration

As AI tools become cheaper and more widely available, government agencies and private companies are rapidly deploying them to perform basic functions and increase productivity. There is every reason to believe we will soon see the same phenomenon in under-resourced and understaffed election offices as well. Indeed, for several years election offices around the country have used AI to perform important, but limited, functions more effectively. That includes using systems that rely on machine learning for tasks as diverse as filtering spam from election worker email, data list management, and preliminary signature matching for mail ballots. Election officials have also begun to turn to AI chatbots to answer basic voter questions, and to AI products to help them make election related material more understandable.

Looking forward, it is easy to imagine additional future uses of AI by election offices, including to assist in identifying new polling place locations through the analysis of geospatial data. AI could be harnessed for such tasks as adjudicating ballots, generating translated materials, analyzing post-election data in order to improve future elections, identifying trends in provisional voting, completing voter registration, and screening for reasons to reject absentee and vote-by-mail ballots. No doubt, as the technology continues to evolve and mature, an array of additional opportunities for AI to aid in the administration of elections will emerge.

While wider adoption of AI in election administration offers promise, it also comes with significant risk. Over the last several years, we have witnessed numerous examples of AI products failing, resulting in amplification of biases, "hallucinations" of false information, and other mistakes that were not caught by human supervisors of the AI systems. Public confidence in the American electoral system is already dangerously low, thanks in large measure to false claims made after the 2020 election. Failures like these in critical election administration functions could cause further, long-term damage to the public's faith in elections.

As AI begins to be used more broadly in elections, it is critical that federal and state governments develop standards, certification, and monitoring regimes for its use in election offices, as should vendors selling products and services to election offices. President Biden's Executive Order on the "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," (the "Executive Order") marks an essential first step in creating standards for the use of AI in elections. But there is more that needs to be done to ensure that AI can be used safely, securely, and transparently. Among the most critical actions are:

- Creating a standards and certification process to allow election officials the ability to identify AI systems that meet baseline federal standards for AI in elections. In addition, these new standards must be complimented by an auditing and monitoring regime once AI is actively deployed. This will allow election offices and vendors to make adjustments to their use of AI and ensure that the mitigations to counter risk are as strong as possible.

- Providing federal funds for state and local election offices to implement the relevant AI guidelines and standards for election projects, as well as for red-teaming, auditing and monitoring of AI systems that will be necessary going forward.
- Providing Congressional funding for the creation of a NIST AI Risk Management Framework profile for AI used in election processes. This builds on the work done by NIST to create an election profile in the Cybersecurity Framework. These profiles help election officials better understand how to apply general recommendations to their election specific circumstances.
- Mandating election security best practices for vendors in the elections space, as Congress has done for vendors in other critical infrastructure sectors. These standards should include guidelines for the use and disclosure of AI in election vendors' work to identify potential security risks.

AI and Election Security

AI poses several threats to the security of the election infrastructure in our election offices and election system vendors. Among the most serious categories of threats is the ability of adversaries to convincingly imitate authoritative voices, whether that is as the fake supervisor at an election system vendor asking an election worker for a sensitive password via phone or email message, or a sham election official making false claims about election processes and results.

AI will also change how software is engineered. It can make cyberattacks bigger, quicker, and sneakier to outsmart existing software security measures. These advantages may become available to a wider set of adversaries, including individual domestic antagonists and nation-states like China, Russia, or Iran, all of whom have meddled in recent American elections and are developing their own AI technologies capable of targeting American networks. Finally, AI can be used to overtax our election infrastructure. As Ron Rivest of the Massachusetts Institute of Technology (MIT) has observed, with AI "any adversary can produce more high-quality output with less effort than before." That could result in more effective denial of service attacks against key on-line systems like voter registration portals. It could also be used to generate massive numbers of requests of election offices, through open records requests or phone calls into election offices, each seeming as if it came from a separate person, but actually as an intentional effort to prevent election officials from doing their jobs.

There is no silver bullet for the additional security risks wrought by AI's rapid advances and increased availability. Broadly, however, there are four areas of potential mitigation against these heightened risks, and in each case a role for the federal government to play: building more secure and resilient election systems; providing election officials with more technical support to safeguard election infrastructure; giving the public ways to authenticate election office communications and detect fakes; and offering election workers AI-specific trainings and resources.

Among the most critical actions Congress and federal agencies can take are:

- Increasing investment in election security. This should include not only providing enough funds to ensure states can build more resiliency into our elections, but also invest in tools to detect and defeat AI generated cyberattacks (which may include new defensive capabilities supported by AI) and the infrastructure that will be necessary to distribute these new tools nationwide;
- Providing more technical support to state and local election offices through investment in state cyber navigator programs and cyber security advisors at the Cybersecurity and Infrastructure Security Agency (CISA);

- Giving election officials more tools to authenticate their communications with the public. This could include having CISA pilot authentication tools with election offices for official election related documents. It could also include having CISA or the Election Assistance Commission (EAC) create government verified handles for election office social media accounts, similar to something Germany has done for government offices there.

AI and Voter Suppression

New developments in AI are primed to exacerbate voter suppression in American elections. Malign actors have long targeted voters —particularly Black Americans and Latino Americans— with deceptions about how, when, and where to vote in attempts to trick groups out of voting. Bad-faith and misguided individuals have also exploited voter challenges and voter purges to attempt to wrongfully toss eligible citizens off the voter rolls.

AI could supercharge these risks in several ways. Malevolent actors can exploit generative AI to increase the speed, scale, sophistication — and potentially the persuasiveness — of voter suppression campaigns that employ deceptive practices. On Election Day, antagonists could use image and video-generated AI to fabricate disasters at polling places or to falsely depict election workers stopping eligible voters from voting. Using the technology underlying generative AI chatbots, they could spread falsehoods about the time, place, and manner of voting at a greater rate and scale than we have previously seen. In time, bad faith actors could also deploy generative AI tools to engage in sophisticated interactive disinformation campaigns geared towards voter suppression— propagating AI systems that adapt in real time to voters, draw from a set of optimally persuasive responses, monitor and respond to voters' emotional states through analysis of tone and mannerisms, and microtarget voters' demographic characteristics to more effectively deter them from voting.

Propagators of vote suppression tactics could also deploy AI tools to lodge mass, frivolous voter challenges and risk disenfranchising eligible voters. In the 2022 mid-term elections, a number of groups fueled by conspiracy theories fielded tens of thousands of voter challenges after combing through reams of voter registration records and miscellaneous public information. Going forward, similar networks could use AI to scrape and download databases of varying quality, from across the web and perform rudimentary, automated data matching against voters' registration records, or auto-prepare voter challenge forms, potentially allowing activists to file mass voter challenges on an alarming scale based on flimsy evidence. Amateur data matching— particularly when relying on flawed and incomplete datasets— is a highly unreliable method of verifying voter eligibility due to transposition errors, outdated data, diverse living arrangements, and other issues. The use of new AI tools to boost voter challenges raises the specter of improper wide-scale voter disenfranchisement, further-burdened election offices, and intensified disinformation campaigns to sow doubt in the validity of election results. New AI developments will allow election deniers to prepare and file voter challenges with increased speed and scale and will run the risk of lending bad faith efforts a false veneer of mystique or sophistication. To mitigate these dangers:

- Congress should pass legislation that restricts the knowing spread of false information about how, when, and where to vote and register to vote, with the intent to prevent votes or deter voters, that incorporates AI-related dangers;
- In states that allow private citizens to challenge voters' registration status, legislatures should shield voters from frivolous challenges and limit the forms of evidence, including AI-assisted evidence, that may be used to substantiate a claim.

AI and Participatory Democracy

In 2017, a deceptive campaign to undermine the federal policymaking process resulted in bots flooding the FCC with millions of comments from fake constituents pressing for the repeal of net neutrality rules. The bots employed natural language processing— not generative AI, but a predecessor to the technology— creating hundreds of thousands of "unique" comments with identical sentence and paragraph structures.

This episode could forecast imminent dangers for agency rulemaking and other policymaking practices informed by public input. It reveals just how vulnerable policymaking and participatory democracy are to malicious distortion through AI. New developments in generative AI greatly increase pre-existing risks by enhancing the sophistication of comments created through AI and making duplicitous campaigns to sway decision-makers much more difficult to detect. State-of-the-art open-source models are now sophisticated enough to launch such deceptive operations.

Conversely, agency officials could deploy AI to augment public participation and government responsiveness: they could use AI tools to incorporate public input more effectively and swiftly into policy, decisions, and implementation practices. But safeguards are needed to protect any such use of AI against bias, ensure accuracy, and guarantee tools' fitness for this purpose. To improve their ability to meaningfully process input from a broader swath of the public, agencies and offices could potentially deploy AI tools to help analyze, summarize, and aggregate public input — if legislatures and agencies require, implement, and enforce necessary safeguards.

Among the actions Congress could take to address these issues are:

- Facilitating investment into research and development focused on more secure, accessible, and privacy-protecting CAPTCHAs ("Completely Automated Public Turing Tests to Tell Humans and Computers Apart") that government bodies can implement where appropriate, with particular attention to the evolving capabilities of generative AI to defeat CAPTCHAs;
- Modifying administrative and public records laws to allow departments and agencies to disregard submissions and requests demonstrated to have been transmitted to the office through fraudulent use of bots or automated systems, including those powered by generative AI. Such a standard should *not capture* the use of generative AI to assist in drafting comments or use of form letters endorsed by actual humans; rather, it should capture the deceptive use of AI to significantly misrepresent the number of humans involved;
- Mandating that federal agencies' use of AI to substantially assist with analysis, summary, or aggregation of public comments, or to facilitate public communications, meet baseline requirements for training data quality, accuracy, reliability, and non-bias.

Conclusion

Professor Bruce Schneier of MIT has noted that artificial intelligence will increase the "speed, scale, scope and sophistication" of threats to our democracy. Put another way, the threats may not be new, but they will become even more dangerous. This, of course, calls for identifying and adopting appropriately tailored responses to mitigate AI's potential harms. But it also calls for addressing the underlying threats themselves: election disinformation, voter suppression, foreign interference, attacks on our election workers and infrastructure, dark money, and the broader efforts to undermine American democracy.