Statement of

JESSICA BRANDT

Brookings Institution

BEFORE THE UNITED STATES SENATE
ARTIFICIAL INTELLIGENCE INSIGHT FORUM

Concerning

ELECTIONS AND DEMOCRACY

November 8, 2023

Thank you Leader Schumer, Senators Rounds, Heinrich, and Young, distinguished members of the Senate, for inviting me to participate in this extremely important forum.

As you are aware, a broad, persistent, asymmetric competition is now underway between democracies and their authoritarian challengers, and the information and technology domains are critical theaters. Against this backdrop, recent advances in Generative Artificial Intelligence (AI) have outpaced expectations and risks that once seemed remote now appear immediate.

One risk lies at the intersection of these trends: that autocratic actors could use generative AI to shape public opinion within the societies they target, to exacerbate division, or to seed nihilism about the existence of objective truth, thereby weakening democratic societies from within.

Both AI and influence operations are evolving rapidly, so any forecast is inherently conjectural. At this juncture, generative AI appears more likely to exacerbate existing threats than to create wholly new ones. With that in mind, there are four, overlapping categories of risk that warrant policymaker attention:

(1) Deepfake technology could be used to, among other things, manufacture an "October surprise" or power robocalls designed to suppress votes or stoke division.
(2) Large language models (LLMs) enable the production of myriad unique content that could be used to create misleading perceptions of constituent sentiment, undermining democratic responsiveness; turbocharge state-backed propaganda; and make influence campaigns less discoverable.
(3) LLMs could increase the personalization, and therefore the persuasiveness, of information campaigns.
(4) In the face of these changes to the information environment, skepticism about the existence of objective truth is likely to increase.

Deepfake Content

This campaign season, for the first time, widely available, low-cost generative AI tools will enable users to produce audio clips in another person's voice, create realistic images of just about anyone, and deploy social media bots or robocalls with extraordinary conversational abilities. One possibility that has received considerable attention is that these tools could be used by foreign state actors to manufacture an "October surprise" that might have an impact on electoral dynamics – or raise concerns that it did. Importantly, the targets for this kind of operation are not just major political party candidates, but journalists and election officials, which themselves represent institutions of democracy.

Some have expressed fear over the possibility that deepfake robocalls could target voters with inaccurate polling instructions to suppress turnout within a population likely to have an outsized impact on the outcome of a race. In this and numerous other cases it is not clear why an actor would choose to use AI-generated content when more simply produced (or decontextualized) audio clips, images, or text are likely to be just as effective.

A potentially more novel possibility is deepfake robocalls that could engage voters about their concerns and respond in real time with convincing replies designed to polarize or mislead. Such an approach would be in keeping with Russia's longstanding goals of exacerbating division to weaken the United States from within. It would pose a challenge democratic deliberation, since the marketplace of ideas model only works when candidates and issue groups know -- and can therefore contest -- opposing arguments.

Large Volumes of Unique Content

Large volumes of semantically distinct, generated content could be used to overwhelm systems that take input from the public, including inboxes of elected officials, such as members of Congress who are themselves candidates for elected office. That could produce a misleading picture of constituent sentiment on any number of pressing policy issues, undermining efforts to discern -- and subsequently act on -- voters' views, and ultimately eroding democratic responsiveness. Such a development would be a problem in its own right. It could also weaken the United States's position in the broader geopolitical competition, since the ability to accurately gauge and adapt to citizen demand is a key asset in the political struggle between democracies and authoritarians. Such an effort would be of a piece with Russia's strategy of making it more difficult for democracies to govern themselves.

China, meanwhile, could leverage the technology to manufacture the appearance of consensus around pro-Beijing positions, for example by creating convincing -- and potentially even interactive -- personas at scale. Such an approach would be in line with Beijing's goals of portraying itself as an attractive global leader and drowning out criticism that would suggest otherwise.

Both Russia and China could use LLMs to create long-form propaganda content, either for state media or for proxy websites. In recent years, the Kremlin has repeatedly hired unwitting freelance journalists from within a target society, including the United States, to produce content for its online outlets as a means of disguising information operations. AI-generated content could be even harder to identify. Meanwhile, my own recent, co-authored research has demonstrated that the ability of Russian and Chinese state media to produce large volumes of fresh, relevant material on topics of concern for Putin and Xi have enabled them to dominate search results for key terms. It is not yet clear how generative AI will impact this dynamic. Although LLMs could enable defenders to level the playing field, authoritarian state media are likely more organized and motivated than their counterparts. It therefore appears more likely that LLMs will exacerbate the problem.

More broadly, generative AI could make influence campaigns less discoverable by decreasing the likelihood that propagandists will use repetitive or identical language across networks and accounts -- signals many bot detection systems rely on.

Content that is more personalized, and therefore more persuasive

Generative AI could increase the persuasiveness of information campaigns by enabling propagandists to test numerous messages at scale before subsequently proliferating those that are most resonant. This is a particular concern when it comes to China, since a major factor inhibiting the success of Beijing's information operations is that it has at times had trouble reading the societies it targets.

Chatbots, deepfake audio, and persuasive generated text could also enable phishing operations that are more personalized and convincing. Phishing was essential to arguably the most consequential element of Russia's "sweeping and systematic" 2016 influence operation: the hack and leak. Here again, journalists and election officials are important potential targets. So too are congressional campaigns -- which, importantly, are often small, start-up organizations without substantial resources to devote to cybersecurity practices.

Nihilism about the existence of objective truth

Public awareness of the proliferation of misinformation and the existence of inauthentic, generated content can exacerbate distrust, even in trustworthy sources. Already the possibility that AI-generated material could be circulating is leading people to disregard real images of the Israel-Gaza conflict, for example. This dynamic creates a "liar's dividend" that malign actors may capitalize on, dismissing genuine content as fake to deflect blame for misdeeds. Meanwhile, LLMs, which are increasingly being incorporated into search engines, are known to periodically generate authoritative seeming, but fabricated, sources -- even to support false information -- further deepening the morass.

Democracies depend on the idea that objective truth exists and can be used as the basis for self-government. Autocrats have no such need for a healthy information environment to thrive. In fact, quite the opposite -- autocrats benefit from widespread skepticism that the truth exists at all. It enables them to draw false equivalences, legitimate repressive practices, and consolidate power at home. Meanwhile pervasive skepticism makes it difficult for democracies to govern themselves -- potentially constraining their foreign policies and dampening the appeal of democratic systems. Perhaps this is why exacerbating this skepticism has been a goal of Russia's influence campaigns targeting the United States over a considerable period.

THE PATH FORWARD

Building resilience to this challenge will require action from all corners. Governments, AI developers, and members of civil society -- from journalists to political parties -- can work together to build norms around the responsible use of generative AI. The technical community can continue its work developing content provenance approaches that help users understand the origins of material they encounter online. Platforms can develop moderation policies that accomplish the same purpose. Journalists and researchers need to hold everyone to account. Those communities can also continue to expose information operations, whether or not they are generative AI-enabled.

Importantly, elections are a flashpoint, but not the start or endpoint, for authoritarian influence operations. Policymakers should develop safeguards to protect elections from the many risks that advancements in generative AI pose. But they cannot stop there, recognizing that the threat is ongoing and targets a wide range of polarizing or contentious political events and entities, using multiple tools of interference, often in conjunction with one another.

Thank you again for inviting me. I look forward to the discussion.