CENTER FOR DEMOCRACY & TECHNOLOGY

Election Integrity Recommendations for Generative AI Developers

July 2024

Authored by
Tim Harper, *Senior Policy Analyst, Democracy and Elections*

With over 80 countries and more than half of the world's population going to the polls this year, 2024 represents the largest single year of global elections since the advent of the internet. It has also been dubbed the 'First AI Election,' in light of the boom in widely accessible generative AI tools that have the potential to accelerate cybersecurity and information integrity challenges to global elections this year.

Addressing the risks that generative AI poses to elections requires an ecosystem approach. In part, that requires focusing on the *distribution* of deceptive AI-generated election content on social networks and private messaging services, and through robocalls, TV, and radio. While identifying solutions to the distribution of this content is absolutely necessary — and CDT has supported several initiatives to create voluntary standards for technology companies that help to prevent these risks — it is also necessary to consider the policies and product interventions that generative AI developers should adopt in order to prevent harmful content from being created on or spread through their apps and services.

Although we are halfway through this election year, it remains imperative for AI developers to quickly develop election integrity programs employing a variety of levers including policy, product, and enforcement to protect democratic elections this year and beyond.

Summary of Recommendations

Usage Policies

- Prohibit the generation of realistic images, videos and audio depicting political figures or political and electoral events.
- Prohibit users from conducting political campaign activities or demographic targeting — at least in the short term — and develop transparent goals for longer-term ethical development of political uses of AI.
- Prohibit the use of generative AI ad tools for political advertisements.
- Prohibit any conduct that interferes with elections, including actions that prevent someone from voting; mislead someone into voting differently or not voting at all; or incite, support, or encourage violence against election processes or workers.
- Refrain from using stored memory or other methods of personalization in generating responses to electoral and political queries.
- Refrain from releasing text-to-speech cloning tools that allow users to replicate the natural voice of real people, including political figures.

Product Interventions

- Develop user interface pop-ups or labels relating to known narratives of election mis- and disinformation.
- Disclose how recently a chatbot's training data was updated when providing responses to time-sensitive election queries.
- Promote, and direct users to, authoritative sources of election-related information.
- Allow users to report policy-violating answers in chatbots and policy-violating apps built using an API.
- Include an appeals option for enforcement actions.
- Commit to develop and embed machine-readable watermarks and metadata into image, video, and audio content using a common standard that social platforms can detect.

Enforcement

- Proactively enforce usage policies on elections at all times, not just during active election periods.
- Consistently deploy product interventions for the most common election lies, and create protocols to quickly deploy product interventions to address newly emerging, election-specific mis- and disinformation.
- Proactively test model answers to common election queries.
- Create escalation channels to accelerate leadership's visibility into emerging issues, particularly during high risk election periods.
- Adequately resource and staff policy and enforcement teams.
- Institute actor-level enforcement for election integrity policy violations.

Transparency

- Be transparent about election policies.
- Publish regular transparency reports on election mis- and disinformation and deceptive AI usage.
- Consult with civil society and facilitate researcher access to usage data.
- Develop relationships and communication channels with election administrators.

Usage Policies

Generative AI developers should implement usage policies to prevent users from generating deceptive election-related content. These should include:

1. **Prohibit the generation of realistic images, video, audio that deceptively depict political figures or political and electoral events.** This should include the generation of content depicting realistic-looking political and electoral events that did not occur (such as political protests, instances of alleged election fraud, and natural or man-made disasters), as well as deceptive depictions of political and electoral events that *did* occur. Political figures covered by this policy should include, at least, national and statewide (or equivalent) candidates for public office, national and statewide (or equivalent) public office holders, and politically appointed and confirmed members of state and national government such as members of the Cabinet or Supreme Court.

   We have already seen examples of deceptive AI-generated content that has the potential to impact elections. Media organizations, such as WIRED and Rest of World, and academics are tracking the use of generative AI in elections globally. For example, a deepfake robocall of President Biden sought to discourage New Hampshire voters from participating in the Democratic primary. Fake images of Trump being arrested have circulated online, and the NRCC has produced fake images and videos of migrant encampments in National Parks. A Ukrainian official shared a fake video of Paris being bombed. AI-generated images, videos, and audio recordings of fake or manipulated political events could play a role in spreading false information. A deepfake showing a candidate conceding or being killed, or a major city in their country being attacked by terrorists or foreign armies, could cause voters to stay home or change their voting behavior if seen on or near election day.

   Prohibiting this content will require proactive detection at the query-level to prevent the content from being created in the first place. This can be done by applying classifiers to prompts and outputs in order to detect likely policy violations. Input classifiers evaluate a prompt submitted by a user for any policy violations, while output classifiers and blocklists review the generated image, video, or audio for policy violations before output is displayed to the user. Prompt and output

risk mitigations should be audited in advance of elections and updated regularly to account for emerging mis- and disinformation narratives in order to prevent additional harmful content from being generated. Companies will have to consider whether differentiating between deceptive and non-deceptive realistic images, video, and audio is currently infeasible, in which case they may want to prohibit such content more broadly at least for now.

2. **Prohibit users from conducting political campaign activities or demographic targeting — at least in the short term — and develop transparent goals for longer-term ethical development of political uses of AI.** This prohibition should include efforts to collect, process, disclose, or infer data about elections and voting, as well as efforts to generate text, images, videos and speech for campaign messaging, fundraising, and advertisements.

   Generative AI is a powerful tool for political campaigning. It can perform analysis of massive datasets and generate content at a scale that would be impossible for most political campaigns. Using these tools on a voter file or a dataset of public opinions could enable campaigns to craft extremely targeted strategies for persuasion and get out the vote efforts. Indeed, several studies have shown that using off-the-shelf generative AI to automate and scale microtargeting can increase the persuasiveness of political messaging and advertising. These uses can be positive and negative. If used responsibly, these tools could enable challengers to more affordably contact constituents and compete with a well-funded incumbent, giving voters better information about their choices. However, if misused, microtargeting has the potential to threaten people's privacy — one of several reasons many platforms have limited targeting functions for political advertising — further misinformation, polarize or radicalize people, and endanger vulnerable groups.

   While campaigns and candidates are the most likely actors to use generative tools for electoral activities, especially for generating messages and ads, this prohibition should apply to all users because of generative tools' ability to infer and process data about elections in ways that could easily be misused. Some AI software, like Eagle AI, have already been developed to analyze voter registration records and generate and submit challenges to those registrations at scale in an attempt to remove eligible voters from voter rolls. In addition to creating misinformation about bloated rolls, this can result in overburdening elections offices by forcing them to unnecessarily review and validate hundreds of thousands of legal registrations. Preventing similar uses of their own products and services should be a goal of all major AI developers.

   We simply don't yet know enough about how political actors could use these tools. This technology has emerged and evolved quickly. As a result, research is only just beginning to reveal how campaigns are experimenting with generative AI. More study will be needed to effectively develop mitigation strategies capable of preventing harmful uses. Until then, developers should be wary of allowing their products to be used for campaign purposes.

3. **Prohibit the use of generative AI ad tools for political advertisements.** Some AI distributors, like Meta and Google, have begun to release new generative AI tools in their advertising suites to enable advertisers to adjust the content and appearance of their ads in ways that may be hard for viewers to detect. For example, they can remove distracting elements of a photo, expand an image using generative fill, or generate multiple versions of ad copy or images in order to A/B test their effectiveness.

   For the purposes of this recommendation, if a social media company would consider an advertisement to require disclosure as a political ad according to their existing policies, they should prohibit those ads from using generative AI ads tools. Political advertisements are defined differently across the major social platforms. Meta's Social Issue, Electoral, and Political (SIEP) Advertisements policy captures: paid content made by, on behalf of, or about politicians, political figures, political parties or political committees or that advocates for the outcome of an election; ads about elections, referendums, ballot initiatives or get out the vote campaigns; and, in some places, ads about sensitive social issues that may come up in political campaigns or influence the ways that people vote in elections. Google's policies are more focused, limited to what they term "election ads." Depending on the country, these include ads about candidates and officeholders as well as ads about elections, ballot initiatives, plebiscites or referendums.

   These tools offer powerful ways to manipulate the content in an advertising campaign. In a political or social issue ad, they could be used to alter the context for a statement made by a political opponent, or to modify a candidate's voice, skin complexion, or other characteristics in intentionally misleading ways. Studies have shown that darkening skin complexion in political ads can negatively bias opinions of a candidate. While tactics like these have been employed long before the era of AI, new ad tools make it easier than ever and could increase their prevalence.

   Fortunately, banning creation of this content in a suite of ad tools can be accomplished simply by disabling access to generative AI tools if the advertiser discloses that an ad is political in nature. However, this is not a perfect solution, as it doesn't prevent bad actors seeking to evade transparency or authenticity requirements from using these tools. Therefore, fully effective prohibitions also would require classifier-based detection.

4. **Prohibit any conduct that interferes with elections, including actions that interfere with someone's ability to vote, mislead someone into voting differently than they would with accurate information (or not voting at all), or incite, support, or encourage violence against election processes or workers.**

   Election interference can have many different goals — from affecting voter behavior or opinions to sowing chaos or polarization. But the three goals listed here cause the most urgent and significant harm to democracy. Therefore, election integrity policies should address these harms most directly.

   First, policies should prohibit conduct that could prevent someone from voting. This should include the generation of content misrepresenting the time, place, or manner of voting (e.g. "Republicans vote on Tuesday, Democrats vote on Wednesday"), the processes of administering elections (e.g. "Your ballot won't count if you vote by mail"), or the qualifications for voting or registering to vote (e.g. "You can't register to vote if you are under 21 years old"). While these specific examples apply to the United States, similar policies should apply anywhere AI developers' products and services are available.

   Second, policies should prohibit conduct that could mislead someone into changing their voting behavior — for instance, by seeking to deter someone from voting by misleading them into believing elections are useless, meaningless, or rigged. This is especially misleading in countries with a strong history of free, fair, and competitive elections with independent and credible election authorities. Policies should also prohibit users from generating content deceptively misrepresenting or falsifying candidate statements, endorsements or other content that seeks to discredit candidates, parties, or other institutions. Election watchers have already documented this sort of content, including deepfakes of Taylor Swift (supporting Donald Trump while claiming the 2020 election was rigged) and of Morgan Freeman (criticizing Joe Biden).

   Finally, these policies should prohibit generating content that directly or indirectly seeks to cause political violence relating to elections. Policies should prohibit conduct that incites, praises or facilitates violence at polling locations or elections facilities, such as content advocating for bringing guns to those locations or organizing rallies with signals of violence or intimidation. These policies should also prohibit targeting and harassment of election officials. Use of generative AI tools to identify election officials phone numbers or addresses, or to generate audio clips with death threats or other intimidating language could enable bad actors to more easily dox or 'SWAT' election officials. Attacks like these have already happened to many officials, including the Missouri Secretary of State Jay Ashcroft, and according to recent surveys have become more common since the 2020 election. Lastly, policies should prohibit the generation of malicious code that could be used to spoof elections websites or create targeted phishing campaigns aimed at election offices.

5. **Do not use stored memory or other methods of personalization in generating responses to electoral and political queries.** Recently, some AI developers have announced that they will start customizing responses based on information gathered from a user's previous chats. While there is always some level of stochastic difference in responses to queries — as is the nature of an LLM — using stored memory could produce larger variability that could be hard to predict. Employing this technique for political and electoral content could have some significant ramifications. For one, it may further polarize the electorate by feeding users answers that align with their pre-existing political leanings or expressed sentiments. For example, a user who queries a model to understand how climate change works may cause the chatbot to predict the user is left-leaning, and therefore provide them with left-leaning answers to questions about other topics like immigration or healthcare in the future.

   In addition to creating potential bias in responses to queries about political issues, it could also change a chatbot's answers to questions about ballot initiatives or a political candidate's conduct or positions on the issues. For instance, if a user has expressed support for Trump or the Big Lie, a chatbot may be more disposed to criticize Trump's political rivals or mislead users with biased answers about their voting history.

   Finally, using stored memory would also make it much harder for both the AI developer and independent researchers to detect patterns of non-violating but harmfully polarizing responses, since the answers may differ for each user. Testing for these patterns would require manually constructing complex personas to determine how this would affect responses in a political context — a costly and likely fruitless prospect. As a result, developers should take extreme care before deploying personalization for political topics until they can be confident that doing so does not produce unintentional harm.

6. **Do not release text-to-speech cloning tools that allow users to replicate the natural voice of real people, including political figures.** Recently, some AI developers have announced testing text-to-voice tools that use a seed file of a person's voice, which can be as short as 15 seconds, to generate a realistic voice impersonation that can then read any text submitted to the tool. The risks of this technology are vast.

   Already, we are seeing examples of the use of voice cloning for harm. It has been used to impersonate President Biden's voice to tell New Hampshire residents not to vote in the presidential primary, and by financial scammers posing as family members and corporate executives to steal identities. These tools pose unique risks for elections as well. With these tools, bad actors could impersonate election officials to spread disinformation that an election is rigged, tell staff to break the chain of custody for ballots, gain access to secure elections databases

and systems, or mislead voters about candidates conduct or positions. Influential international figures or celebrities could also be impersonated to further target or spread disinformation to language minority communities, which could have a major effect in swing states. Even with strong detection and enforcement efforts, content translated into other languages — the primary positive use case that has been posited for these tools — is also much more difficult for automated detection systems to identify accurately due to a number of factors, including a lack of high-quality training data and robust benchmarks to test these systems across languages.

Product Interventions

Develop product interventions to reduce the spread mis- and disinformation and mitigate the risk of hallucination. These should include:

7. **Develop user interface pop-ups or labels relating to known narratives of election mis- and disinformation.** While generative AI can supercharge the spread of election lies, many of the narratives we're seeing in this year are the same tired claims that have existed for years. For queries that contain common narratives that may contribute to harm against elections but do not violate user policies (such as anecdotal claims of wrongdoing by election officials or of unqualified voting by a noncitizen), AI developers should create canned responses containing authoritative information from election officials or fact checkers that can be deployed in labels and pop-ups year-round following design best-practices for misinformation interventions. During active election cycles, when new false narratives may emerge (like, for instance, "sharpiegate," which claimed that ballots marked with sharpies could not be read by ballot scanners in Arizona during the 2020 election,) AI developers should actively track the emergence of harmful narratives about elections, and should regularly update lists of keywords and classifiers used to trigger labels. Of course, maintaining these lists of keywords, tracking these narratives, and updating these labels will require a surge in staffing during major election years, and developing clear protocols for their deployment will help to speed implementation.

8. **Disclose how recently a chatbot's training data was updated when providing responses to time-sensitive election queries.** LLMs are trained on batches of data, so their responses may only reflect information that was updated months or even years ago. This can be problematic during an election when details about registration deadlines, voting locations, and candidate positions can change rapidly. Clearly stating when the LLM's training data was last updated, refusing to answer questions about events that occurred after the last update, and providing users with options to use more appropriate search tools can help prevent the dissemination of outdated and potentially false information. One example of putting

this policy into practice is Anthropic's announcement that the Claude system prompt will soon include information about its knowledge cutoff date.

9. **Promote and direct users to authoritative sources of information.** Platforms can provide authoritative information through canned responses, election information hubs, and hyperlinks to authoritative information in responses. When developing authoritative content to share with users, AI developers should prioritize official elections websites at the state and local level over third party voting information aggregators. Third-party websites like Vote.org, which provide vital services by helping to register voters, can struggle to keep up-to-date on the latest changes to voting laws in every jurisdiction across the country. Chatbots should also utilize resources from their parent companies, where they exist. For example, Google's Gemini should direct people to Election Information Panels that appear in Google search, and Meta's AI Assistant on Facebook and Instagram should direct voters to the Voting Information Center. Similarly, AI developers and social media companies that have generative AI tools should ensure that chatbots are connected to, and provide hyperlinks to, other election integrity products, such as third-party fact-checking, ads libraries, and voting alerts products, in order to make sure users receive similar misinformation interventions as they would on their social media feeds.

10. **Allow users to report policy-violating answers in chatbots and policy-violating apps built using an API.** Reporting violations should be easy for users — there should be a reporting option in the user interface of the chat so that users don't have to track down an email address or fill out a form in another browser window — and developers should make efforts to inform users how to do it. OpenAI's ChatGPT, Google's Gemini, and Anthropic's Claude already have this option in their consumer products. Users of apps built by third-parties using APIs should be able to report these apps for violating the developer's third-party API terms of service. User reporting not only acts as a way to identify bugs or unintended model behavior, but also as a way to incorporate the feedback of elections stakeholders like researchers, journalists, fact-checkers, civil society organizations, and election officials, who might ask a chatbot questions that the developer would not have considered.

11. **Ensure that enforcement actions have an appeals option.** No matter how accurate automated detection and enforcement systems are, they will still capture unintended content. In those instances, users should have a way to raise the issue for review. This might occur if a user is seeking to get more information about historical political movements or candidate positions in ways that do not request the chatbot produce harmful content that advances mis- or disinformation. In addition to giving model developers more feedback about the sort of content their classifiers are catching, providing an appeals option helps to vent user frustration and can also present an opportunity to educate users by directing them to more information about content policies.

12. **Commit to develop and embed machine-readable watermarks and metadata into image, video, and audio content using a common standard that social platforms can detect.** As AI becomes more realistic and difficult to detect, people are increasingly likely to believe that real content is AI-generated, and vice versa. This phenomenon, known as the Liar's Dividend, has eroded trust in facts. Watermarking is one potential step in helping to address this problem. By providing users with information about the origin of AI-generated content, provenance data and watermarking can help prevent misinformation and give people more confidence in trustworthy information.

    Watermarking and other forms of labeling have become a major focal point of AI developers and policymakers. Laws have been passed in the EU, the UK, and in several US States to require disclosure of AI-generated content. Earlier this year, social media companies announced an accord to invest in developing provenance methods such as watermarking technology, which was joined by many AI developers including Anthropic, OpenAI, Adobe, Meta, and Google among others.

    As a result, there are several provenance standards in use today, but most, including those developed by the Coalition for Content Provenance and Authenticity (C2PA), rely on the metadata of an image to confirm the provenance of a piece of content. Unfortunately, metadata can be easy to remove (e.g., being stripped if a user screenshots an image).

    In the immediate term, watermarking is unlikely to be a full solution to the use of generative AI to spread disinformation due to the ease of stripping watermarks and the availability of many LLMs that do not add watermarks to their outputs. Instead, these watermarks should be seen as one technique to help prevent the unintended spread of misinformation.

    Another approach involves synthetic content detection. Some AI developers, like OpenAI, have developed tools that can identify their own AI-generated content without relying on metadata. But these tools are still quite inaccurate for content made by others (98% effective when identifying AI content generated using Dall-E, but only 5-10% effective when identifying content generated using a competitor like Midjourney) and aren't yet available to the public. Google has also recently announced new investments in AI detection through its SynthID standard, which will potentially enable detection of AI-generated text in addition to audio, video and images — a promising but still new approach.

Enforcement

13. **Proactively enforce usage policies on elections at all times, not just during active election periods.** Usage policies on elections should be enforced at all times, rather than for discrete periods of time around the highest risk periods of elections (such as immediately before and after Election Day). Usage policies prohibiting the generation of deceptive election-related content are most valuable if they prevent the content from being generated in the first place, which will happen even in years without a major national election in a given country.

    Maintaining enforcement will require regularly updating lists of keywords and classifiers relating to policy violations, and, for some policies, maintaining and auditing a civic graph — a record of candidates, political figures, and other entities that should trigger enforcement of relevant policies. Creating and maintaining these lists is a full time job to which companies like Meta dedicate substantial resources and can't be easily achieved without consistent effort.

14. **Consistently deploy product interventions for the most common election lies, and create protocols to quickly deploy product interventions to address newly emerging, election-specific mis- and disinformation.** Many product interventions discussed in this report should be deployed regardless of the election cycle. Providing transparency about a model's knowledge cutoff date, including a feature to allow users to report inaccurate or harmful responses, directing users to authoritative sources of information about voting, and embedding machine-readable watermarks in AI-generated images, video, and audio are all common-sense tactics that should be active at all times. Labels and pop-ups designed to address the most common election lies that are used across different elections and different countries should also be enforced consistently. This can make it easier to track trends as new issues emerge in countries or if there is a surge of requests seeking to translate this content into new languages.

    However, it may make sense to deploy some authoritative information, labels, and pop-ups as break-the-glass measures, rather than as steady-state interventions. Specifically, authoritative information or other labels designed to address newly emerging narratives that undermine elections, which are more likely to be country- and context-specific, should be deployed as needed and should be layered on top of steady-state usage policies and product interventions. But it is more challenging to detect these urgent issues, and more time consuming to mitigate their risks, if core usage policies and product interventions are not being consistently enforced. AI developers should create clear protocols for determining how and when an intervention is needed in order to speed these updates through leadership approval, as these interventions are very time sensitive. Creating clear communication channels and protocols for these break-the-glass measures is one reason why AI developers should surge capacity to teams working on elections during active elections periods.

15. **Proactively test model answers to common election queries.** AI developers should proactively test for model responses to election-related queries, including in languages other than English. Several studies have shown that LLMs have continued to relay harmful content about elections despite recent announcements that AI developers plan to limit responses to election-related queries. Proactively testing model responses to election-related queries should be a standard element of any product or election policy release.

    Red-teaming efforts have also shown that election-related safeguards can be circumvented when models are queried in languages other than English. In one case, a model trained to redirect users seeking information related to ISIS recruitment was easily circumvented when the same prompts were written in Farsi. Therefore, model developers should invest in funding internal and external researchers to test their systems for responses to these questions in widely spoken languages in the United States, and should publish their findings so they can be held to account for their policy and enforcement decisions. Some model developers have made recent strides in publishing information about their testing. For example, in June Anthropic published a report on its election integrity testing efforts, probing how well the Claude models respond to questions about election administration. Similar efforts should be undertaken across the industry and in languages other than English.

16. **Create escalation channels to speed leadership visibility, particularly during high risk election periods.**

17. **Adequately resource and staff policy and enforcement teams.** Elections are dynamic events that require agility and quick reactions from tech companies to identify and remove harmful content as new trends emerge. This is true for AI developers as well. As new narratives of mis- and disinformation emerge, teams should have the resources and staffing necessary to handle the influx in queries seeking to generate harmful content seeking to influence an election. Some companies have recently emphasized that advances in AI have enabled faster training of classifiers and therefore better algorithmic detection of policy violations, but coordination during major elections requires more than automated systems — in-country expertise and relationships with key stakeholders, protocols to ensure effective escalation to leadership, and assistance from product, policy, legal, engineering and others are all necessary to address major risks during the height of an election, which can't be replicated with technology.

18. **Institute actor-level enforcement for these policy violations.** Users that repeatedly evade content policies and enforcement should face escalating restrictions on their accounts.

Transparency

19. **Be transparent about elections policies.** AI developers should publish their policies in clear language. Policies should be centralized in one place, and changes to these policies should be clearly flagged and dated. Policy clarifications released on a blog, in a press release, or through a spokesperson on social or traditional media should also be linked on the formal policy page.

    It can be difficult to get a complete sense of the policies an AI developer enforces when these announcements, clarifications, and other statements are not collected in one place or are not dated or version controlled. Clearly flagging changes and consolidating updates will make it easier for AI developers to respond to claims that they do not have or are not enforcing certain policies and also assists researchers in studying chatbots. For companies that are not exclusively AI developers, these attributes should apply not only to the main policy pages, but to AI policies specifically. For instance, while Meta and Google both have consolidated policy pages and version control for their primary policies, neither Gemini or Llama have version control for their specific policies.

20. **Publish regular transparency reports on election mis- and disinformation and deceptive AI usage.** Transparency reporting has become a common practice for social media companies in part due to regulations like the Digital Services Act. But this has not become a commonplace requirement for generative AI companies as of yet. Reporting should include information about the policies in place and how they are enforced, as well as disclosure of details about the number of queries that violate those policies, the accuracy and error rate of your classifiers and prompt refusals, and the composition of your trust and safety teams responsible for these policies and enforcement systems. The frequency of these reports can vary, but during major election years these reports should be at least quarterly. In addition to providing transparency reports about enforcement, companies should also invite input from government and civil society on questions that would be helpful to answer, and run analyses to answer them in transparency reports.

21. **Consult with civil society and facilitate researcher access to usage data.** Involvement and oversight by civil society organizations can provide important subject matter expertise and, when provided with access to the proper data, researchers can independently verify that companies are enforcing their elections policies. Consultation can also allow companies to build goodwill and explain their decision making processes. Consulting with civil society and ensuring researchers and journalists have access to affordable data tools can also provide an opportunity for companies to gather additional insights into the threats occurring on and off platform that may emerge throughout an election. Companies can facilitate researcher access to data by allowing users to donate their data to researchers. There are three ways developers can do this: first, allow users to opt to donate their data to a dataset that the company then allows certain vetted researchers

to access; second, build APIs that allow users to share their data with individual researchers; third, tacitly allow users to donate their data to researchers via browser extensions (a la the NYU Ad Observer).

22. **Develop relationships and communication channels with election administrators.** Building these relationships can help to educate election administrators about the opportunities and uses of generative AI in running elections, as well as the cybersecurity and information integrity risks. Open communication channels provide a route for election administrators to report concerns about model behavior or chat responses, as well as mis- and disinformation about their jurisdictions that may have originated from generative AI.

Find more about CDT's work on elections & democracy at *cdt.org/elections*.

Acknowledgements
This report was also contributed to by CDT's Miranda Bogen, Samir Jain, Kate Ruane, Aliya Bhatia, Gabriel Nicholas, and Laura Kurek.

This report is made possible with support by grants from the David and Lucile Packard Foundation and the Open Society Foundations.

The Center for Democracy & Technology (CDT) is the leading nonpartisan, nonprofit organization fighting to advance civil rights and civil liberties in the digital age. We shape technology policy, governance, and design with a focus on equity and democratic values. Established in 1994, CDT has been a trusted advocate for digital rights since the earliest days of the internet.