Report

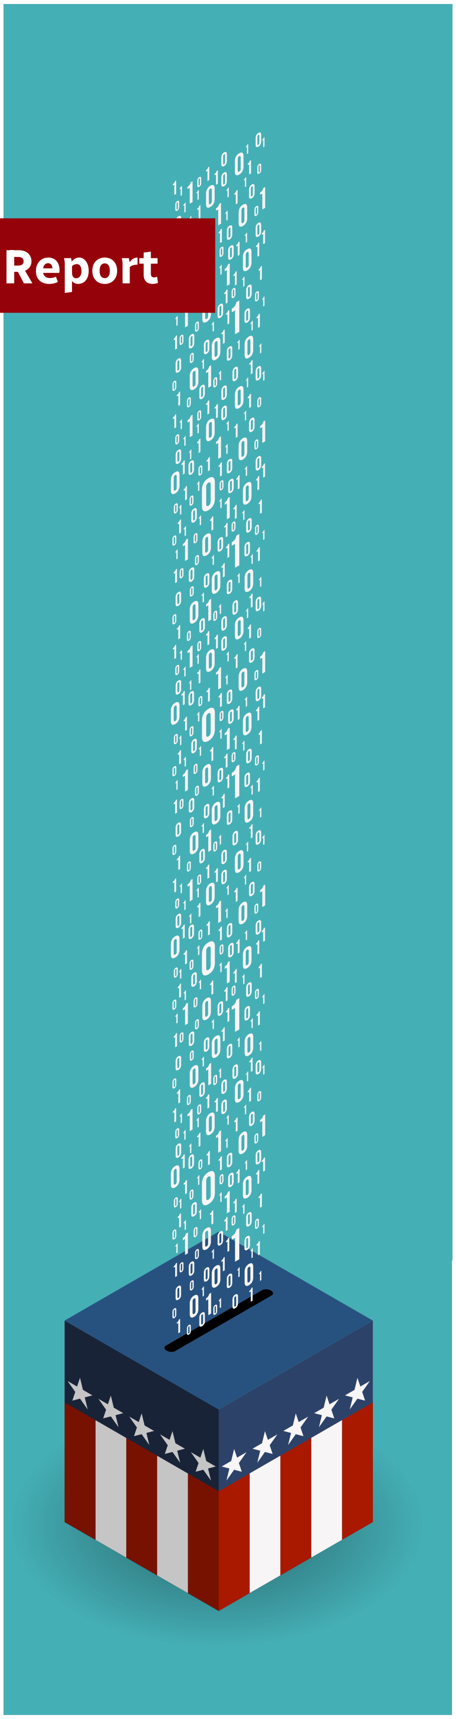

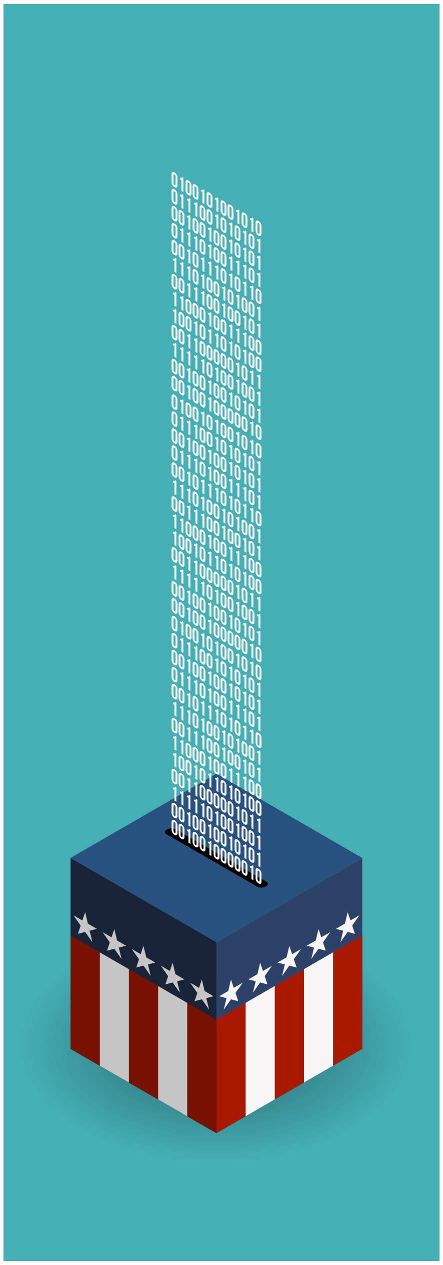

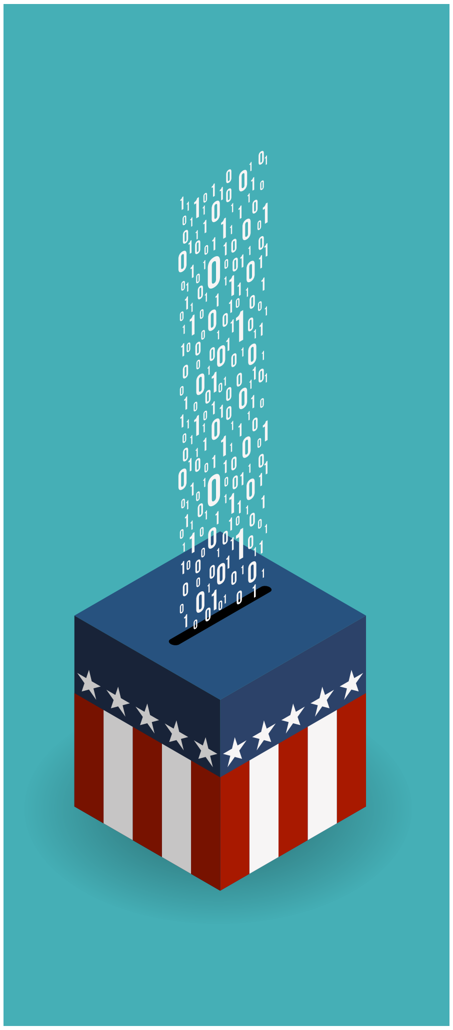

G | M | F Alliance for Securing Democracy

The ASD AI Election Security Handbook

Lindsay Gorman | David Levine

Table of Contents

Introduction

The typical local election official in the United States is a white woman who is older than 50, earns more than $40,000, does not have a college degree, and is less likely today than in the past to have extensive experience in election administration. She is responsible for protecting her jurisdiction's elections from a growing list of security challenges, from well-funded authoritarian regimes to conspiracy theorists, and is also continuously being asked to take on additional responsibilities in her role, often with little or no additional funding. Election officials have always been public administrators and logisticians, but after the 2000 presidential election, they were expected to be communication specialists and legal analysts. After the 2016 presidential election, they were expected to be cybersecurity experts, and during the 2020 campaign, they were even looked to for public health guidance. Now, it's 2024, and it appears they are expected to provide expertise on artificial intelligence too.

As the 2024 US presidential cycle hits its stride, election officials face a new technological challenge amid declining public confidence in their work: artificial intelligence (AI). Trust in the US political system today is alarmingly low. Almost a third of the US electorate still falsely believes the 2020 presidential election was stolen, a belief rooted in disinformation and arguments that have no evidence to support them. The prevalence of AI technologies magnifies the challenges for maintaining trust in election integrity and democratic governance, the success of which depends on numerous dedicated actors who contribute to the health and vibrancy of democratic life—poll workers, election officials, political campaigns, the media, and the voting public. This handbook is directed specifically at election officials, who shoulder an enormous burden on behalf of the American people. Our objective is to relieve some of the burden when it comes to navigating AI in the elections space. The handbook explores how AI tools could exacerbate vulnerabilities that malign actors may exploit to undermine the integrity of the 2024 presidential election—and future elections—and suggests steps for further protecting them.

Conceptualizing AI Threats to Elections

The US Department of Homeland Security (DHS) recently assessed that foreign adversaries like Russia, the People's Republic of China (PRC), Iran, and North Korea, as well as domestic extremist actors, are likely to try to disrupt the 2024 election with techniques that are different from previous elections, including the use of AI technologies:

"Nation-states seeking to undermine trust in our government institutions, social cohesion, and democratic processes are using AI to create more believable mis-, dis-, and malinformation campaigns, while cyber actors use AI to develop new tools that allow them to compromise more victims and enable larger-scale, faster, efficient, and more evasive cyberattacks."

AI amplifies challenges we have seen in previous elections. This handbook focuses on AI-enabled cyberattacks on electoral infrastructure and AI-powered audio and visual disinformation from foreign and domestic actors alike.

At issue are not only the risks posed by AI-enabled information manipulation, but equally importantly, the perception hacking that AI can facilitate. In perception hacking, a malign actor does not need to successfully execute a cyberattack, for example, to sow discord; giving the impression that an attack has been carried out—and spreading that view across the media landscape—could cause widespread damage to Americans' trust in elections. In an atmosphere of mistrust, AI-enabled perception hacking risks sowing even greater division and doubt in the voting system.

AI-enabled cyberattacks on election infrastructure

Advances in generative AI—a class of machine learning techniques that can create computer-generated content that mimics real-world content, such as images, audio files, videos, and text—have provided tools that make it much easier to execute damaging cyberattacks that could compromise electoral infrastructure. Hackers are successfully leveraging these developments in business e-mail compromise attacks, and foreign and domestic actors targeting electoral infrastructure may find them relevant as well. According to the 2023 Comcast Business Cybersecurity Threat report, nine out of ten attempts to breach customer networks started with a phishing attempt. Between January and February 2023, cybersecurity firm Darktrace observed a 135% increase in "novel social engineering" attacks, coinciding with the widespread adoption of OpenAI's generative AI chatbot ChatGPT.

AI can help cybercriminals perform large numbers of repetitive cyberattacks at little cost, map out election systems and their vulnerabilities, and generate compelling and realistic targeted content to dupe potential victims with more convincing phishing e-mails. Large Language Models (LLMs), such as ChatGPT and the WormGPT tool, have broadened the ability for attackers to generate fluid e-mail messages with correct grammar in non-native languages. For example, a foreign adversary could spoof emails from state election officials to local election officials in order to compromise their networks with malicious links. These tools could give them the ability to craft messages in English without the diction and syntax errors that can be tip-offs to their origin.

For election officials that use longstanding safeguards such as paper ballots, strong chain custody of protocols, and robust post-election audits, AI-enabled cyberattacks still appear unlikely to prevent voting or the accurate tabulation of votes. However, new AI-enabled attacks as well as the automation and streamlining of traditional techniques risk exacerbating cyber threats to election infrastructure. AI technologies could ultimately make it easier to access election systems and other components that feed into the voting process that are network connected, such as digital voter registration systems and cellular modems for transmitting election results.

Here too, the very perception of potential compromise could pose an even larger threat to trust in free and fair elections in the United States. An objective for many malicious actors attempting to weaken US democracy is to degrade confidence in the electoral process itself. While there is no mechanism that can fully secure election systems from cyber threats, there is also no evidence that any technical vulnerability has been exploited to alter the outcome of a US election. Simply citing the existence of a vulnerability does not prove a cyberattack occurred, much less that it changed an election result. If voters in the United States increasingly believe that electoral systems are insecure and can be easily manipulated—especially when there is no evidence to support these assertions—false and malign narratives about voter fraud, voting by mail, voting machines, and other aspects of elections may catch on more readily, further eroding confidence in US elections. As such, election workers must continue to walk a careful line of proactively communicating about and shoring up election infrastructure against new threats from AI-enabled cyberattacks, while at the same time not catastrophizing them.

AI-powered audio and visual disinformation

AI content generators exploded in the last year with the democratization of tools to create realistic AI images, the rise of AI voice-cloning for audio deepfakes, and the honing of realistic deepfake videos for manipulation. These technologies and their products are already making waves in democracies across the world. For example, just two days before Slovakia's closely contested 2023 parliamentary election in September, an AI-generated audio recording circulated on social media depicting Progressive Slovakia leader Michal Šimečka and local journalist Monika Tódová discussing how his party would rig the election. Slovak language fact-checkers did their best to debunk the clip, but not before millions heard the recording across Facebook, TikTok, and Telegram.

In a polluted information environment, the United States is highly vulnerable to political and electoral disinformation. As political candidates themselves experiment with AI tools to reach constituents through voter engagement or attack ads, social media platforms are trying to erect guardrails on requiring labeling of AI-generated content. However, enforcement of comparable standards has historically been mixed. In addition, the frequency of such content appearing on social media platforms is growing sharply. The AI communications firm Deep-Media estimated that 500,000 video and voice deepfakes would be shared on social media in 2023 alone. A well-timed AI image or audio recording conveying false information about voting locations or election results in a critical jurisdiction risks affecting voter turnout, the administration of the election, confidence in the election in 2024, or even the election outcome itself.

What to Watch Out for in 2024: How AI Could Affect the Elections

In the United States, election denialist narratives predate the explosion of generative AI and had already evolved into a sprawling, nationwide movement in the aftermath of the 2020 presidential election. Heading into 2024, election officials have these lessons from 2020 top of mind. In this section, we describe how AI-powered disinformation, cyberattacks, and perception hacking could supercharge these narratives and threats, and how that could impact election officials' work before, on, and after Election Day.

The use of AI could further fuel election subversion narratives and attempts at interference

Many efforts to overturn the 2020 election failed due to a lack of evidence of fraud and a willingness by election officials, judges, and others across the political spectrum to uphold legitimate election results and fairly adjudicate various claims of evidence falsely alleging that 2020 was stolen. These challenges continued to a lesser extent after the 2022 midterm elections, but were unsuccessful as well. In an age of AI-driven disinformation, the ability to readily fabricate images, audio, and video to support election denialist narratives risks lending credence to—or at least further confusion around—such claims and inspiring real-world action that undermines elections. These might include:

- **Candidates who refuse to concede after losing a race:** The vast majority of election-denying candidates who lost in the 2022 midterms conceded defeat shortly after it became clear that they lost, but some did not. Whereas in 2020 and 2022 such claims of fraud hit a dead end due to lack of evidence, candidates who lose in 2024 could rely on fabricated AI-generated images or audio to claim election fraud and contest the results.

- **Election jurisdictions refusing to certify election results:** Many local election officials have a ministerial duty to certify elections. In an effort to protest elections perceived as fraudulent, however, officials in some localities in 2022 refused to certify or unnecessarily delayed the certification of their elections. Fortunately, past efforts to stop the certification of elections results failed, due in large part to quick, decisive intervention from courts across the country that dismissed cases for lack of evidence. Although a few of those individuals who sought to block certification have faced serious and long-term consequences, certification disputes are likely to continue in 2024. In an era of AI-powered disinformation, jurisdictions refusing to certify elections could point to fabricated images and messages as a justification for refusing to certify the results. Judges or election officials adjudicating these claims could discover faked "evidence" of fraud and rely on false text, images, or video to cancel or overturn a legitimate election result. Even if the truth eventually came out, the damage to the public trust may be harder to recover from.

- **Election deniers discrediting voting machines:** Perhaps the most widespread claim of election fraud in 2020 surrounded voting machines themselves, with actors falsely claiming that Dominion voting machines had been compromised to alter the election results. In 2022, attacks on the legitimacy of voting machines continued, often lobbed in an attempt to bar the use of the machines and force a hand count, which is often a lengthier and less accurate process. AI-generated images or videos could potentially be used to fabricate evidence and support claims of voting system compromise or blow regular and inconsequential issues with voting systems—such as occasional jamming—out of proportion to advance this narrative. In addition to undermining public confidence in voting systems, such attacks could confuse judges or persuade elected officials to take legislative action based on these fallacies.

Across all these potential narratives, malign actors could also claim the 'liar's dividend', casting doubt on the veracity of evidence that buttresses trust in US elections and painting that evidence as faked by AI.

The use of AI could supercharge the deluge of public records requests

Since the 2020 election, public records requests directed towards election offices have increased dramatically, driven in large part by prominent election deniers. When such requests are made in good faith, they serve an important role in increasing government transparency and accountability. But when weaponized, they can overly burden and even undermine the work of election officials. In the run up to the 2022 midterms, these requests often forced election offices to divert valuable, already constrained resources that would have otherwise been devoted to administering the elections, to addressing these requests. They also imposed significant additional costs on election jurisdictions. Fortunately, election officials largely succeeded in preventing these requests from turning into election administrative crises. In 2024, however, election officials in many jurisdictions are still working with scarce resources. An ability to automatically generate, tailor, and send records requests at massive scale by leveraging generative AI could turn an increasingly onerous burden into an overwhelming distributed denial of service (DDOS) attack on election resources.

The use of AI could increase harassment and threats against election officials and workers

Election "rigging" rhetoric continues to inspire threats of violence, harassment, and abuse against election workers, which in turn is leading many election administrators and poll workers to leave the field. Just before the 2022 midterm elections, the Federal Bureau of Investigation (FBI) warned there were high levels of violent threats against election workers and officials in seven states. Following the midterms, social media threats directed at Maricopa County, Arizona's Supervisor Bill Gates forced him to move to an undisclosed location over safety concerns. Fortunately, these threats did not undermine the integrity of the 2022 midterms, and several states have recently tried to take steps to better protect their election workers. AI-powered disinformation that paints election officials in a negative or salacious light could be used to harass election officials online and offline—creating seemingly endless content and threats. In a combustible environment where discerning fact from fiction becomes even more difficult, the spread of mis- and disinformation could indeed lead to demonization of public officials and more overt threats to their lives and livelihoods.

Recommendations

The proliferation of accessible AI tools gives foreign and domestic actors additional means to undermine public trust in the 2024 presidential election. Malign actors seeking to undermine democratic processes and erode trust in democratic institutions are using AI to create more believable mis- and disinformation campaigns. Electoral infrastructure already at risk of cyberattacks or the perception thereof faces the prospect of larger, faster, and more effective intrusions from actors using AI tools to compromise more victims and harder, more secure targets. Election officials need to act with urgency in response. We recommend election officials take the following measures:

Incorporate AI risks into training and election planning

1) Add external cyber and AI expertise to your office

After hackers affiliated with the Russian government successfully penetrated Illinois' voter registration database during the 2016 presidential election cycle, the state started the nation's first cyber navigator program, employing cybersecurity experts with responsibility for geographic zones across the state to work with local election officials to conduct comprehensive election security risk assessments of each jurisdiction. Since then, other states have launched similar programs to assist under-resourced jurisdictions by managing their cyber risks, sorting through the onslaught of risk information, advice and available services, and fast-tracking mitigation efforts.

States that have cyber navigator programs should supplement existing programs with AI expertise around AI-enabled phishing campaigns and malware attacks. States that do not already have a cyber navigator program should strongly consider creating one and adding AI experts or expertise. If that is not a possibility, states should work with their local jurisdictions to ensure the latter have the requisite AI and cyber expertise in their offices. States could do this by providing targeted training to localities who need it, helping vulnerable localities find individuals with AI and cyber expertise, facilitating cross-jurisdiction programs that allow jurisdictions with more technical expertise to assist those with less, and or providing greater awareness of resources that already exist, such as those provided by the Cybersecurity and Infrastructure Security Agency (CISA).

2) Provide internal AI-specific training to ensure there is adequate awareness of threats

Election officials, particularly full-time officials, should take advantage of available AI-specific training. They could carry out CISA's trainings on how to spot and prevent AI generated phishing attacks, reach out to see what their state's Chief Information Officer offers, and or shadow local information technology (IT) office staff (if a jurisdiction has them) that are assigned to monitor AI and other cyber-related issues. The sooner election officials can learn about AI, the more readily this knowledge can be bolstered across their offices and shared with voters. All election administrators, at a minimum, should seek to learn what generative AI is, how it works, and best practices for engaging with it. For example, election workers should be trained on identifying the signs of electronic communication fraud to help reduce the risk of falling victim to AI, as well as how to verify the accuracy of information received from AI-based sources, how to communicate about AI-generated content, and how to report suspicious content or activity. Subsequent testing could ensure that such training is being successfully understood and applied.

As the Brennan Center for Justice has noted, AI training for election officials could also include how to remove data from their websites that could be used to personalize AI-generated communications, such as phishing emails or cloned voice messages;help election workers identify AI-generated content; and counter AI-fueled false narratives. It is also important that AI training occurs in a timely manner; otherwise, it could risk becoming outdated since generative AI is evolving so quickly.

3) Form a local election security working group with an AI component

If they have not done so already, local election officials should consider forming working groups with representatives from the private sector and local law enforcement and emergency management services, as well as experts in cybersecurity and IT. One of the authors previously served as the Ada County, Idaho elections director and held regular meetings in the run-up to an election with a similar group of individuals to flag any potential problems with administering the election and develop contingency processes to resolve them in case they did arise.

The members of these groups could help election officials grapple with the rapid changes being created by an acceleration in AI technologies. For example, the working groups could help localities determine whether AI-based cyber solutions should be utilized to help counter AI-powered threats to breach voter registration networks, steal voter data, and cause other damage. They could also examine whether traditional cybersecurity solutions such as multifactor authentication, regular patching, endpoint security and well-implemented firewalls suffice. Moreover, a working group could help ensure that local election officials receive up-to date information and alerts about threats and vulnerabilities to their systems, as well as build or deepen communicative relationships among key stakeholders and local officials. DHS has taken meaningful steps to improve information sharing and raise cyber threat and incident awareness with state and local election officials. These efforts are only likely to grow as it hires additional regional election security advisors to work directly with state and local election officials nationwide to strengthen election security and impenetrability. However, DHS is highly unlikely to have the staffing and resources to provide tailored assistance to every local jurisdiction in the country ahead of the 2024 presidential election. As such, local election officials should form working groups, identify problems, and then appeal to state and federal entities for assistance as needed.

4) Prioritize trustworthy communication, including on AI itself

Generative AI can be used to make the job of election officials more difficult, either by quickly generating content with inaccurate information to deceive voters or misrepresenting an election official's action to attack their integrity. One of the best ways election officials can mitigate these challenges is to ensure that their offices are trusted sources for constituents. This can be done in a myriad of ways, many of which are outlined in The Election Group's Telling Our Story: An Elections Communications Guide. The guide suggests establishing a social media presence to reach more constituents with facts and connecting with local reporters and partners in the community before issues, including AI-enabled ones, arise. Elections officials can build trust with reporters, community stakeholders, and voters by offering these constituents opportunities to see how an election is administered and creating an open channel of communication on topics such as AI's potential impact on the election. The latter could help election officials identify which community organizations, businesses, and reporters to call on to help counter election threats that arise from AI.

AI also presents a messaging challenge for election officials. They want to be honest and transparent about the threats posed without causing alarm or eroding trust in the system. Election officials should incorporate potential AI threats into their crisis communications plans so that if any such threats arise, election officials are ready to respond proactively. Election officials should have talking points ready; some suggested points include:

- **Define AI:** “There is a lot of talk about AI and elections right now. Let’s be clear about what exactly we mean by AI. Artificial intelligence has been used for a long time; most of the chatter right now is around advances in generative AI—a class of machine learning techniques that can create new fake content, such as fake images, audio files, videos, and text. You can see the potential problems as it is increasingly easier for anyone to make realistic fake content.”
- **Share how your office is approaching AI:** “I’m thinking about this challenge in two ways: First, as a supercharged cybersecurity threat—that means we have doubled-downed on our cyber protocols and cyber hygiene. Second, as a factor that could make it easier for false election narratives to be shared widely. So that’s why I’m talking to you. I want to be clear about how the process works in [insert jurisdiction here] and who you should turn to for trusted election information.”
- **Welcome inquiries:** “I’d encourage you/viewers/voters with any election questions to come to the election office for a tour and to stop and consider any information you come across online that disparages our election system—like allegations of cheating or a rigged system—before you share and if you are troubled by what you see, check our website, [Twitter handle/Facebook page] send us an e-mail, or give us a call.”

5) Encourage the creation of public record requests reading rooms to address the possibility of voluminous records requests due to AI

As elections in the United States become more contentious, the burden on election administrators is also growing. Many are saddled with greater responsibilities in a climate of inadequate funding and face increased threats. Since 2020, a growing number of election officials have faced huge numbers of record requests that perpetuate specious allegations of election rigging and overburden already under-resourced election offices. Through generative AI, such requests could be produced in real-time at a massive scale.

If resources allow, one way election officials could try to counter this threat is by setting up a public records reading room, in person and or virtually, as the US Election Assistance Commission has previously suggested, that proactively offers members of the public access to copies of public records that are frequently requested. This room could include the policy for requesting records, as well as a list of frequently asked questions with subsequent explanations for documents that are often asked for, but that are not publicly available. For a virtual reading room, election offices could also consider using a version of a CAPTCHA test to ensure that only actual persons are using the reading room. If members of the public cannot find the documents they are seeking, they could then be directed to fill out a public records request form that would subsequently be reviewed by a staff member. If records are proactively offered to the public in this manner, election offices may be able to cut down on the time, money, and labor spent processing the most common public records requests.

6) Simulate potential AI threats in the cyber and information domains to election infrastructure

When there are new or emerging threats to election infrastructure, exercises to help prepare for these threats can help ensure that the policies, procedures, and equipment in place to secure the election are adequate. This includes: voter registration databases and associated IT systems; IT infrastructure and systems used to manage elections, such as for the counting, auditing, and displaying of election results, as well as post-election reporting to certify and validate results; voting systems and associated infrastructure; storage facilities for election and voting system infrastructure; and voting locations, including election day polling places, early voting locations, and drop boxes.

One example of an exercise that could help determine which systems and processes work well against AI-enabled threats is a mock election where white-hat security researchers are invited to try to attack election systems using AI-driven phishing campaigns. If localities do not have adequate budgets to fund these mock elections themselves, they could perhaps look to partner with their states, their vendors, federal departments like CISA, local civic engagement groups, and or local universities to make up the difference. Mock elections can also help with election education and research efforts, in addition to testing the election infrastructure.

Another way to test election infrastructure against potential threats, such as the use of open access AI tools to create and amplify disinformation, is through tabletop exercises like the one Arizona recently completed. Tabletop exercises allow election officials to meet in an informal, classroom-esque setting to discuss their roles and responses under various threat scenarios. In 2020, election officials across the country gamed out how malign actors could disrupt the 2020 election through cyberattacks and information operations, and many touted these efforts as critical to helping ensure the security of the 2020 election. Conducting such exercises against AI-driven threats can help prepare election offices to respond effectively and in real time to AI-driven election hoaxes, particularly in states that have become magnets for election rigging conspiracy theories or are expected to be closely contested in 2024. Michigan is planning to hold a similar election exercise this spring, and additional states should follow suit.

Double down on cybersecurity

7) Secure official election websites

In the run up to the 2020 election, the FBI identified numerous fake election websites imitating federal and state elections sources using .com or .org domains. Ahead of the 2022 midterm elections, the FBI and CISA warned voters to be cautious around websites that solicit voting information without using .gov due to the risk of foreign information manipulation campaigns. With generative AI increasingly able to produce sophisticated fake images and even fake web pages, the ability to rely on visual cues, such as website quality or formatting, to distinguish authentic sites from spoofs will decrease, making authenticity symbols such as .gov domains even more vital. In an increasingly confused information environment, these clear symbols establish benchmarks for trust.

Unfortunately, only 25% of election websites currently use a .gov domain, which means that the federal government has not verified their authenticity and that voters cannot clearly tell whether the information on them is from an actual government agency or a fake website. It is imperative that as many localities as possible transition to a .gov web address before the 2024 presidential election to help the public more easily identify accurate election information. CISA has increasingly made it easier to transition to .gov by eliminating registration fees for the .gov top-level domain and offering federal funds to help jurisdictions transition to new domains. States could also mandate the use of a .gov, or at minimum .us, domain for their local election offices websites, as Ohio's secretary of state did in 2019.

8) Prioritize good cyber hygiene

While hackers are increasingly using generative AI tools for cyberattacks—including phishing emails, keystroke monitoring malware, and ransomware attacks—basic cybersecurity hygiene can still protect against the vast majority of attacks. Yet, some election jurisdictions are still lagging behind in basic cyber hygiene. In the face of AI-enabled cyberattacks, strong identity verification and authentication mechanisms are necessary.

One way jurisdictions can improve their cyber hygiene is to ensure that passwords for important accounts and data are in line with current best practices. For example, passwords that consist of a few words put together as a passphrase are stronger and easier for users to remember than longer, more complicated passwords. It is also important for election workers to avoid password reuse, use password managers to store unique passwords (and generate new ones), and periodically change their passwords to reduce their cyber and AI risk.

Additionally, some jurisdictions can improve their cyber hygiene when communicating electronically with voters, candidates, and the general public by enforcing multifactor authentication, which requires users to present a combination of something they know (e.g., a password), something they have (e.g. an authenticator app on a cell phone that receives a code or request to verify), or some form of biometric identification (e.g., a fingerprint). For five straight years, the leading cause of breaches has been stolen passwords. However, multifactor authentication can blunt the impacts of many of these potential phishing attacks by requiring attacks to crack more than two factors of verification to gain access to someone's system, and attackers are far less likely to have access to a user's other authentication factors, such as a cell phone.

9) Encourage election officials to keep hard copies of voter registration cards in an organized manner that can quickly and easily be accessed in case of an AI-enhanced ransomware attack

In the 2020 presidential election, a ransomware attack reportedly disabled Hall County, Georgia's voter signature database for authenticating mail ballots. Ransomware attacks disable affected computer networks with encryption that can only be unlocked with keys provided once the victim pays the attacker's requested ransom. This ransomware attack did not stop Hall County employees from verifying mail ballot signatures, but it did reportedly slow down the process by forcing employees to manually pull hard copies of voter registration cards to compare the signature on the cards with the signature for the mail ballot in many cases.

In most states, signatures are used to validate mail ballots returned by voters. Written on the envelopes that sheath the ballots, they are often matched by election workers against digital signatures on file with state and local election offices. As a result of these technological innovations, some election offices have moved away from having their voter registration cards stored in a manner that they can quickly access. These offices thought, understandably, that the need for having these cards easily accessible had diminished, and that the space could be better used for other election tasks.

AI technology is now beginning to drive ransomware to new heights by making these attacks more effective and increasing their volume. The frequency of ransomware attacks continues to increase, and attackers are likely to look to generative AI to craft increasingly effective attacks. For example, the number of reported attacks on municipalities across the world has already doubled since 2022 and more than quadrupled since 2021. If an AI-enhanced ransomware attack successfully disables another jurisdiction's voter signature database, that could lengthen the time it takes for that jurisdiction to process mail ballots, a delay malign actors could try to exploit.

Looking to the future: leverage AI and new technologies for responsible use

10) Consider piloting content authenticity technologies to build trust

As AI-generated images and video become more sophisticated, new classes of technologies to defend against these deepfakes are rapidly being developed. While additional road-testing may be needed for use in an election setting, content authenticity technologies are one such promising class of technologies. When an image or video is captured by a camera or generated by another tool, content authenticity technologies embed data about how, when, and potentially by whom this content was created into the content itself. This metadata is updated with each subsequent change to the file, so that any user can view its modification history. In other words, content authenticity technologies show how an image was created and how it has been altered over time. In this way, viewers can track for themselves the lifecycle of an image or video from point of capture or creation to when it appears on their screens through a browser or social media platform. While not a silver bullet, creating independent verification trails for images and videos can help information consumers distinguish quality visual content from fabrications.

The White House's October 2023 Executive Order on AI has called for the development of standards and best practices for content authentication technologies so that Americans know that government communications are authentic. The Coalition for Content Provenance and Authenticity (C2PA), a multistakeholder group of leading technology developers, media outlets, social media and civil society actors pioneering content authenticity technologies, has also released an open internet standard for their implementation. Voluntary AI commitments from industry in partnership with the White House have also included the need to authenticate and disclose AI-generated content.

Some election offices, particularly those in well-funded, tech-savvy jurisdictions, could consider using content authenticity technologies in official communications deeper in the 2024 election cycle to pilot new approaches for increasing trust in the visual information environment. Such approaches would be even more impactful when these technologies are adopted on a larger scale by technology companies and social media platforms. OpenAI, for example, has announced that it will implement the C2PA's digital credentials in its image generation tool DALL-E 3 as part of its 2024 elections plan. By building content authenticity technologies into their communications platforms when they put out image and video content, election officers can help guard against AI-enabled disinformation. One way to do so concretely is for election offices that already have Adobe Photoshop or Microsoft Designer to use the "Content Credentials" feature. Areas where election offices can experiment with authenticated content could include filming the logic and accuracy testing of voting equipment, the group of pre-election procedures that ensure that the voting equipment and ballots used in an upcoming election can properly display the ballot, collect votes, and tabulate results. These tools could help voters distinguish quality and authoritative information from spoofs and fakes that could be used to undermine an election.

11) Add generative AI guidance to election office security policies to ensure responsible use

An election office's security policies play an important role in protecting the office from reputational and data losses. One consideration for updating these policies is when new technologies, such as generative AI, could be adopted. As the Center for Internet Security has previously noted, election offices that have yet to create guidelines on generative AI platforms in their organization's cybersecurity policies should begin doing so. These guidelines should specify the types of data that can and cannot be entered into publicly-available generative AI tools to avoid exposing sensitive information to unauthorized parties, how to review content that is published that includes AI-generated material, and where to report suspicious messages that may have been created with the assistance of generative AI. As election officials know firsthand, security transformations generally do not happen in days or even weeks. They often take place over months, or even years, and the constantly evolving nature of AI justifies action sooner rather than later. If election offices fail to review and update their policies and procedures when they implement AI technologies, they could possibly expose their organization to an increased risk of data loss.

There may be instances where election officials want to consider using AI tools to make their election operations more efficient. For example, as the US Election Assistance Commission noted in its AI Toolkit for Election Officials, "Text-based AI tools will become more widely used in the workplace to make work easier. Some future application could be drafting emails, populating spreadsheets, or making first drafts of presentations and reports." The National Institute of Standards of Technology's AI Risk Management Framework provides a basis for how to use AI technologies while mitigating risk, with many suggestions that are relevant to election officials.

It would not be particularly surprising for election officials to express interest in using AI to help with tasks like drafting emails to prospective poll workers, populating spreadsheets with poll worker assignments, and helping write first drafts of voter education materials, particularly in light of the resource constraints that many election officials face ahead of 2024. However, before election officials go full speed ahead with such efforts, it is important for them to have a framework in place for managing potential risks. Jurisdictions should consider consulting CISA, which is accelerating its efforts on AI in the 2024 election cycle and supporting the use of AI-enabled tools to strengthen federal cyber defenses after the White House's Executive Order on AI. Since DHS designated election infrastructure as critical infrastructure in 2017, CISA has offered free resources and services to help election officials improve their resiliency against cyberattacks, and an increasing number of elections officials are pointing to these resources as a source of key supplemental support for successfully securing their elections.

Moreover, it is important to note that any use of generative AI in an election office also requires robust human oversight. For example, election administrators must review any outputs they receive from LLMs for accuracy, bias, and hallucinations (the generation of false or misleading text) to avoid the accidental dissemination of inaccurate or controversial information in a high-profile setting. If an election official wants to use generative AI to help put together a voter education guide for its voters, for example, it is imperative, at a minimum, that election officials review the content in the guide to verify its accuracy. Election officials could also consider providing ethical AI training for full-time workers so that they are prepared to assess AI tools, with priority given to testing models used in the highest-stakes environments.

Conclusion

The US national security and intelligence community have stated that foreign adversaries and domestic actors are likely to attempt to undermine the integrity of the 2024 presidential election, just as they did in 2020, and that the proliferation of accessible AI tools will likely improve their capabilities. It is critical that election officials have the necessary personnel, tools, and resources to repel attacks from increasingly sophisticated malign actors. Implementing the measures recommended in this handbook will improve resilience to AI-related threats ahead of November 2024.

About The Alliance for Securing Democracy at GMF:

The Alliance for Securing Democracy (ASD) at the German Marshall Fund of the United States (GMF) is a nonpartisan initiative that develops comprehensive strategies to deter, defend against, and raise the costs on autocratic efforts to undermine and interfere in democratic institutions. ASD has staff in Washington, DC, and Brussels, bringing together experts on disinformation, malign finance, emerging technologies, elections integrity, economic coercion, and cybersecurity, as well as Russia, China, and the Middle East, to collaborate across traditional stovepipes and develop cross-cutting frameworks.

securingdemocracy.gmfus.org | gmfpress@gmfus.org

About the Authors:

Lindsay Gorman (@LindsayPGorman) is the senior fellow and head of the technology and geopolitics team at ASD at GMF and a venture scientist with Deep Science Ventures. Her research focuses on the emerging technology competition between democracies and autocracies. Lindsay previously served as a senior adviser at the White House on AI and national security and has testified as an expert on these matters before the US Senate and House of Representatives. She is the former CEO of a technology consulting firm and has served as an expert contributor to the Cyberspace Solarium Commission, a technology adviser to US Senator Mark Warner, and a fellow with National Academy of Sciences. She holds a BA in Physics from Princeton University and an MS in Applied Physics from Stanford University.

David Levine (@davidalanlevine) is the senior elections integrity fellow at ASD at GMF, where he assesses vulnerabilities in electoral infrastructure, administration, and policies. David is also an adjunct professor at George Mason University, an advisory committee member for the Global Cyber Alliance's Cybersecurity Toolkit for Elections, an advisory council member for The Election Reformers Network, a member of the Election Verification Network, and a contributor to the Fulcrum. Previously, he worked as the Ada County, Idaho Elections Director, managing the administration of all federal, state, county, and local district elections.

With additional research from Krystyna Sikora and Gabriele Sava.

Acknowledgements:

The authors would like to recognize the following individuals and organizations for their invaluable feedback and support: Andy Parsons, Josh Lawson, Larry Norden, Noah Praetz, TJ Pyche, Mounir Ibrahim and Maggie Munts. The authors consulted publications by some of these individuals, which were invaluable resources in developing the AI Election Security Handbook.

The authors also co-hosted a private roundtable under the Chatham House Rule on December 13, 2023 that reinforced the importance of some of the report's recommendations broadly and specifically helped shape the election official talking points in recommendation four. However, the contents of this report reflect exclusively the views of the authors.

Disclaimer:

The views expressed in GMF publications and commentary are the views of the author(s) alone.

As a non-partisan and independent research institution, The German Marshall Fund of the United States is committed to research integrity and transparency.

Cover photo credit: idspopd | Adobe Stock