1

# Artificial Intelligence and Democracy:

## *Campaigns, Elections, Movements, and Deliberation*

Bailey Flanigan, Florian Foos, Archon Fung, and Charles Stewart III

**Abstract:** This chapter examines how generative artificial intelligence (AI) is transforming four domains of democratic practice: political campaigns, election administration, social movements, and citizen deliberation. Using a task-based framework from labor economics, the authors analyze which political functions AI can substitute for, leave unchanged, or augment. In campaigns, AI accelerates content production, enables multilingual outreach, and demonstrates persuasive capabilities exceeding traditional methods. In election administration, adoption remains limited, with AI's tendency toward hallucination restricting its use to back-office tasks. For social movements, AI lowers barriers to professional content production and enables mobilization tools in hostile environments. In citizen deliberation, pioneering applications – Taiwan's vTaiwan and Google DeepMind's Habermas Machine – use AI to facilitate large-scale deliberation and bridge linguistic divides. Authors identify tensions between AI's capacity to expand participation and the risks of centralizing control among platforms and well-resourced actors. They emphasize that AI's democratic impact depends upon organizational integration and deployment of these technologies, competitive political and economic dynamics, regulatory frameworks, and political leadership rather than technological capabilities alone.

Every major advance in information and general purpose technologies changes democracy because that technology frequently enables political actors, as well as economic and social actors, to do their work more effectively. Technology can enable insurgents to displace incumbents, as some resources and capabilities (e.g. hegemonic support from mainstream media actors) lose value and other capabilities (e.g. ability to command attention on social media) increase in political value. On the other hand, technological change can also enable incumbents to tighten their grip on power if new technologies amplify political capabilities according to previously existing asymmetries of resources and power. In this chapter, we look into the near future to venture some reasoned speculation regarding the ways in which the proliferation of some artificial intelligence technologies may change democratic practices.

"Artificial intelligence" and "democracy" denote two enormous domains of activity and social inquiry. We therefore limit the scope of these reflections in both domains. Within "Artificial Intelligence," we focus on the political applications of Large Language and Multi Modal Models, and generative artificial intelligence. For the most part, we disregard other parts of AI, such as robotics, expert systems, and the possibility of Artificial General Intelligence. "Democracy" is an even larger domain. We limit our consideration to four of its aspects: campaigns, election administration, social movements (and governments' efforts to control them), and citizen deliberation. Furthermore, we look into the relatively near term and do not address further term possibilities, such as AI agents enabling each individual to become a super-citizen on the one hand[1] or, on the other hand, Foucauldian total domination through hyper-surveillance by authoritarian state actors.

## Enhancing Political Capabilities: Into the Near Future

Our approach to thinking about the effects of artificial intelligence (AI) is inspired by the task-based conceptual framework that Daron Acemoglu and Pascual Restrepro developed for reasoning about the effects of technology on labor and production.[2] In this framework, a job is a bundle of tasks. Rather than trying to figure out whether a technology will eliminate (automate) a job, we should think about the tasks which the technology performs without labor (substitution), which tasks are left untouched, and which tasks might be done better with a combination of labor and technology (augmentation). In the analysis below, we reason about the tasks that can be done

[1] See Andy Hall, "Building Political Superintelligence" substack (March 26, 2026). (Accessed on March 28, 2026): URL: https://substack.com/home/post/p-192214946

[2] Acemoglu, Daron, and Pascual Restrepo. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33, no. 2 (May 1, 2019): 3–30. https://doi.org/10.1257/jep.33.2.3.

with, or done better, with generative AI, in political campaigns, election administration, social movements, and citizen deliberation.

But what does "done better" with generative AI mean? The core capability of generative AI is to facilitate the manipulation of information in the form of text, images, and videos through actions such as searching, summarizing, translating, recombining, and generating. Humans also do these things; but, as Schneier and Sanders write, generative AI may help humans do them better, or do them better without humans, in four ways: by increasing speed, scale, scope, sophistication – or all of these.[3]

While enhancing the quality of election administration seems unambiguously good for democracy, increasing the effectiveness ("doing better") of the other three kinds of political activities – campaigns, social movement, and citizen deliberation – could be good or bad for democracy. For example, enhancing the capabilities of political campaigns to win by misleading or mobilizing citizens' darker impulses may harm the quality of democracy. Amplifying the capabilities of ethnonationalist or illiberal social movements may also harm democracy.

## Political Campaigns

Generative AI is a new technology, and its use by political campaigns is in its infancy. Nevertheless, political campaigns in countries as diverse as India, South Africa, Mexico, the United Kingdom, and the United States already utilize large language models.[4]

People who work in political campaigns devote enormous amounts of labor to producing campaign materials for email, postal, and video platforms, fundraising proposals, and other written and video communication. Generative artificial intelligence augments capabilities of campaign workers to produce a greater variety of potentially higher quality text, audio, and visual materials much more quickly. These capabilities are especially valuable to down-ballot, resource strapped, campaigns. For example, Democratic strategist Mike Nellis has founded Quiller, a startup that harnesses generative AI to produce drafts of fundraising emails for campaign workers to tailor and improve.[5]

---

[3] Schneier, Bruce, and Nathan E. Sanders. *Rewiring Democracy: How AI Will Transform Our Politics, Government, and Citizenship.* Cambridge, Massachusetts London: The MIT Press, 2025.
[4] Foos, Florian. "The Use of AI by Election Campaigns." *LSE Public Policy Review* 3, no. 3 (November 4, 2024). https://doi.org/10.31389/lseppr.112.
[5] "How a Startup Is Using AI to Write Fundraising Emails." Accessed March 28, 2026. https://thehill.com/homenews/campaign/4146154-how-a-start-up-is-using-ai-to-write-fundraising-emails/.

Generative artificial intelligence already features high quality, cheap, and speedy voice and text language translation capabilities. In the 2024 Indian elections, Prime Minister Narendra Modi[6] used a generative AI tool called Bhashini to translate his campaign speeches[7] – delivered in Hindi – to Tamil and other regional languages in real time. When he was mayor of New York City, Eric Adams used a translation tool from ElevenLabs to send messages to some four million people through the city's robocall system in languages he does not speak, including Spanish, Yiddish, Mandarin, Cantonese, and Haitian Creole.[8] Critics accused Adams of using "deepfakes" through these calls because he did not disclose that it was not actually him speaking, while the mayor responded that it was an effective way to reach the city's diverse population. Partially in response to this episode, the New York State Legislature enacted legislation that requires any political communication in the state that uses AI to disclose that the message has been manipulated by artificial intelligence.[9]

While generative AI will continue to increase the efficiency of such "back-office" tasks as fundraising and language translation, it may have a more profound impact on political campaigns by increasing the scope and quality of a campaign's interactions with voters, relating to targeting, listening, persuasion, and mobilization. At the positive end of outcomes, generative AI can facilitate dialogue between campaigns and voters and make campaigns more responsive to voters' priorities.

**Message Development**

In the domain of targeted messaging, political consultants are developing technologies of "synthetic audiences," at large and small scale, that enable politicians to conduct polls and focus groups much more cheaply, and to test political messages and discover political values with greater speed, at scale, and with greater scope. Adjacent to the politics industry, *The Times* (of London) reported in 2025 that it engaged a company called Electric Twin to create a synthetic audience of the British media public using the newspaper's database of several hundred thousand

[6] "AI and Deepfakes Played a Big Role in India's Elections - New Lines Magazine." Accessed March 28, 2026. https://newlinesmag.com/spotlight/ai-and-deepfakes-played-a-big-role-in-indias-elections/.
[7] "AI Amplifies Political Reach but Magnifies Disinformation in India Elections." Accessed March 28, 2026. https://www.asiapacific.ca/publication/indian-election-use-of-ai-political-campaigns-voter-engagement.
[8] "AI and You: NYC Mayor Can't Really Speak Mandarin, the AI Money Trail, Who Sets the Rules - CNET." Accessed March 28, 2026. https://www.cnet.com/tech/computing/ai-and-you-nyc-mayor-cant-really-speak-mandarin-the-ai-money-trail-who-sets-the-rules/.
[9] New York State Senate, "Senate Bill S9678, 2023–2024 Legislative Session: Relates to materially deceptive media in political communications," accessed March 29, 2026, New York State Senate.

actual subscribers.[10] They used this model and synthetic audience to test marketing and communication strategies to, evidently, very good effect. Companies are starting to advertise similar synthetic audiences, focus groups, and polling services to political clients.

**Mobilization**

Two fundamental aims of political campaigns are to (i) mobilize people who are already inclined to support their candidates and (ii) persuade those who are not yet on board to become supporters. Get-Out-The-Vote (GOTV) campaigns are essential to a modern ground game. However, some GOTV contact modes that have been shown to be effective at increasing turnout, such as door-to-door canvassing and telephone calls,[11] [12] are also difficult to scale-up. Based on a meta-analysis of more than 700 GOTV interventions, Foos et al. find that the larger the sample size of a GOTV campaign, the smaller the size of its reported effect.[13] This association holds for interventions that contain an important human element in them, door-to-door canvassing and telephone calls, but not for GOTV interventions that are easy to scale-up, such as mail shots, SMS text messages, and social media messages, making it unlikely that the association stems predominantly from publication bias in smaller studies. The authors attribute this smaller effect size in larger GOTV interventions that rely on humans for delivery to the well-known organizational challenges of scaling-up social interventions, including the recruitment, training, and monitoring of human volunteers and workers, which is required to deliver an intervention at scale and at consistent quality.[14] [15] While human canvassers will be difficult to replace at the doorstep, randomized trials have shown that AI chatbots are effective at mobilizing voters.[16] Already, candidates' phone calls can be delivered in a conversational manner by AI voice agents, substituting for human telephone canvassers and enabling campaigns to contact voters at scale via telephone calls.[17] It is not hard to imagine that those calls will be as – or even more –

[10] "The Times Is Using Synthetic Research for Its Editorial Product Plans." Accessed March 28, 2026. https://digiday.com/media/how-the-times-is-using-ai-to-model-synthetic-focus-groups-from-human-audiences/.
[11] Green, Donald P., and Alan S. Gerber. *Get out the vote: How to increase voter turnout*. Brookings Institution Press, 2019.
[12] Issenberg, Sasha. *The victory lab: The secret science of winning campaigns*. Crown, 2013.
[13] Foos, Florian, Peter John, Francisco Tomas-Valiente and Asli Ceren Cinar. Scaling-Up Mobilization Campaigns. Chapter 4: Intervention Size, Manuscript under development. (Publication pending.)
[14] List, John A. The voltage effect: How to make good ideas great and great ideas scale. Currency, 2022.
[15] Al-Ubaydli, Omar, John A. List, and Dana Suskind. *The science of using science: Towards an understanding of the threats to scaling experiments*. No. w25848. National Bureau of Economic Research, 2019.
[16] Mann, Christopher B. "Can conversing with a computer increase turnout? Mobilization using chatbot communication." *Journal of Experimental Political Science* 8.1 (2021): 51-62.
[17] Greschler, Gabe. "Voters Are Being Called by Campaigns' AI Bots. I Talked to One for 10 Minutes." *San Francisco Standard*, February 27, 2024. https://sfstandard.com/2024/02/27/peter-dixon-congress-ai-robocalls-civox/.

effective as calls by humans and more effective than robo calls of the past. Another GOTV mode that can benefit from the involvement of AI chatbots are SMS text messages. For a long time, campaigns sent one shot text messages to voters to remind them to vote; but, more recently campaigns have introduced interactive elements to those text exchanges, that now predominantly take place via messenger apps. AI chatbots are likely to replace human volunteers in interactive text messaging; that shift is already evident in other industries for many customer service interactions.

**Persuasion**

Beyond mobilization, multiple recent studies found that conversations with LLM-powered chatbots can be politically persuasive – even more persuasive than traditional (non-AI) methods of political persuasion. Lin et al. (2025) conducted an experiment in which participants engaged in conversation with a chatbot that was instructed to persuade them to support a political candidate (e.g. Donald Trump or Kamala Harris) or a ballot measure (legalizing psychedelic drugs in Massachusetts).[18] They found that those exposed to the chatbot supported the bot's candidates and positions more than they did in the control conditions, at quite significant levels. The researchers conducted the experiment for multiple contexts: 2024 national elections in the United States, Canada, and Poland, as well as a ballot campaign in Massachusetts. Offer-Westort et al. (2026) presents suggestive evidence from a randomized trial, replicating the intervention in Kalla and Broockman (2016), that AI chatbots integrated into Facebook messenger were effective at reducing exclusionary attitudes among users.[19]

In a recent study published in *Science*, Hackenburg et al. (2025) tested the political persuasiveness of 19 different large language models.[20] Participants were assigned to chat with an LLM that was instructed to persuade them on an issue for two to ten rounds, while control participants read an LLM-generated static message. Both the static message and multi-round conversations produced statistically significant persuasion effects. Interestingly, large LLMs (more expensive) were more effective than smaller ones. However, "training" LLMs on conversations that were more persuasive increased their effectiveness more than model size,

[18] Lin, Hause, Gabriela Czarnek, Benjamin Lewis, et al. "Persuading Voters Using Human–Artificial Intelligence Dialogues." *Nature* 648, no. 8093 (2025): 394–401. https://doi.org/10.1038/s41586-025-09771-9.

[19] Offer-Westort, Molly, et al. "Deep canvassing with automated conversational agents: Personalized messaging to change attitudes." *Research & Politics* 13.1 (2026): 20531680251414926. https://www.semanticscholar.org/paper/Deep-canvassing-with-automated-conversational-to-Offer-Westort-Liu/91a56830097d287d1d5b34b65682a253fdaed6a8

[20] Hackenburg, Kobi, Ben M. Tappin, Luke Hewitt, et al. "The Levers of Political Persuasion with Conversational Artificial Intelligence." Science 390, no. 6777 (2025): eaea3884. https://doi.org/10.1126/science.aea3884.

suggesting that it may be possible to create low-cost persuasive LLMs that are accessible to less well-resourced political campaigns.

**Listening and dialogue**

Most political campaigns operate in a one-to-many, one-way mode: A candidate broadcasts their appeals to a mass audience of potential supporters and swing voters. Potential supporters sometimes have their say in smaller venues, such as house parties and town halls, and politicians also try to glean what voters want through surveys and focus groups. Over the last decade, listening has become an important task for campaigns in order to poll individual voters on political issues and to collect personal contact data. Listening campaigns involve human canvassers visiting voters at the doorstep or calling them on the phone. Listening campaigns often accompany the rollout of a campaign and have been applied widely, for instance by Emmanuel Macron's 2017 presidential campaign[21] and by Hillary Clinton in her 2016 presidential campaign.[22] While information still flows from the candidate to the audience in most campaigns, generative artificial intelligence applications may alter this dynamic by *increasing* the flow of information from constituents to politicians through listening and dialogue. AI can augment a political campaign's abilities on the listening side in two ways: It can help with the systematic analysis of text and voice responses, and it can conduct conversations at scale. While listening campaigns collect a lot of texts from citizens, they have struggled in the past to analyze these open text responses effectively. Large Language Models are easily able to identify issues, positions, and tone of open text or voice answers and provide campaigns with real-time means to aggregate voter sentiment and issue priorities.

Large Language Models combined with online political discussion platforms and apps also have the potential to enable campaigns and politicians to "listen" to citizens and constituents at scale, with increased sophistication, by engaging almost unlimited numbers of citizens in many-round discussions, identifying the most common themes and points of agreement and locating issues of majority or super-majority agreement.

[21] Reuters. "France's Macron Launches Door-to-Door Campaign as Presidency Battle Looms." Accessed March 28, 2026. https://www.reuters.com/article/world/frances-macron-launches-door-to-door-campaign-as-presidency-battle-looms-idUSKCN0YJ09E/.

[22] "For This Hillary Clinton Listening Tour, Few Hugs but Much Note-Taking" The New York Times. Accessed March 28, 2026. https://www.nytimes.com/2015/05/28/us/politics/hillary-clinton-tells-voters-shes-listening-and-learning.html.

"AI Steve" may have been the first LLM-powered politician to run for office.[23] "AI Steve" is the digital avatar of Steve Endacott who ran as an independent seeking to represent what was the Brighton Pavilion constituency in the United Kingdom Parliament in the 2024 General Election. Electoral authorities made clear that, if he won, Endacott – not "AI Steve" – would be the member of parliament. But Endacott proposed that the candidate, and the legislator, if elected, would be a cyborg combination of Endacott and "AI Steve." "AI Steve" was directed to have conversations on policy issues with people in Endacott's constituency. It was capable of carrying on 10,000 simultaneous conversations and was available 24/7. The avatar would record and analyze all conversations, identify positions with majority support, and then Endacott would advocate for those positions. The election took place in July 2024, and "AI Steve" lost, having received only 179 votes (less than one percent of the total votes cast).[24] But, as this experience illustrates, it seems likely that different kinds of human+AI chimeras will become more common in politics.

A more prominent recent experience of a human politician + AI chimera comes from Pakistan. In the election of early 2024, candidates who were aligned with the Pakistan Tehreek-e-Insaf (or PTI) party won the most votes for seats in parliament. Ruling authorities have imprisoned PTI leader Imran Khan since 2022. Khan's supporters used AI technologies to construct a video of Khan delivering a victory speech, and they disseminated it widely.[25] They also campaigned on behalf of Khan, using AI to construct speeches from notes that Kahn wrote from prison.

Dialog as a political activity is closely related to listening. Dialog represents a two-way exchange between politicians and constituents and may help politicians and constituents be more informed about one another and increase the confidence and trust that constituents have in their representatives.[26] But for candidates and representatives, dialog is time-consuming and expensive – so expensive that scholars who study money and politics have argued that campaign contributions "buy" access[27] or the opportunity to engage in dialogue with one's political representative.

Applications of generative AI have the potential to decrease dramatically the cost of dialog by enabling constituents to engage in conversations with virtual LLM-powered avatars that are

[23] "There's an AI Candidate Running for Parliament in the UK." in *Wired*. Accessed March 28, 2026. https://www.wired.com/story/ai-candidate-running-for-parliament-uk/.
[24] "Election Result for Brighton Pavilion (Constituency) - MPs and Lords - UK Parliament." Accessed March 28, 2026. https://members.parliament.uk/constituency/3945/election/422.
[25] "Imran Khan Uses A.I. To Give Victory Speech in Pakistan" *The New York Times*. Accessed March 28, 2026. https://www.nytimes.com/2024/02/11/world/asia/imran-khan-artificial-intelligence-pakistan.html.
[26] Neblo, Michael A., Kevin M. Esterling, and David Lazer. *Politics With the People: Building A Directly Representative Democracy*. Cambridge, United Kingdom ; New York, NY: Cambridge University Press, 2018.
[27] Kalla, Joshua L., and David E. Broockman. "Campaign Contributions Facilitate Access to Congressional Officials: A Randomized Field Experiment." *American Journal of Political Science* 60, no. 3 (July 2016): 545–58. https://doi.org/10.1111/ajps.12180.

capable of answering questions about the politician and their positions and perspectives. Takahiro Anno used these technologies in his bid to become Tokyo's mayor (known as governor). Schneier and Sanders write that "Anno, a software engineer by trade, invented new political technologies and leveraged AI to amplify his individual capacity to listen. He used an AI avatar trained on his political manifesto to respond to 8,600 questions from voters over a seventeen-day continuous livestream."[28] Though he was an outsider candidate and lost the mayoral race in 2024, he ran for a seat in the upper house of Japan's National Diet in 2025 and won. In the course of his campaigns, he founded a new political party called the "Future Party" that promises to utilize AI and other technologies to usher in new methods of digital democracy.

Dean Phillips, a long-shot candidate for U.S. president in 2024, also utilized an LLM-powered avatar; he called it "Dean.Bot." His supporters commissioned a startup called Delphi to build the bot using ChatGPT. They reportedly trained the Dean.Bot on the candidate's speeches, podcasts, and interviews – totaling some 240,000 words.[29] Users could ask "Dean.Bot" questions verbally and receive spoken responses in a voice that sounded like Phillips, but "Dean.Bot" was a very short-lived effort. It went live in early January 2024, but OpenAI – the company that runs ChatGPT – shut the effort down by late January 2024 because their terms of use prohibit political campaigning.[30]

**Citizens' private use of AI chatbots**

One under-researched area is citizens' private use of AI chatbots to obtain political information. Citizens are increasingly using AI chatbots in their everyday life to obtain information about everyday issues, including politics and elections. These interactions with AI chatbots include users fact-checking political claims that friends are making in conversation, querying the political positions of candidates and parties, and asking about political events and practical voting matters. As more citizens rely on chatbots and other forms of AI to search for political information, the effects on the quality of the information they receive, and thus the information environment for democracy, strikes us as radically underdetermined. On one hand, AI can provide highly relevant information in responsive and accessible ways – more so than a web search or perhaps even newspaper and magazine articles. On the other hand, we do not know whether this shift will exacerbate the polarization of information spheres or whether it will aid or

[28] Schneier, Bruce, and Nathan Sanders. "Rewiring Democracy Now." Accessed March 28, 2026. https://therenovator.substack.com/p/rewiring-democracy-now.
[29] "Silicon Valley Insiders Are Trying to Unseat Biden with Help from AI - The Washington Post." Accessed March 28, 2026. https://www.washingtonpost.com/elections/2024/01/18/ai-tech-biden/.
[30] "OpenAI Bans Developer of Dean Phillips Bot - The Washington Post." Accessed March 28, 2026. https://www.washingtonpost.com/technology/2024/01/20/openai-dean-phillips-ban-chatgpt/.

hamper truth seeking efforts. Those future patterns depend in part on choices and behaviors of individuals, political leaders, AI companies, and others. Will individuals seek out biased information to confirm or disconfirm their views? Will the AI that companies build provide accurate information and how will they address viewpoint diversity?[31] In the provision of information through AI, how will companies balance truth-seeking versus user engagement? For example, will profit motives lead companies to push information to individuals based upon how much campaigns, politicians, and economic actors are willing to pay? The answers to these questions will be determined not in the first instance by the arc of technological development but rather by political, economic, and regulatory decisions.

**AI and personal data regulation**

The scale at which campaigns can use generative AI in their campaigning depends on the willingness, and the ability, of countries and other political entities to regulate campaigns' use of generative AI systems[32] and their ability to harvest and share citizens' personal data legally.[33] While the United States has so far abstained from regulating the political use of generative AI at the federal level and have comparatively permissive data privacy and data sharing laws, the European Union has taken concrete steps to regulate the use of AI in elections via the 2024 EU AI Act. Classifying AI systems intended to influence voting behavior or electoral outcomes as higher-risk, the AI Act imposes conditions for AI-enabled campaign interventions. Generative AI outputs need to be marked as AI-generated in a machine-readable format, and, in some cases, need to be visibly labeled. Similar regulations, which require campaigns to disclose whether content was produced by generative AI and if voters are interacting with AI systems, have, as of 2026, been adopted by twenty US states.[34] California has taken a particularly strict approach to regulating the use of AI in election campaigns. The European Union has already previously, in 2016, restricted organizations' ability to harvest and share personal data via the General Data Protection Regulation (GDPR), which regulates the collection, processing and sharing of personal data. Obtaining personal contact details at scale, and particularly telephone numbers, is a necessary requirement for AI-driven GOTV and persuasion tactics that use phone or SMS text

[31] Westwood, Sean, Justin Grimmer, and Andrew Hall. "Measuring Perceived Slant in Large Language Models Through User Evaluations." Working Paper, May 8, 2025. https://www.gsb.stanford.edu/faculty-research/working-papers/measuring-perceived-slant-large-language-models-through-user.

[32] Foos, Florian, Peter John, Francisco Valiente-Jorda and Asli Ceren Cinar. Scaling-Up Political Campaigns. Manuscript under development.

[33] Foos, Florian. "The Use of AI by Election Campaigns." *LSE Public Policy Review* 3, no. 3 (November 4, 2024). https://doi.org/10.31389/lseppr.112.

[34] https://ai-law-center.orrick.com/us-ai-law-tracker-see-all-states/

messages. Richer data, such as individuals' consumption and on-line behavior, are also important for targeted and tailored persuasion efforts.[35] These regulations, hence, directly pertain to the scalability of GOTV interventions where AI-based systems are used to replace human volunteers or to micro-target voters. AI-based gains in scalability will therefore, at least to some extent, depend on favorable national and supranational AI and data regulation environments. There is evidence that voters are in favor of stricter AI regulation if presented with campaigns' use of generative AI that could be labeled as deceptive, but that they are unlikely to punish parties for using these tactics.[36] This speaks to the incentives that different types of political actors face when deciding whether or not to deploy generative AI legally in voter contact, and illegally, for instance via the use of deepfakes or by not labelling content in line with regulations. Besides the regulatory environment, as Foos (2024) argues, smaller parties that have no membership base and are less bound by established political norms and inertia, should have a larger incentive to rely on AI-bots in their campaigning than mainstream parties.[37]

## Election Administration

Like many organizations whose work involves communicating, interacting with the public, and responding to requests for information, local election officials (LEOs) have the potential, with the use of artificial intelligence, to operate more effectively, with less effort and fewer resources. This is especially important because election administration offices are often highly strapped for resources, particularly in the United States. At the same time, perhaps because many LEOs are operationally conservative, few have begun to utilize the novel technology of generative artificial intelligence in any significant way. According to surveys conducted by the Brennan Center, just eight percent of local election officials[38] had used AI in their offices by 2025 (up from five percent in 2024).[39] The most common uses were drafting social media content and press releases, creating graphics, and language translation.

---

[35] Fung, Archon, and Lawrence Lessig. "How AI Could Take over Elections – and Undermine Democracy." The Conversation, June 2, 2023. http://theconversation.com/how-ai-could-take-over-elections-and-undermine-democracy-206051.

[36] Jungherr, Andreas, Adrian Rauchfleisch, and Alexander Wuttke. "Artificial intelligence in election campaigns: Perceptions, penalties, and implications." *Political Communication* (2026): 1-22.

[37] Foos, Florian. "The use of AI by election campaigns." *LSE Public Policy Review* 3.3 (2024).

[38] "Local Election Officials Survey — July 2025 | Brennan Center for Justice." Accessed March 28, 2026. https://www.brennancenter.org/our-work/research-reports/local-election-officials-survey-july-2025. Also see National Conference of State Legislatures, "Artificial Intelligence in Election Administration," Accessed March 29, 2026, https://www.ncsl.org/elections-and-campaigns/artificial-intelligence-in-election-administration.

[39] "Local Election Officials Survey — May 2024 | Brennan Center for Justice." Accessed March 28, 2026. https://www.brennancenter.org/our-work/research-reports/local-election-officials-survey-may-2024.

Like people working on political campaigns, local election officials must produce many different kinds of materials to inform voters: educational materials about how to register to vote; where and when to vote; voter ballot guides about offices, candidates, and issues; "I voted" stickers and other civic "swag;" and, materials about the process of validating and counting votes. The latter is increasingly needed to defend against claims of voter fraud and other forms of election malfeasance. In addition, LEOs are responsible for producing training materials that translate state-mandated requirements about voting procedures for the hundreds of thousands of workers who staff polling places in each election. Generative AI can greatly speed the work of producing such materials, though election officials must be careful to edit and verify the accuracy of final versions lest they unwittingly propagate misinformation that disenfranchises voters or erodes confidence in the election machinery.

Generative artificial intelligence could also be used to respond to the many routine requests that election officials receive. These include information requests from voters and other residents: Am I eligible to vote? How do I register to vote? Where should I be registered? How and where do I vote? How do I become a poll worker? Others, if they are concerned about the integrity of how the election process is conducted, seek information about the processes and personnel involved in election administration. And election officials report increases in records requests that have become a substantial burden in some jurisdictions.[40]

As in many other industries that have demanding customer support requirements, some election officials have created chatbots and other technologies that interact via voice or text or a website to respond to high-volume and frequent inquiries from citizens. Many jurisdictions offer web forms that enable residents to check on their voter registration status and correct errors. Idaho election officials, for example, work with IBM to use its Watson technology to create and deploy a chatbot to respond to voter inquiries about, for example, where to vote or how to acquire an absentee ballot.[41] At this writing, publicly documented examples remain limited, and election offices that deploy AI chatbots appear to favor retrieval-augmented systems grounded in official materials.[42]Many election officials are concerned that LLMs would produce hallucinations and other errors.

In a 2024 Columbia University workshop, participants tested the voting information veracity of several unmodified commercial chatbots, including OpenAI's ChatGPT-4, Meta's Llama 2,

[40] "Election Officials & the Misuse of Public Records Requests | The Center for Election Innovation & Research." Accessed March 28, 2026. https://electioninnovation.org/research/public-records-requests/.
[41] "Election Questions? Idaho's New Watson-Powered Chatbot Has Answers | StateScoop." Accessed March 28, 2026. https://statescoop.com/election-chatbot-idaho-watson-ai/.
[42] National Conference of State Legislatures, "Artificial Intelligence in Election Administration," accessedMarch 24, 2026, https://www.ncsl.org/elections-and-campaigns/artificial-intelligence-in-election-administration

Google's Gemini, Anthropic's Claude, and Mistral.[43] Participants rated half of the chatbot responses as inaccurate and 40 percent as harmful. For example, according to PBS News (2024), "In Nevada, where same-day voter registration has been allowed since 2019, four of the five chatbots tested wrongly asserted that voters would be blocked from registering to vote weeks before Election Day."

The current common wisdom in the election administration field is to refrain from using generative technologies to automate responses to requests for information and rely instead on rule-based systems and other such technologies.

In recent years, a different kind of information query has become more common in election administration: the public records request. Local election officials have fielded an increasing number of public records requests since the battles over the integrity of the 2020 US presidential election.[44] Whether motives for these requisitions stem from misinformation, malicious intent to paralyze election machinery, or standard inquiries from academics and reporters, these requests consume enormous amounts of time and resources in some jurisdictions. The Center for Election Innovation and Research reports that:

> …in Wake County, North Carolina, officials reported a sevenfold increase in the number of election-related records requests received from 2020 to 2022, leading the county to increase its budget to hire a new staff member dedicated to processing incoming requests. In 2022, officials in Maricopa County, Arizona, reported one request that required nearly half the election office's staff to spend four days sorting and scanning 20,000 documents… Local election officials in states like Florida and Michigan have reported spending anywhere from 25% to over 70% of their time just on processing public records requests in recent years. Officials in states like Virginia have complained that they have been unable to engage in routine tasks like voter outreach due to the volume of requests received. (Yoder & Tan, 2024, p. 2)

The proliferation of generative AI may exacerbate this problem by making it easier for citizens to make such burdensome records requests. On the other hand, LLMs may make those requests less burdensome by enabling election officials to respond more quickly and with less effort. Though we could not find cases of local election officials who have built and used generative AI

[43] "Chatbot Info on U.S. Elections Is Inaccurate, Misleading and Could Keep Voters from Polls, Report Finds | PBS News." Accessed March 28, 2026. https://www.pbs.org/newshour/politics/chatbot-info-on-u-s-elections-is-inaccurate-misleading-and-could-keep-voters-from-polls-report-finds.

[44] "Trump Backers Inundate Election Offices with Requests for 2020 Records - The Washington Post." Accessed March 28, 2026. https://www.washingtonpost.com/nation/2022/09/11/trump-election-deniers-voting/.

to speed records requests, there are promising examples from other governments. For example, a team in the US State Department trained an AI on diplomatic documents in order to see whether it could accurately classify documents that should be declassified.[45] The classifier turned out to be highly accurate, and the team expanded the project to respond to Freedom of Information requests by identifying the documents that should be returned, based in part on similar requests in the past.

There are several high-risk areas of election administration that could be sped up through partial AI automation. But, at this stage in the development of generative AI – and other forms of artificial intelligence – mistakes are harmful (e.g. could cause disenfranchisement). Such tools should be used experimentally and with human oversight. For example, many jurisdictions validate mail-in ballots by matching handwritten signatures. Many jurisdictions utilize pattern discriminating AI (not generative AI) as a first pass to flag questionable signatures for human assessment called Automatic Signature Verification.[46]

Generative and other forms of artificial intelligence have the potential to improve the operations of too-often under-resourced agencies of local election officials. Some of the tasks that election officials must do can be enhanced by consumer level generative artificial intelligence – chatbots such as OpenAI's ChatGPT and Google's Gemini. These chatbots, for example, are well suited to producing drafts of press releases, informational brochures, and poll worker training materials. They are also already quite capable of translating election materials into different languages. We expect generative AI to proliferate in these domains. Other tasks require generative AI applications that are customized – "trained" – with specialized information and domain specific materials. For example, chatbots that provide information to voters will need to accurately convey the specific rules and provisions for particular states, and even counties and cities. The spread of these generative AI applications depends in large part on supply side questions: Will technology vendors develop capable applications of these technologies? That depends, in turn, on whether local election officials will have access to resources sufficient to purchase such applications and so incentivize technology firms to develop them.

[45] "Leadership Program Inspires an AI Revolution at the State Department • Partnership for Public Service." Accessed March 28, 2026. https://ourpublicservice.org/about/impact/leadership-program-inspires-an-ai-revolution-at-the-state-department/.
[46] Brennan Center. 2023. "Safeguards for Using Artificial Intelligence in Election Administration | Brennan Center for Justice." URL: https://www.brennancenter.org/our-work/research-reports/safeguards-using-artificial-intelligence-election-administration. Accessed on April 11, 2026; Electoral Assistance Commission. "Signature Verification and Cure Process" URL: https://www.eac.gov/sites/default/files/electionofficials/vbm/Signature_Verification_Cure_Process.pdf. Accessed on April 11, 2026.

From a larger vantage, well functioning democracy depends on systems of election administration that are both trusted and trustworthy.[47] Generative AI may help to increase trustworthiness by increasing the efficiency, reliability, and service quality of these agencies. It may increase trust by enabling local election officials to communicate more effectively with voters and media audiences in their jurisdictions. But on the other side, skeptics and critics of election administrators also have access to generative artificial intelligence tools that enhance their capabilities to communicate and persuade.

## Social Movements

Artificial intelligence creates great challenges and potential opportunities for civil society organizations and social movements. Civil society organizations are struggling to constrain and mitigate the negative effects of artificial intelligence on many fronts: algorithmic bias, privacy protection and surveillance, preventing and mitigating harms to people and the environment that stem from AI training and the construction and operation of data centers, and so on. On the other hand, artificial intelligence can help to make social and democracy movements more effective and, potentially, reduce power asymmetries between movement organizations and the larger, better resourced governments and private organizations against which they often struggle.

Like political campaigns, social movements depend upon building popular support and raising funds in order to grow and accomplish their goals. Public communication – in person, through mass media, and in social media; through spoken word, audio, images, video, etc. – is a principal method to build such support. As in many other fields, people in social movements can use generative AI technologies to produce text, images, and videos with greater speed and sophistication.

Freddy Guevara is a democratic activist and a leader of the Venezuelan opposition in exile. He writes, "I use AI almost daily in my strategizing, activism, and research.… you don't need advanced skills – just access to a trusted LLM and basic knowledge for asking the proper questions and creating the right prompts."[48] Creating well-produced content – especially images and videos – can require expensive professional services whose costs are out of reach for many

[47] Charles Stewart III, "Trust in Elections," *Daedalus* 151, no. 4 (2022): 234–53, https://doi.org/10.1162/daed_a_01953

[48] Guevara Freddy, "Using AI Now to Improve Movements' Effectiveness: A Basic Introduction for Social Activists" Commentary, Ash Center for Democratic Governance and Innovation. (March 4, 2025) URL (accessed on January 17, 2026): https://ash.harvard.edu/articles/using-ai-now-to-improve-movements-effectiveness-a-basic-introduction-guide-for-social-activists/

social movement organizations and activists (but readily available to governments and private corporations). But generative AI tools – such as DALLE and Midjourney for images and Synthesis and Sora for video – make it possible to produce sophisticated content for much less money, in much less time.

An early defining moment in social movements' use of AI-generated content was the quick and somewhat controversial spread of an image called "All Eyes on Rafah" in 2024. On May 26, 2024, Israeli military raided a large displacement camp, holding some one million people, in the city of Rafah in the southern Gaza Strip. The raid killed several dozen people.[49] In the immediate aftermath, eyewitnesses shared photos of the devastation on social media, but platforms' content moderation algorithms (especially Instagram) blocked the images because they were too graphic.

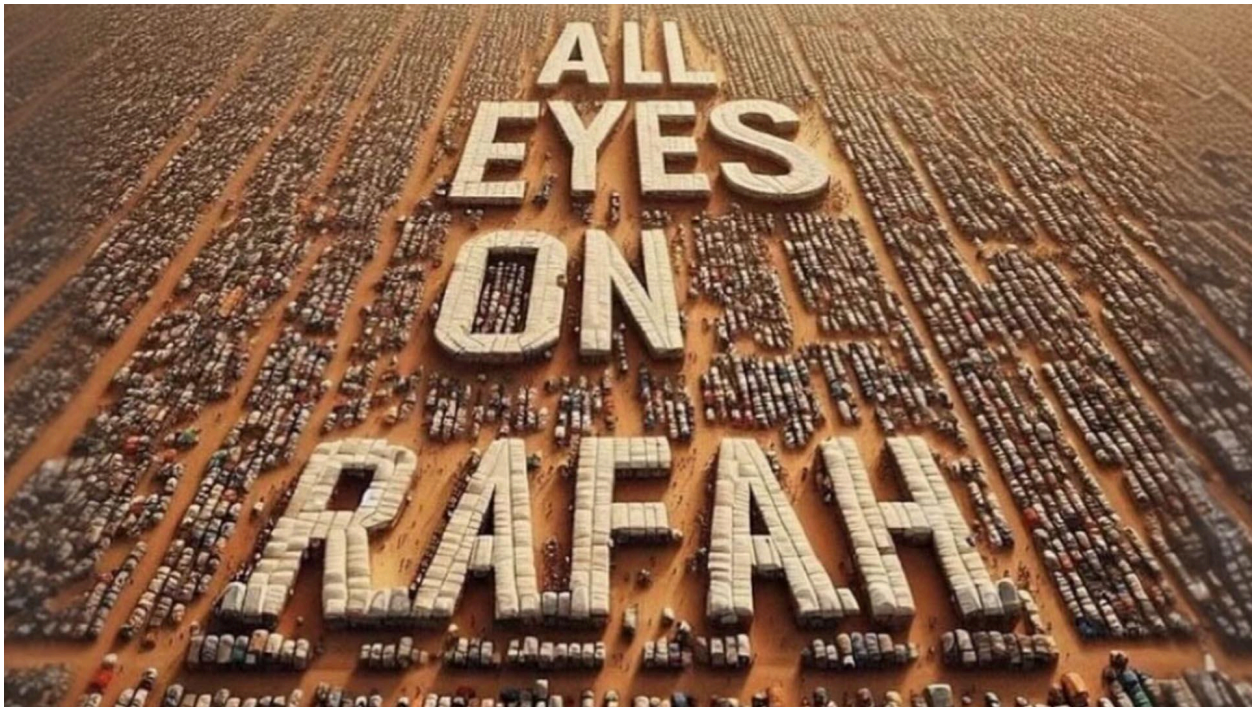


Figure 1.1. "All Eyes on Rafah" image

Soon after, an image now known as "All Eyes on Rafah" began to circulate on Instagram and other platforms. The AI-generated image depicted a sea of tents with large white tents spelling "All Eyes on Rafah" (See Figure 1.1 above). At its peak in late May 2024, the image had been

[49] Burgess, Annika and Morris-Grant, Brianna. 28 May 2024. "As the world reacts to Israel's deadly air strike on Rafah, here are key events from the Gaza war leading up to the attack" ABC News. URL (accessed on January 17, 2024): https://www.abc.net.au/news/2024-05-28/israel-gaza-war-time-line-before-fatal-rafah-air-strike/103901144

shared by between  47 million to 50 million users on Instagram. It also spread widely on X and TikTok. Its spread was fueled in part by celebrities, such as Bella Hadid, Dua Lipa, Pedro Pascal, and Alia Bhatt (a major Bollywood star).[50] The image spurred several controversies. First, two different Malaysian individuals claimed credit for creating the image.[51] Critics complained that the image sanitized the war by omitting the blood and carcasses in favor of clean rows of tents. The image also overshadowed, and perhaps thereby de-valued, the work of actual journalists who documented the attack.

The Centre for Responsible Union AI conducts training programs for hundreds of union staff and members in the United Kingdom, Europe, Australia, and the United States about how to utilize artificial intelligence tools. In a survey of their program participants, members wanted AI largely to help them with back office tasks, the mundane rather than the revolutionary: writing, managing email and meetings, research and analysis, training, and member engagement.[52] For example, Belgian ABVV-Metaal created an AI-backed process to help union staff write letters (all containing standard information) requesting leave for members.[53]

One emerging bespoke use of AI by labor unions – using generative AI, natural language processing, and rule-based systems – is to create chatbots running through apps and messaging platforms.These chatbots help union members understand collective bargaining agreements and their rights and prerogatives under law. They also provide a means to file grievances, to report misconduct or workplace violations, and to recruit new members. For example, the UK Public and Commercial Services Union (PCS) created a chatbot to help train and simulate recruiting conversations for existing members to use in recruiting new members.[54] A large teachers' union

[50] Burga, Solcyre. May 29, 2024. "Why the 'All Eyes on Rafah' AI Post Is Going Viral on Social Media" *Time*. URL (accessed on January 17, 2026): https://time.com/6983344/all-eyes-on-rafah-ai-post-viral/; Dunworth, Liberty. 29 May 2024 "Dua Lipa shares support for #AllEyesOnRafah: "Burning children alive can never be justified"" URL (accessed on January 17, 2026): https://www.nme.com/news/music/dua-lipa-shares-support-for-all-eyes-on-rafah-3760399; ABP News Bureau. 29 May 2024 "Alia Bhatt, Kareena Kapoor And Other Celebs Support Gaza As "All Eyes On Rafah" Trends After Israel's Attack" URL (accessed on January 17, 2026): https://news.abplive.com/entertainment/celebrities/all-eyes-on-rafah-indian-celebs-alia-bhatt-kareena-kapoor-priyanka-chopra-support-gaza-after-israel-attack-on-rafah-1691055

[51] Allyn, Bobby. June 3, 2024. "All eyes on Rafah' is the Internet's most viral AI image. Two artists are claiming credit" in New England Public Media. URL (accessed on January 18, 2026):https://www.nepm.org/national-world-news/2024-06-03/all-eyes-on-rafah-is-the-internets-most-viral-ai-image-two-artists-are-claiming-credit

[52] Centre for Responsible Union AI. "Beyond the Hype: What Union Staff Actually Want from AI." Center for Responsible Union AI. Accessed January 18, 2026. https://www.agileunions.ai/p/what-union-staff-actually-want-from-ai.

[53] "How ABVV-Metaal Used AI to Automate Creating and Sending Union Leave Letters for Reps." Accessed March 28, 2026. https://www.agileunions.ai/p/automate-creating-sending-union-leave-letters.

[54] "How PCS Created an AI Chatbot to Support Unions Reps Practice Conversations to Recruit New Members." Accessed March 28, 2026. https://www.agileunions.ai/p/case-study-repcoach-pcs-union-reps-practice-recruitment-conversations.

in Northern Ireland and Scotland (NASUWT) created a chatbot trained on its large internal library to help members deal effectively with many different kinds of challenging workplace situations.[55]

In addition to crafting communications to amplify social movement framing of their narratives[56] and accelerating "back office" tasks, movement activists and organizations have utilized AI (both generative and rule-based) to help recruit supporters and engage them in political activities.

An important example of such engagement comes from the 2024 Venezuelan elections. In that context, government authorities could not be counted on to run a free and fair election and to provide ample information and opportunities for Venezuelans to participate. Responding to this hostile electoral environment, the democratic opposition movement sought to create in-effect a parallel election apparatus that would, among other things, inform voters about where and how to vote, enable members of the public to report election improprieties, and – amazingly – create a parallel vote-counting apparatus to challenge what turned out to be fraudulent official results.[57]

Venezuelan social movement and political opposition activists worked with technologists to develop several chatbots to provide information and engage people in the 2024 elections.[58] They developed Cheo, a natural language chatbot operating on top of Facebook and Instagram and then later WhatsApp. Cheo provided real time information about polling locations – which can be difficult to ascertain in Venezuela because there are more than 3,000, often in people's homes, and can shift rapidly. Activists developed capacities for Cheo to provide poll worker training for individuals and to help them get certified. Some 158,000 users signed up and interacted with the poll-worker training bot.

Activists also developed a Toma El Control (TEC) chatbot to enable observers to report incidents of suspected election malfeasance, complete with geocodes and the ability to upload photographic and video documentation. Researcher and participant Isabella Picón reported that "Over 200,000 messages were exchanged, receiving 88,885 and sending 112,080 responses with instructions to address interference at the polls."

---

[55] "How NASUWT Used an AI Chatbot to Make Their Knowledge Base More Accessible to Members." Accessed March 28, 2026. https://www.agileunions.ai/p/case-study-nasuwt-chatbot-members-resignation-notice-periods.
[56] Benford, Robert D., and David A. Snow. "Framing Processes And Social Movements: An Overview And Assessment." *Annual Review of Sociology*, 2000, 611–39.
[57] Schmidt, Samantha, Steven Rich, Ana Vanessa Herrero, and María Luisa Paúl. "Maduro Lost Election, Tallies Collected by Venezuela's Opposition Show." *The Washington Post*, August 4, 2024. https://www.washingtonpost.com/world/2024/08/04/maduro-gonzalez-election-actas-analysis/.
[58] Picón, Isabella. March 6, 2025. "AI-powered bots and electoral participation: a view from the Venezuelan experience" Ash Center for Democratic Governance and Innovation. URL (accessed on January 17, 2026): https://ash.harvard.edu/articles/ai-powered-bots-and-electoral-participation-a-view-from-the-venezuelan-experience/

A third broad domain in which artificial intelligence can boost the efforts of social movements is research. Activists engage in research for many purposes: to root out corruption; to develop better advocacy and campaign strategies; and to identify government abuses and rights violations. AI tools can augment the capabilities of human researchers and analysts in social movements by speeding up research activities and, through dramatic speed increases, reach scales that are not possible for small organizations to achieve with human effort alone.

Some efforts to identify corruption, bias, and abuse involve analyzing reams (or gigabytes) of government records. The Global Investigative Journalism Network, for example, has developed a method to identify and expose official gender based discrimination and violence. The method involves using ChatGPT to scan many thousands of pages of court records to identify suspicious cases and incidents that merit further investigation and can constitute broader patterns of abuse.[59]

In another example, Nathan Sanders developed tools that use generative AI to analyze oral and written testimony in state legislatures in the United States. He worked with an organization called Health Resources Action to utilize the tool to analyze more than a thousand oral testimonies in state legislatures to track emergent legislation. He used it also to fact-check testimony and characterize the values and ideological orientations appearing in this vast body of public testimony.[60]

Will social movements be able to harness the potential of artificial intelligence to increase the speed, scale, scope, and sophistication of their activities – and hence their effectiveness and power – in these and other ways? They face distinctive challenges when it comes to harnessing the emerging technologies of artificial intelligence. While some of the tasks of activists can be enhanced through off-the-shelf consumer technologies like chatbots, other tasks require tailored tools – such as the large-scale analysis of government documents or bespoke chatbots to engage supporters in elections or protests.

Whereas political campaigns and election officials are likely to rely on commercial vendors to develop custom AI applications and tools, social movements may be more reliant on networks of movement-friendly developers who provide open-source tools.[61] First, many social movements and civil society organizations face even tighter resource constraints than political campaigns or election administrators. Thus, an ecosystem of vendors providing high quality tools that these

[59] Ulrich, Sarah. February 3, 2025. "Using Data to Expose Systemic Gender-Based Violence — in 10 Steps". Global Investigative Journalism Network. URL (accessed on January 17, 2026): https://gijn.org/stories/10-steps-data-expose-gender-violence-eswatini/. Discussed in Sanders and Schneier. Rewiring Democracy.
[60] Sanders and Schneier. 2025. Rewiring Democracy, chapter 35 "Watching the Government."
[61] Jason Jie. March 2025. "Unlocking AI's potential for social movements," Ash Center for Democratic Governance and Innovation. URL (Accessed on January 17, 2026): https://ash.harvard.edu/articles/unlocking-ais-potential-for-social-movement/

organizations can afford is less likely to develop. Second, in the face of hostile governments with formidable surveillance capabilities, many movement organizations place a high premium on privacy and data security. These activists may not trust private sector vendors who sell applications and tools based on code that cannot be inspected, audited, or easily changed.

## Citizen Deliberation

Practices of citizen deliberation are as old as democracy itself: from discourse in the Athenian agora to town meetings in New England and open meetings today. In the contemporary political context, many see increased citizen participation in politics and public affairs as a way to strengthen democratic practices at a time when trust and confidence in democratic institutions and political leaders are in steep decline. At the same time, public involvement beyond casting a vote has long faced important practical limits – around scale, participation, and inclusion. There is growing optimism that some of these limits may be alleviated by the careful use of artificial intelligence.

This section illustrates how AI for citizens participating in the democratic process is best understood – not as a single intervention but – as a diverse toolkit for addressing recurring bottlenecks in deliberative processes. Different AI technologies can play different roles at different points in the democratic pipeline: facilitating registration, summarizing data, translating legal terms and election language, clustering, or mediation. In doing so, they potentially improve some aspects of enabling voter participation and entirely new deliberative formats while introducing distinct tradeoffs. The key question is often not *whether* "AI helps deliberation" in general but *which AI tool* is best suited to address a given need, how does it shift discretion and control within the process, and what forms of accountability remain.

We begin by outlining four recurring bottlenecks in citizen deliberation. Then, we present three case studies, which offer examples of how AI tools have been deployed toward alleviating some of these bottlenecks. Finally, we synthesize across these case studies to unpack the key opportunities and risks associated with using AI to support citizen participation in democracy.

### Four Key Bottlenecks

One of the most widely-discussed challenges faced by direct public involvement in the political process is its potential to ***scale,*** i.e., to include many people as participants. In *After the Revolution*, Robert Dahl describes how a town meeting cannot remain both inclusive and

participatory as the number of citizens attending any given opportunity grows. In many, if not most public meetings, only a small fraction of attendees can speak, or else a routine meeting or special hearing runs impractically long. Either way, "the greater the numbers, the more the town meeting runs the risk of becoming unrepresentative."[62] In addition to these logistical tensions, a major barrier to scaling is the need for high-quality facilitation: face-to-face deliberations often feature trained facilitators to keep discussions focused and make sure that participants feel included and engaged. Many public meetings (town halls, public hearings) lack sophisticated facilitation and, consequently, suffer from excess conflict, polarization, and the lack of meaningful exchange. Of note, non-technical solutions have already made important progress on the scale problem: Since the 1990s, organizations such as AmericaSpeaks[63] have created methods such as the 21st Century Town Meeting[64] to conduct face-to-face public deliberation at scale. Its two-level deliberations – in which many groups of 8-10 people simultaneously discuss a single topic (e.g., a city's budget) – allow several thousand individuals to participate in a single deliberative event.

A second significant barrier to direct public deliberation is ***sensemaking***: With so many ideas being exchanged and refined among many people, a core challenge is how to synthesize this information, suggest avenues of common ground, and aggregate ideas into intermediate or final conclusions. Traditionally, the former two tasks are done by participants and facilitators. For example, methods like the 21st Century Town Meeting rely on humans to read and analyze notes from deliberations in real time and spot emerging themes for subsequent deliberative consideration and group convergence. Similarly, aggregation is often done by asking participants to vote on resolutions or by polling their agreement with relevant statements. These methods can be dissatisfying because they can obscure the reasons why participants support the positions they do, how and why their views change over the course of public deliberations, and areas and opportunities for expanding agreement or minimizing excess conflict.

The third key bottleneck in citizen deliberation is ***unequal participation and inclusion***. This relates to Dahl's worry about whether someone can meaningfully impact – and thus be meaningfully included in – a large conversation. It also encompasses many other issues, including limited and unequal participation rates in synchronous deliberative processes like citizens' assemblies,[65] and discursive and circumstantial differences (e.g., differences in language

[62] Dahl, Robert Alan. *After the Revolution? Authority in a Good Society*. Yale University Press, 1990.
[63] See article on AmericaSpeaks here: https://participedia.net/organization/americaspeaks (Accessed on March 26, 2026)
[64] See article on the 21st Century Town meeting here: https://participedia.net/method/145 (Accessed on March 26, 2026)
[65] Flanigan, B., Gölz, P., Gupta, A., Hennig, B., & Procaccia, A. D. (2021). Fair algorithms for selecting citizens' assemblies. *Nature*, *596* (7873), 548-552.

or lived experience) that mean that people are best able to contribute to the "deliberative sphere" in different deliberative formats.

Deliberative forums also face a fourth bottleneck around ***trust and perceived legitimacy***, in both the general concept of deliberative forums and in specific deliberative events. Public trust in deliberative forums is not automatic, and it does not appear to depend only on whether deliberation occurred; rather, it is shaped by factors such as who is perceived to have participated,[66] whether the process is understandable to non-participants, and whether decision makers visibly respond to the resulting recommendations.[67] More broadly, the conditions under which deliberative processes generate trust or legitimacy gains remain only partly understood, especially as these processes are adapted across institutional settings. Theoretical legitimacy, deliberative quality, and positive political outcomes may be insufficient if the process and its outputs are not legible to participants, observers, and downstream decision makers.

**Three Case Studies**

Case #1: vTaiwan and Pol.is.

Audrey Tang – Taiwan's first Minister of Digital Affairs – is perhaps the most prominent public official using digital tools to enable public deliberation and gather direct input from constituents. In the mid 2010s, Tang deployed the *Pol.is* platform as part of her *vTaiwan* initiative to engage Taiwan's population in policy deliberations. The Pol.is platform is similar to a wiki-survey: Participants can contribute policy positions in text form and indicate their agreement or disagreement with a series of others' policy positions. Then, the platform uses clustering techniques and visualizations to highlight regions of agreement and convergence based on these agree/disagree voting patterns.

The policy application of this technique that has received the most attention concerns the regulation of Uber and other ride sharing services.[68] Visualization of initial positions revealed polarization, but ultimately the differing sides actually reached agreement on several core principles: fair liability insurance, transparent pricing, and no undercutting of market rates. These principles were incorporated into the government's regulation. Some observers have credited this participatory process with increasing support for, and perceived legitimacy of, the resulting policy.

[66] Pow, J., Van Dijk, L., & Marien, S. (2020). It's not just the taking part that counts: 'Like me'perceptions connect the wider public to minipublics. Journal of Deliberative Democracy, 16(2), 43-55.
[67] Germann, M., Marien, S., & Muradova, L. (2024). Scaling up? Unpacking the effect of deliberative mini-publics on legitimacy perceptions. Political Studies, 72(2), 677-700.
[68] See "Lessons From Consensus Building in Taiwan." Accessed March 28, 2026. https://democracy-technologies.org/participation/consensus-building-in-taiwan/.

While vTaiwan's original launch predated modern Large Language Models (LLMs), recent deliberative processes in Taiwan increasingly incorporate these technologies. For example, the "Talk to the City" platform uses LLMs to aid in facilitation, synthesize open-ended input, identify areas of overlap, and map potential paths to convergence.[69]

Case #2: The Habermas Machine.

A second case comes from researchers at Google DeepMind and multiple universities. They have proposed an AI-powered tool called the "Habermas machine" that is designed to improve public deliberations on complex and controversial issues. This tool works as follows: Participants submit their individual views in text format on the issue at hand and, in each round, an LLM drafts a set of statements intended to reflect competing perspectives while maximizing broad endorsement. Participants then evaluate these statements (e.g., approve/disapprove and provide feedback), and the system iteratively revises them. Unlike Pol.is – which primarily maps opinion clusters from voting patterns – the Habermas Machine uses generative AI in the mediator role to synthesize and refine candidate common-ground statements.[70]

Researchers evaluated the Habermas Machine in a multi-experiment study in the United Kingdom, which included randomized online experimental tasks and a virtual deliberative process. In the deliberative process, 200 participants met for three weekly one-hour sessions and deliberated over nine contentious policy topics using the Habermas Machine. They found that the LLM-produced common-ground statements were often positively endorsed (somewhat more often than those from human mediators), and within-group agreement increased on multiple issues from pre- to post-deliberation. This study's findings are suggestive that LLMs could usefully assist with mediation and common-ground discovery in deliberation. On the other hand, some critics argue that these observed opinion changes may simply result from people's tendencies to defer to algorithms, so more investigation is warranted.[71]

Case #3: Digital Dialogues in Libya

A third example comes from Libya in 2020–2021. After years of civil war in Libya, a team working with the United Nations sought to engage many Libyan citizens in a series of public deliberations about topics central to the country's transition to peace.[72] Working with Remesh AI,

[69] See "Amplifying Voices: Talk to the City in Taiwan — AI • Objectives • Institute." Accessed March 28, 2026. https://ai.objectives.institute/blog/amplifying-voices-talk-to-the-city-in-taiwan.
[70] Tessler, Michael Henry, Michiel A. Bakker, Daniel Jarrett, Hannah Sheahan, Martin J. Chadwick, Raphael Koster, Georgina Evans, et al. "AI Can Help Humans Find Common Ground in Democratic Deliberation." *Science* 386, no. 6719 (October 18, 2024): eadq2852. https://doi.org/10.1126/science.adq2852.
[71] Cohen, Joshua, and Henrik D. Kugelberg. eLetter "TRUST IN AI MEDIATORS MAY CHANGE DELIBERATIVE OUTCOMES." *Science*, May 5, 2025. https://philpapers.org/archive/COHTIA-2.pdf.
[72] Irwin, C., Masood, D., Wählisch, M., & Konya, A. (2021). Using artificial intelligence in peacemaking: the Libya experience. In *A. WAPOR 74th Annual Conference*.

they ran large-scale "Digital Dialogues" on an online, mobile-accessible platform. Participants provided both open-ended responses and answers to multiple-choice prompts. The key technical contribution was an AI-assisted analysis pipeline: The system rapidly organized open-text input into themes, estimated support for proposals, and compared patterns across participant subgroups (e.g., minority-group membership). Though its implementation was different, the AI system's purpose here was similar to that of the Habermas Machine, contributing sensemaking across multiple kinds of information. These dialogues were conducted in local dialects, and they were broadcast on both social media and live TV so that the rest of the population could witness them.

The team used this platform to conduct five digital dialogues between October 2020 and January 2021, with 1,000 Libyans participating in each round. Recruitment was handled by a local polling company. These dialogues addressed impacts of the war; militias and other fighters; economic issues; oil revenues; human rights; and elections. In one dialogue, participants proposed questions that they would like candidates for the new Government of National Unity (GNU) to publicly address. Organizers write that, "The questions collected through the Dialogue were later put to the GNU candidates who answered them on live television. These events achieved social media audiences of 1.7 million, a third of the Libyan population." As such, organizers present these Digital Dialogues as having informed – and potentially strengthened the legitimacy of – the process of forming Libya's interim government in 2021.

**Zooming Out: Opportunities and Challenges**

In the case studies above, we saw AI helping increase the ***scale*** of participation by connecting people across long distances, allowing lower-cost online participation, reducing language barriers, and reducing the need for extensively trained human facilitators. In Taiwan, tools such as Pol.is (and more recently Talk to the City) helped structure and process input across large numbers of participants, while in Libya the Digital Dialogues platform supported repeated large-scale, mobile-accessible deliberations with rapid analysis of participant input. Beyond these case studies, AI methods have also supported scaling in other ways, including in recruitment/sortition pipelines, agenda management, and information summarization.

While AI-based deliberation workflow tools promise to enable deliberation at greater scale, we risk centralizing control over key elements of the *substance* of deliberation, placing it in the hands of whoever designs the models, their inputs, and their outputs. Such risks can be exacerbated by *overreliance* on AI, a documented phenomenon where people tend to defer by default to AI recommendations, even when these recommendations are suboptimal or unreliable. A key question, then, is how to use automation to expand the scale and frequency of deliberation while ensuring that power over the process is distributed intentionally and remains so even after accounting for how participants will change their behavior in response to the presence of these tools.

The above case studies also illustrate how AI can contribute to ***sensemaking***: in Taiwan, Pol.is helped map agreement patterns and identify consensus in participants' responses; in the Habermas Machine, an LLM more explicitly synthesized across participants' inputs by iteratively generating and revising potential common-ground statements; and in Libya, an AI-assisted analysis pipeline organized open-text input into themes. Beyond these examples, there has recently been a rapid proliferation and deployment of AI models for performing sensemaking tasks.[73,74]

While sensemaking tools are highly promising, they are also embedded in the substance of deliberation and thus run the risk of becoming implicit architects of key choices in deliberation. For example, they can subtly steer attention, reframe issues, introduce biases, or collapse minority perspectives, all while appearing "objective" and "technically replicable." These risks pose another key challenge to future work: to empower participants with richer information and more incisive interpretations of others' ideas, while preserving their ability to inspect, challenge, and revise the interpretive choices that shape which ideas are elevated.

Across these cases, we also see AI being used to mitigate our third bottleneck, ***unequal participation and inclusion.*** In Libya, the Digital Dialogues were mobile-accessible and conducted in local dialects. In Taiwan, citizens could contribute online in a diversity of formats. These case studies allude to a future in which multimodal interfaces and AI assistance lower barriers to participation access and make participation less costly for citizens by fostering communication across language and geography. AI helps participants learn about issues more quickly and make inferences across diverse input formats, allowing people to contribute in the mode they find most natural.

When trying to elicit the input of a more diverse set of people at a lower participation cost, it is important to distinguish between techniques that make giving input easier, and those which use more sophisticated inference techniques to infer what people *would* say, without directly asking them. Proposals in the latter vein range from empowerment to *replacement*, with some proposals approaching the idea of replacing participants with AI-based proxies as a way of avoiding the time and effort costs of human participation. Aside from the question of whether it is technically possible for AI to act as a faithful human proxy, such proposals raise the question of the very purpose of democratic participation: Is it simply a transfer of information or is it a personal act of investment in the public sphere? As our technological toolkit encompasses more tools aimed at

[73] Fish, S., Gölz, P., Parkes, D. C., Procaccia, A. D., Rusak, G., Shapira, I., & Wüthrich, M. (2023). Generative social choice. *arXiv preprint arXiv:2309.01291*.

[74] Kabbara, J., Phan, T. M., Rakhilin, M., Detwiller, M. E., Dimitrakopoulou, D., & Roy, D. (2025, April). AI-assisted sensemaking: Human-AI collaboration for the analysis and interpretation of recorded facilitated conversations. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (pp. 1-8).

more equal inclusion, one can question which elements of democratic participation a given tool empowers versus replaces – and how may this differ along the same (or different) societal dimensions on which inequality appears today.

Finally, these case studies offer a glimpse into the future of how AI can be used to enhance ***trust and perceived legitimacy***. In both Taiwan and Libya, AI tools helped produce outputs that were more legible to broader publics and more usable by decision makers. In Libya in particular, the dialogues were also broadcast on social media and live television, extending visibility beyond direct participants. Looking ahead, AI may help make citizen deliberation more accessible to the wider public by allowing members of the public to directly query detailed information about the process. On the other hand, the opacity and inherent technicality of AI models can shift trust from democratically accountable institutions and facilitators to inscrutable systems and their designers. When recommendations come from an AI-infused deliberation pipeline, there is the risk that trust becomes more a matter of accepting the semi-automated pipeline that produced them, rather than considering the reasons of deliberation participants.

Taken together, these cases – and more broadly, the exploding landscape of deliberative technologies – illustrate why there is so much excitement about AI in this space: AI can offer toolkits that can help address some of the most stubborn bottlenecks in enmeshing more deliberative participation in our democratic institutions. Ultimately, however, whether AI deepens democracy will be decided not just by the technical sophistication of our tools but by the political choices that govern their use – whether they augment judgment or automate it; whether they make deliberation more legible or simply relocate trust from publics and institutions to technical systems; whether they reduce barriers to participation or reproduce inequality behind a more inclusive interface; and whether they strengthen public reasoning or bypass it in the name of efficiency.

## Conclusion

The integration of generative artificial intelligence in political campaigns, election administration, social movements, and citizen deliberation hold both great promise but also the potential to damage the quality of our democracies. On one hand, AI potentially broadens and deepens democracy by engaging and empowering citizens, enriching the connections between citizens and their leaders and facilitating sophisticated large-scale deliberation. These capabilities can empower resource-strapped campaigns, enhance the service quality of under-funded election offices, and provide marginalized social movements with new tools for mobilization and

accountability. By addressing long-standing bottlenecks in participation and sensemaking, AI holds the promise of making democratic institutions more responsive and inclusive. However, these same tools can be weaponized to mislead voters, amplify illiberal politics, or centralize control in the hands of well-resourced platforms and technical elites.

No one can now say what the impact of AI on elections, campaigns, movements, and public deliberation will be. That future depends in part on the development of AI technologies and in part on how a vast array of political actors utilize these technologies to enhance their capabilities and achieve their ends. Will AI be open or closed? Biased and narrow in perspective or encompassing? Will the information they provide serve truth-seeking impulses or serve those who can pay the most or who seek domination? Will effective technologies amplify or flatten pre-existing asymmetries of resources and power in the political sphere? The answers to these questions do not depend in the first instance on pure technological questions such as the sophistication of AI algorithms, the speed with which advanced chips can execute matrix-multiplication, and the scale of data centers. Rather, the future of the democratic environment depends on regulatory frameworks like the EU AI Act, the competitive maneuvers of political and economic actors, and the ways in which leaders integrate these tools into existing organizations and create new organizations. Ultimately, the impact of AI upon the quality of democracy in the future depends upon how democracies decide to shape AI and its political uses now.

## References

ABP News Bureau. (2024, May 29). Alia Bhatt, Kareena Kapoor and other celebs support Gaza as "All Eyes On Rafah" trends after Israel's attack. *ABP News*. https://news.abplive.com/entertainment/celebrities/all-eyes-on-rafah-indian-celebs-alia-bhatt-kareena-kapoor-priyanka-chopra-support-gaza-after-israel-attack-on-rafah-1691055

Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, *33*(2), 3–30. https://doi.org/10.1257/jep.33.2.3

Allyn, B. (2024, June 3). 'All eyes on Rafah' is the internet's most viral AI image: Two artists are claiming credit. New England Public Media. https://www.nepm.org/national-world-news/2024-06-03/all-eyes-on-rafah-is-the-internets-most-viral-ai-image-two-artists-are-claiming-credit

Al-Ubaydli, O., List, J. A., & Suskind, D. (2019). *The science of using science: Towards an understanding of the threats to scaling experiments* (NBER Working Paper No. 25848). National Bureau of Economic Research. https://doi.org/10.3386/w25848

Banerjee, D., Dutta, S., & Dutta Choudhury, S. (2024, June 5). AI amplifies political reach but magnifies disinformation in India elections. Asia Pacific Foundation of Canada. https://www.asiapacific.ca/publication/indian-election-use-of-ai-political-campaigns-voter-engagement

Benford, R. D., & Snow, D. A. (2000). Framing processes and social movements: An overview and assessment. *Annual Review of Sociology*, *26*, 611–639. https://doi.org/10.1146/annurev.soc.26.1.611

Cortés, E., Norden, L., Frase, H., & Hoffmann, M. (2023). *Safeguards for using artificial intelligence in election administration*. Brennan Center for Justice. https://www.brennancenter.org/our-work/research-reports/safeguards-using-artificial-intelligence-election-administration

Brennan Center for Justice. (2025a). *Local election officials survey — July 2025*. https://www.brennancenter.org/our-work/research-reports/local-election-officials-survey-july-2025

Brennan Center for Justice. (2025b). *Local election officials survey — May 2024*. https://www.brennancenter.org/our-work/research-reports/local-election-officials-survey-may-2024

Burga, S. (2024, May 29). Why the 'All Eyes on Rafah' AI post is going viral on social media. *Time*. https://time.com/6983344/all-eyes-on-rafah-ai-post-viral/

Burgess, A., & Morris, B. (2024, May 28). As the world reacts to Israel's deadly air strike on Rafah, here are key events from the Gaza war leading up to the attack. *ABC News*. https://www.abc.net.au/news/2024-05-28/israel-gaza-war-time-line-before-fatal-rafah-air-strike/103901144

Centre for Responsible Union AI. (n.d.-a). Beyond the hype: What union staff actually want from AI. https://www.agileunions.ai/p/what-union-staff-actually-want-from-ai

Centre for Responsible Union AI. (n.d.-b). How ABVV-Metaal used AI to automate creating and sending union leave letters for reps. https://www.agileunions.ai/p/automate-creating-sending-union-leave-letters

Centre for Responsible Union AI. (n.d.-c). How NASUWT used an AI chatbot to make their knowledge base more accessible to members. https://www.agileunions.ai/p/case-study-nasuwt-chatbot-members-resignation-notice-periods

Centre for Responsible Union AI. (n.d.-d). How PCS created an AI chatbot to support union reps practice conversations to recruit new members. https://www.agileunions.ai/p/case-study-repcoach-pcs-union-reps-practice-recruitment-conversations

Cohen, J., & Kugelberg, H. D. (2025, May 5). Trust in AI mediators may change deliberative outcomes [eLetter]. *Science*. https://philpapers.org/archive/COHTIA-2.pdf

Cushing Rodriguez, S. (2023, October 4). Consensus building in Taiwan, the poster child of digital democracy. Democracy Technologies. https://democracy-technologies.org/participation/consensus-building-in-taiwan/

Dahl, R. A. (1990). *After the revolution? Authority in a good society*. Yale University Press.

Dunworth, L. (2024, May 29). Dua Lipa shares support for #AllEyesOnRafah: "Burning children alive can never be justified." *NME*. https://www.nme.com/news/music/dua-lipa-shares-support-for-all-eyes-on-rafah-3760399

Elliott, V. (2024, June 11). There's an AI candidate running for parliament in the UK. *Wired*. https://www.wired.com/story/ai-candidate-running-for-parliament-uk/

Fish, S., Gölz, P., Parkes, D. C., Procaccia, A. D., Rusak, G., Shapira, I., & Wüthrich, M. (2023). *Generative social choice* (arXiv:2309.01291) [Preprint]. arXiv. https://arxiv.org/abs/2309.01291

Flanigan, B., Gölz, P., Gupta, A., Hennig, B., & Procaccia, A. D. (2021). Fair algorithms for selecting citizens' assemblies. *Nature*, *596*(7873), 548–552. https://doi.org/10.1038/s41586-021-03788-6

Foos, F. (2024). The use of AI by election campaigns. *LSE Public Policy Review*, *3*(3). https://doi.org/10.31389/lseppr.112

Foos, F., John, P., Tomas-Valiente, F., & Cinar, A. C. (n.d.-a). *Scaling-up mobilization campaigns. Chapter 4: Intervention size* [Manuscript in preparation].

Foos, F., John, P., Valiente-Jorda, F., & Cinar, A. C. (n.d.-b). *Scaling-up political campaigns* [Manuscript in preparation].

Freed, B. (2020, September 17). Election questions? Idaho's new Watson-powered chatbot has answers. StateScoop. https://statescoop.com/election-chatbot-idaho-watson-ai/

Fung, A., & Lessig, L. (2023, June 2). How AI could take over elections — and undermine democracy. *The Conversation*. https://theconversation.com/how-ai-could-take-over-elections-and-undermine-democracy-206051

Germann, M., Marien, S., & Muradova, L. (2024). Scaling up? Unpacking the effect of deliberative mini-publics on legitimacy perceptions. *Political Studies*, *72*(2), 677–700. https://doi.org/10.1177/00323217221137444

Green, D. P., & Gerber, A. S. (2019). *Get out the vote: How to increase voter turnout* (4th ed.). Brookings Institution Press.

Greschler, G. (2024, February 27). Voters are being called by campaigns' AI bots. I talked to one for 10 minutes. *San Francisco Standard*. https://sfstandard.com/2024/02/27/peter-dixon-congress-ai-robocalls-civox/

Guevara, F. (2025, March 4). Using AI now to improve movements' effectiveness: A basic introduction for social activists [Commentary]. Ash Center for Democratic Governance and Innovation.

https://ash.harvard.edu/articles/using-ai-now-to-improve-movements-effectiveness-a-basic-introduction-guide-for-social-activists/

Guglielmo, C. (2023, October 21). AI and you: NYC mayor can't really speak Mandarin, the AI money trail, who sets the rules. CNET. https://www.cnet.com/tech/computing/ai-and-you-nyc-mayor-cant-really-speak-mandarin-the-ai-money-trail-who-sets-the-rules/

Hackenburg, K., Tappin, B. M., Hewitt, L., et al. (2025). The levers of political persuasion with conversational artificial intelligence. *Science*, *390*(6777), Article eaea3884. https://doi.org/10.1126/science.aea3884

Hall, A. (2026, March 26). Building political superintelligence. Substack. https://substack.com/home/post/p-192214946

Imran Khan uses A.I. to give victory speech in Pakistan. (2024, February 11). *The New York Times*. https://www.nytimes.com/2024/02/11/world/asia/imran-khan-artificial-intelligence-pakistan.html

Irwin, C., Masood, D., Wählisch, M., & Konya, A. (2021). *Using artificial intelligence in peacemaking: The Libya experience*. WAPOR 74th Annual Conference.

Issenberg, S. (2013). *The victory lab: The secret science of winning campaigns*. Crown.

Jie, J. (2025, March). Unlocking AI's potential for social movements. Ash Center for Democratic Governance and Innovation. https://ash.harvard.edu/articles/unlocking-ais-potential-for-social-movement/

Jungherr, A., Rauchfleisch, A., & Wuttke, A. (2026). Artificial intelligence in election campaigns: Perceptions, penalties, and implications. *Political Communication*. Advance online publication.

Kabbara, J., Phan, T. M., Rakhilin, M., Detwiller, M. E., Dimitrakopoulou, D., & Roy, D. (2025). AI-assisted sensemaking: Human-AI collaboration for the analysis and interpretation of recorded facilitated conversations. In *Extended abstracts of the CHI Conference on Human Factors in Computing Systems* (pp. 1–8). Association for Computing Machinery.

Kalla, J. L., & Broockman, D. E. (2016). Campaign contributions facilitate access to congressional officials: A randomized field experiment. *American Journal of Political Science*, *60*(3), 545–558. https://doi.org/10.1111/ajps.12180

Lin, H., Czarnek, G., Lewis, B., et al. (2025). Persuading voters using human–artificial intelligence dialogues. *Nature*, *648*(8093), 394–401. https://doi.org/10.1038/s41586-025-09771-9

List, J. A. (2022). *The voltage effect: How to make good ideas great and great ideas scale*. Currency.

Mann, C. B. (2021). Can conversing with a computer increase turnout? Mobilization using chatbot communication. *Journal of Experimental Political Science*, *8*(1), 51–62. https://doi.org/10.1017/XPS.2020.5

National Conference of State Legislatures. (2026, March 24). Artificial intelligence in election administration. https://www.ncsl.org/elections-and-campaigns/artificial-intelligence-in-election-administration

Neblo, M. A., Esterling, K. M., & Lazer, D. M. J. (2018). *Politics with the people: Building a directly representative democracy*. Cambridge University Press.

New York State Senate. (2023). *Senate Bill S9678, 2023–2024 legislative session: Relates to materially deceptive media in political communications*. https://www.nysenate.gov/legislation/bills/2023/S9678

Offer-Westort, M., et al. (2026). Deep canvassing with automated conversational agents: Personalized messaging to change attitudes. *Research & Politics*, *13*(1), Article 20531680251414926. https://doi.org/10.1177/20531680251414926

OpenAI bans developer of Dean Phillips bot. (2024, January 20). *The Washington Post*. https://www.washingtonpost.com/technology/2024/01/20/openai-dean-phillips-ban-chatgpt/

Orrick. (n.d.). U.S. AI law tracker. https://ai-law-center.orrick.com/us-ai-law-tracker-see-all-states/

Participedia. (n.d.-a). 21st-century town meeting. https://participedia.net/method/145

Participedia. (n.d.-b). AmericaSpeaks. https://participedia.net/organization/americaspeaks

Partnership for Public Service. (n.d.). Leadership program inspires an AI revolution at the State Department. https://ourpublicservice.org/about/impact/leadership-program-inspires-an-ai-revolution-at-the-state-department/

PBS News. (2024, February 27). Chatbot info on U.S. elections is inaccurate, misleading and could keep voters from polls, report finds. *PBS News*. https://www.pbs.org/newshour/politics/chatbot-info-on-u-s-elections-is-inaccurate-misleading-and-could-keep-voters-from-polls-report-finds

Peterson, T. (2025, November 4). How The Times is using AI to model synthetic focus groups from human audiences. Digiday. https://digiday.com/media/how-the-times-is-using-ai-to-model-synthetic-focus-groups-from-human-audiences/

Picón, I. (2025, March 6). AI-powered bots and electoral participation: A view from the Venezuelan experience. Ash Center for Democratic Governance and Innovation. https://ash.harvard.edu/articles/ai-powered-bots-and-electoral-participation-a-view-from-the-venezuelan-experience/

Pow, J., Van Dijk, L., & Marien, S. (2020). It's not just the taking part that counts: 'Like me' perceptions connect the wider public to mini-publics. *Journal of Deliberative Democracy*, *16*(2), 43–55. https://doi.org/10.16997/jdd.395

Reuters. (2016, May 28). France's Macron launches door-to-door campaign as presidency battle looms. *Reuters*. https://www.reuters.com/article/world/frances-macron-launches-door-to-door-campaign-as-presidency-battle-looms-idUSKCN0YJ09E/

Sakunia, S. (2024, July 12). AI and deepfakes played a big role in India's elections. *New Lines Magazine*. https://newlinesmag.com/spotlight/ai-and-deepfakes-played-a-big-role-in-indias-elections/

Samuels, B. (2023, August 14). How a startup is using AI to write fundraising emails. The Hill. https://thehill.com/homenews/campaign/4146154-how-a-start-up-is-using-ai-to-write-fundraising-emails/

Schmidt, S., Rich, S., Herrero, A. V., & Paúl, M. L. (2024, August 4). Maduro lost election, tallies collected by Venezuela's opposition show. *The Washington Post*. https://www.washingtonpost.com/world/2024/08/04/maduro-gonzalez-election-actas-analysis/

Schneier, B., & Sanders, N. E. (2025). *Rewiring democracy: How AI will transform our politics, government, and citizenship*. The MIT Press.

Schneier, B., & Sanders, N. E. (2026, January 11). Rewiring democracy now: A new kind of political engagement emerges in Japan. The Renovator. https://therenovator.substack.com/p/rewiring-democracy-now

Silicon Valley insiders are trying to unseat Biden with help from AI. (2024, January 18). *The Washington Post*. https://www.washingtonpost.com/elections/2024/01/18/ai-tech-biden/

Stewart, C., III. (2022). Trust in elections. *Daedalus*, *151*(4), 234–253. https://doi.org/10.1162/daed_a_01953

Tessler, M. H., Bakker, M. A., Jarrett, D., Sheahan, H., Chadwick, M. J., Koster, R., Evans, G., Campbell-Gillingham, L., Collins, T., Parkes, D. C., Botvinick, M., & Summerfield, C. (2024, October 18). AI can help humans find common ground in democratic deliberation. *Science*, *386*(6719), Article eadq2852. https://doi.org/10.1126/science.adq2852

Trump backers inundate election offices with requests for 2020 records. (2022, September 11). *The Washington Post*. https://www.washingtonpost.com/nation/2022/09/11/trump-election-deniers-voting/

Tseng, Z. (n.d.). Amplifying voices: Talk to the City in Taiwan. Objectives Institute. https://ai.objectives.institute/blog/amplifying-voices-talk-to-the-city-in-taiwan

UK Parliament. (n.d.). Election result for Brighton Pavilion (constituency). https://members.parliament.uk/constituency/3945/election/422

U.S. Election Assistance Commission. (n.d.). *Signature verification and cure process*. https://www.eac.gov/sites/default/files/electionofficials/vbm/Signature_Verification_Cure_Process.pdf

Ulrich, S. (2025, February 3). Using data to expose systemic gender-based violence — in 10 steps. Global Investigative Journalism Network. https://gijn.org/stories/10-steps-data-expose-gender-violence-eswatini/

Westwood, S., Grimmer, J., & Hall, A. (2025, May 8). *Measuring perceived slant in large language models through user evaluations* [Working paper]. Stanford Graduate School of Business. https://www.gsb.stanford.edu/faculty-research/working-papers/measuring-perceived-slant-large-language-models-through-user

Yoder, K., & Tan, A. (2024, June). *CEIR focus brief: Election officials & the misuse of public records requests*. The Center for Election Innovation & Research.

https://electioninnovation.org/wp-content/uploads/CEIR_Public-Records-Requests-Brief_June-2024.pdf