POLICY BRIEFS SERIES

AI for Democracy Movements: Toward a New Agenda

Summary Report of an Implementation Workshop Convened by the Nonviolent Action Lab, Harvard Kennedy School April 2026

DECEMBER 2025, ASH CENTER FOR DEMOCRATIC GOVERNANCE AND INNOVATION, HARVARD KENNEDY SCHOOL, CAMBRIDGE, MASSACHUSETTS, USA

HARVARD Kennedy School
ASH CENTER
for Democratic Governance and Innovation

POLICY BRIEFS SERIES

AI for Democracy Movements: Toward a New Agenda

Summary Report of an Implementation Workshop Convened by the Nonviolent Action Lab, Harvard Kennedy School April 2026

DECEMBER 2025, ASH CENTER FOR DEMOCRATIC GOVERNANCE AND INNOVATION, HARVARD KENNEDY SCHOOL, CAMBRIDGE, MASSACHUSETTS, USA

Workshop Participants

Matthew Cebul
Albert Cevallos
Erica Chenoweth
Emily Chatterjee
Jennifer Earl
Cassy Dorff
Julian Fortuna
Nathan Freitas
Kate Gage
Freddy Guevara Cortez
Soha Hammam
Suhan Kacholia
Jonathan Jayes-Green
Hillary Lehr
Jie (Jason) Lian
Evan Perkoski
Jamila Raqib
Max Reynolds
Causten Rodriguez Wollerman
Justin Ruben
Hannah Sassaman
Christopher Shay
Jhanisse Vaca Daza
Dizzy Zaba

About the Ash Center

The Mission of the Roy and Lila Ash Center for Democratic Governance and Innovation at is to develop ideas and foster practices for equal and inclusive, multi-racial and multi-ethnic democracy and self-government.

This paper is one in a series published by the Ash Center for Democratic Governance and Innovation at Harvard University's John F. Kennedy School of Government. The views expressed in the Ash Center Policy Briefs Series are those of the author(s) and do not necessarily reflect those of the John F. Kennedy School of Government or of Harvard University. The papers in this series are intended to elicit feedback and to encourage debate on important public policy challenges.

Acknowledgments

We are grateful to the MacArthur Foundation and the Ash Center for Democratic Governance and Innovation for supporting the workshop. We also appreciate Dana Guterman's editorial assistance.

How to Cite This Report

Erica Chenoweth, Matthew Cebul, and Jie (Jason) Lian, eds. *AI for Democracy Movements: Toward a New Agenda*. Ash Center for Democratic Governance and Innovation Occasional Paper, Harvard Kennedy School, 2026.

A PUBLICATION OF THE

Ash Center for Democratic Governance and Innovation
Harvard Kennedy School
79 John F. Kennedy Street
Cambridge, MA 02138
617-495-0557
ash.harvard.edu

Contents

Glossary of Terms[1]

Artificial Intelligence (AI): The broad field of building computer systems that perform tasks normally requiring human intelligence, such as understanding language, recognizing images, or making decisions. Modern AI usually refers to software that learns patterns from data rather than following hand-written rules.

Generative AI: AI that creates new content—text, images, audio, video, or code—rather than just analyzing existing content. Common examples include AI writing an essay, creating an image from a description, or composing music.

Large Language Model (LLM): A program trained on billions of pages of text that has learned to predict what word comes next. This capability allows the model to hold conversations, write essays, and answer questions. For instance, Claude and ChatGPT are both built on LLMs.

Chatbot: A program designed for text-based conversation. Today's chatbots use LLMs to understand questions and generate natural-sounding responses. They range from simple customer service bots to sophisticated research assistants.

AI Agent: An AI system that can take independent actions—such as browsing the web, writing and running code, sending emails, or making purchases—rather than just answering questions. Unlike a chatbot that waits for your next message, an agent pursues a goal through multiple steps with minimal human direction.

AI Literacy: The ability to understand what AI is, how it works at a basic level, and how it affects your life—including its limitations and risks. Much like digital or media literacy, AI literacy is increasingly seen as a civic skill that schools, organizations, and governments should teach.

Open-Source AI: AI models whose underlying code and learned knowledge (stored as numerical values) are publicly available for anyone to inspect, use, or modify. The term is often contested; critics argue that models restricting commercial use or withholding their training data are not truly open source. This distinction matters for policy, procurement, and transparency debates.

Vibe Coding: Writing software by providing plain-English instructions to an AI and letting it generate the code, rather than writing it yourself line-by-line. The programmer focuses on the "vibe"—the intent and direction—while the AI handles the technical details. The term was coined by Andrej Karpathy in early 2025 and reflects how AI is changing who can build software. The resulting code still needs human review for accuracy and security.

AI Bias: When an AI system produces results that are systematically unfair to certain groups of people. This can happen because the training data reflects existing societal inequalities—for example, a hiring AI trained on historical data may learn to favor male candidates if most past hires were men. Bias can also be introduced through the design choices of those building the system. Addressing bias is a central concern in AI regulation and ethics.

Hallucination: When an AI confidently presents factually wrong or completely invented information as truth. The AI is not "lying"—it is generating plausible-sounding text without checking whether it is true. This is one of the most common and dangerous limitations of current AI systems.

Alignment: The challenge of ensuring AI systems actually do what humans want them to do and avoid causing unintended harm. A misaligned AI might technically achieve its goal but do so in a harmful

way—for instance, a content algorithm that maximizes engagement by promoting outrage. Alignment research is a growing field focused on ensuring advanced AI remains safe and beneficial; it is central to policy debates about AI safety governance.

Guardrails: Rules and limits built into an AI system to prevent it from producing harmful, offensive, or dangerous content. For example, these rules might prevent models from providing instructions for making weapons. Guardrails are a key topic in AI regulation.

Introduction

The swift rise of artificial intelligence (AI) holds the potential to reshape a wide range of fields —including the politics of nonviolent democracy movements. While recent scholarship on AI and democratization has largely focused on how authoritarian regimes use the technology for digital surveillance and repression,[2] an important parallel question remains understudied: How can organizers and activists use AI to empower nonviolent democracy movements?

In December 2024, Harvard Kennedy School's Nonviolent Action Lab (NVA Lab) convened an exploratory workshop on this topic. The resulting report outlined several possible uses of AI by democracy movements and urged practitioners, social scientists, and technologists to forge a new community of AI research and practice.[3] Since then, interest in AI for democracy movements has only grown among activists, organizers, and funders. Activists are increasingly applying AI tools to their work, sparking both curiosity and skepticism about AI as a force for democracy. Given the many ongoing challenges of promoting democracy and the increasing technological sophistication of autocrats,[4] this is a crucial moment to interrogate what AI can and cannot do for the world's democracy movements.

To that end, the NVA Lab convened a second scholar-practitioner workshop in December 2025, including a wider array of technologists and practitioners alongside social scientists, to implement some of the key recommendations emerging from the 2024 meeting. We first sought to take stock of the rapid progress in knowledge about AI integration into social movements over the past year, learn about prominent use cases, and identify key challenges for organizers considering AI adoption. We further aimed to launch collaborative working groups, organized by the NVA Lab, to support the broader democracy movement ecosystem as the AI era progresses.

This report summarizes key insights from the meeting and ongoing working group discussions. We identify broad points of consensus among participants about AI adoption in movements, alongside persistent tensions and uncertainties that scholars and practitioners must navigate in the days ahead. From this discussion, we distill three interventions that are urgently needed in the democracy field: (1) an AI code of conduct for democracy movement groups and organizations, (2) a targeted research agenda on AI in democracy movements, and (3) efforts to expand and strengthen collaborative linkages between the technologist and democracy activist communities. These core interventions correspond to ongoing scholar-practitioner working groups that are driving progress along each of these lines of effort.

Points of Consensus

AI is a dynamic and complex field, and workshop participants found that many issues under discussion defied simple resolution. Participants grappled with serious disagreements, and complete unanimity was rare—even for areas of broad agreement. Generally speaking, however, participants converged on three high-level areas of consensus regarding the incorporation of AI tools into pro-democracy movements.

The Time to Adapt Is Now

First, participants shared a general sense of urgency for movements to adapt to the rapidly shifting technological and political landscape driven by advances in AI. While a healthy skepticism toward AI boosterism is warranted, particularly given numerous ethical and strategic risks surrounding uncontrolled AI usage, most participants concluded that an absolute rejection of the technology would carry significant strategic costs. Failing to engage with AI now would only leave democracy movements further behind.

Participants noted that movement opponents have already embraced AI and are reaping rewards for their early adoption. Authoritarian governments, anti-democratic civic groups, and major corporations are all experimenting with and integrating AI into surveillance, persuasion (disinformation), and anti-democratic organizing infrastructures. Given that AI-enhanced illiberalism is here to stay and may worsen over time, democracy advocates must learn to effectively respond to these developments. That requires them—at a minimum—to understand the state of the field and more systematically track authoritarian uses of AI. Holding AI technologies at arm's length prevents democracy movements from understanding their opponents' evolving capabilities and developing necessary countermeasures.

Beyond defending against digital repression, AI has clear potential to proactively empower democracy movements. Participants highlighted AI's ability to help organizers collect, synthesize, and analyze large volumes of information, handle repetitive or labor-intensive cognitive tasks, and support targeted and sustained engagement with key constituencies. AI tools will not solve every movement problem, and as discussed below, hasty or ill-conceived forms of adoption could exacerbate movement weaknesses. Yet some potential gains are clearly worth exploring, and responsible and intentional adoption could meaningfully expand democracy movements' capacity and reach.

In sum, participants generally recognized that ignoring the rise of AI is not a neutral default position. Instead, it is a consequential strategic choice that likely disadvantages democracy activists and organizers. The central challenge is not whether or not to engage with AI, but how to do so in ways that are both effective and ethically grounded.

Human-Centered Tech Adoption

The second area of high-level consensus was that humans must remain at the center of tech-empowered movements. Organizers should retain responsibility for strategy, judgment, and values, while AI should serve primarily as an accelerator or amplifier. In practice, this means using AI to scale or handle well-defined, labor-intensive tasks—such as information collection, data cleaning, drafting and summarization, and outreach logistics—rather than to make strategic decisions or determine movements' overall direction.

Both practical and ethical concerns motivated this point of consensus. One widespread concern was that AI tools may feature embedded bias and systematic errors. Participants stressed that AI systems are built on uneven and often opaque data, including Western-centric knowledge bases and other inherent biases, and warned that unquestioning adoption of such tools could mislead organizers or produce counterproductive strategies. Moreover, participants emphasized the need to safeguard core human organizing skills in the AI era, including strategic analysis, nonviolent discipline, coalition building, and narrative work. They expressed concern that overreliance on AI could depersonalize activism and hollow out the relational, in-person practices that make movements effective and resilient.

These concerns are matters of degrees. Different participants were more or less comfortable with engaging AI tools in strategic processes, and they would likely draw different red lines for AI misuse. Still, participants were united in the belief that AI cannot replace the core human elements of social movement work and that adoption should be in service of enhancing movement human capacity, not reducing it.

Leading with First Principles

The final area of high-level consensus was that there are significant ethical concerns surrounding the introduction of AI tools into social movements and that navigating these concerns requires a clear, overarching statement of ethical first principles for AI adoption.

The potential normative pitfalls of AI adoption are legion. They include privacy concerns and misuse of personal data, systematic AI biases and errors, and the potential to exacerbate inequalities

between actors with and without access to AI, particularly between the global North and South. Furthermore, AI raises questions around whether it meaningfully abides by duty of care or "do no harm" ethical standards, the human costs of worker displacement, and the potential loss of democracy movements' moral authority. We expand on some of these ethical tensions below.

In attempting to navigate these issues, participants sought to articulate a set of emerging norms that should guide democracy movements' use of AI. Central to these is that AI must only be used in service of the work of building community and movement power, humans must remain accountable for movements' strategic choices, and AI tools should be based on transparent data sources and carefully judged by whether they strengthen movements' collective capacity, discipline, and safety. Moreover, movements should own or control their own AI infrastructure as much as possible, and democracy organizations should develop and abide by explicit AI use policies to mitigate security and privacy risks. Finally, AI design decisions should assume surveillance, adversarial use, and cross-border diffusion rather than benign conditions.

Different workshop participants would likely emphasize different fundamental principles of AI adoption. But the underlying point is the need for a shared, baseline code of conduct to ground movements in a clear ethical framework and help them embed positive AI norms into daily practice. Developing this code is a work in progress, as discussed further below.

Uncertainties and Outstanding Questions

Beyond these general points of agreement, the convening also unearthed areas of persistent tension and uncertainty about emerging technologies and democracy movements. We highlight five such areas: (1) the overwhelming pace of AI innovation, (2) pro-social vs. anti-social AI uses, (3) AI's influence over movement strategy, (4) global variation in AI adoption, and (5) competing models of AI adoption.

Keeping Pace with Change

Perhaps the most fundamental element of uncertainty surrounding AI integration with democracy movements is the astounding rate of AI technological development. These technologies have been characterized by explosive growth, and the field continues to evolve rapidly. AI models are gaining new capabilities quickly and advancing in unexpected ways—agentic models capable of complex, autonomous work have now surpassed "basic" chatbot AI capabilities that were revolutionary just a few years ago. These AI agents are being deployed for all sorts of tasks, from workflow assistants for emails and scheduling (such as OpenClaw) to AI targeting systems in warfare. While some of the AI mania may be profit-seeking hype, the frenetic pace of innovation is undeniable.

This breakneck speed has made it exceptionally challenging for non-specialists—even those who use AI tools regularly—to keep pace with the field. Participants observed that some of their beliefs in the previous workshop about AI's capabilities and likely impact on democracy movements had already been overtaken by events on the ground. With capabilities and innovative applications evolving this rapidly, both the opportunities for and risks of AI adoption for democracy movements remain difficult to pin down.

An important element of AI's relentless dynamism is a highly permissive regulatory environment. Globally, AI is minimally regulated—the 2024 EU AI Act is already outdated, and the Trump administration has explicitly prioritized AI innovation and sought to quash state-level regulations. As a result, the global AI rollout is occurring with minimal governmental rules or guidelines. This has already produced many cases of unethical or irresponsible uses of AI by governments, such as police wielding systematically discriminatory facial recognition software and various data privacy violations.[5] Malicious AI applications can clearly impact movements, but it is important to acknowledge that activists themselves

are not immune to dubious or reckless uses of AI. For instance, should activists be using AI to generate deepfake images or recordings of their opponents, or to spam legislators with AI-generated "constituent messages"? Absent clear regulation of AI technologies, the onus of responsibility falls heavily on civil society and activists themselves to navigate a world of new security, privacy, and usage concerns.

In short, the tremendous pace of AI development and innovation makes it difficult for democracy practitioners to anticipate future trends or offer enduring technical guidance. This uncertainty further motivated practitioners to establish a set of general first principles for AI adoption to help movements navigate the technical frontier as it evolves.

Is AI Pro- or Anti-Social?

At a basic level, a social movement is a network of people mobilized around a shared goal. Social movements are inherently **social** and grounded in relationships. Expanding, strengthening, and activating these relationships is the heart of organizing. By forging connections across differences and strengthening community ties, organizers establish the relational bedrock of people power. Along these lines, social movement scholarship contends that movements with deep community roots and extensive, cross-cutting interpersonal networks are better able to mobilize participants and more likely to succeed.6

Workshop participants agreed that AI technologies cannot replace this crucial social component of movement building; there is no substitute for cultivating personal relationships. In this vein, some participants worried that overreliance on AI and other new technologies could have **anti-social** effects, undermining movement investments in pro-social coalition building. If this is the case, an over-fixation on AI adoption could rob democracy movements of a key source of strength.

Yet it is difficult to predict whether AI—or any new technology—will prove to be a pro- or anti-social force. Technology can be alienating or uniting, and it is sometimes both at once. The rise of social media is an instructive case in point. Social media rapidly expanded interpersonal networks to an almost unfathomable degree, enabling activists to directly reach millions of strangers virtually overnight. Early scholarship lauded these advances in "liberation technology" as unalloyed goods for social movements. But at the same time, overreliance on online activism may have detracted from investments in local, in-person organizing, making movements broad and brittle rather than deep and resilient.[7] In other words, social media has had important pro- and anti-social effects on democracy movements.

Participants identified at least two examples of the pro- and anti-social tension for AI and democracy movements. The first involves **authenticity**. Authenticity is a highly prized quality in political communication that helps both politicians and social movements forge genuine connections and earn popular support. AI products that appear disingenuous or fake may undermine this image, alienating movements from their potential supporters. For instance, AI-generated ad campaigns from McDonald's and Coca-Cola generated intense public backlash to "AI slop," undermining their brands.[8] Likewise, deepfakes and other AI-generated media are already being used as disinformation; movement efforts to co-opt this technology to grab attention or dramatize their actions could merely increase popular distrust in them.[9]

Nevertheless, participants also identified ways in which AI technologies could help movements forge authentic connections. For instance, AI text analysis tools could help organizers synthesize feedback from canvassing, improving their responsiveness to community concerns. More ambitiously, AI chatbots can directly tailor outreach and calls to action in response to constituent engagement and even conduct autonomous follow-ups to ensure people follow through on their plans for political activism. Practitioners who have used these tools were surprised to discover that many respondents engaged positively with chatbots even when they knew they were not human. They suspected that respondents liked the AI's interactive responsiveness—the sense of **being heard**—as compared to non-personalized

mass emails or SMS messages. In both cases, AI tools could plausibly enhance movements' authentic connections to their intended audiences, though these effects could attenuate over time if AI outreach becomes oversaturated.

A second example of this tension involves movements' human capital. On one hand, a core claim from proponents of AI adoption is that the technology will increase productivity by offloading menial labor from skilled workers. In theory, then, movement organizers empowered by AI could devote more time toward field work and meaningful, pro-social engagement in their communities. On the other hand, activists may need repetitive experience with the day-to-day tasks of movement building—or introductory on-ramps into organizations—to become skilled organizers in the first place. If so, automating these basic tasks could inadvertently degrade organizer capacity or the pipeline of experienced organizers, consequently undermining movement efforts to build community power.

Ultimately, AI adoption can have both pro- and anti-social effects on democracy movements, and their cumulative impacts will be difficult to assess in advance. Activists must take care that new technologies strengthen ties between organizers and their communities rather than depersonalizing those relationships, ensuring that technology is a complement to, not a replacement for, human connections.

AI and Movement Strategy

Another area of tension involves the appropriate role of AI in democracy movement strategy. Here, participants identified both promises and perils, and while they generally agreed that AI should not authoritatively dictate movement strategy, they retained differing degrees of comfort for the impact of AI on decision-making processes.

AI models could aid movements' strategic development in many ways. One benefit could be to facilitate widespread education on nonviolent resistance. The basic theory of change behind existing movement training programs is that activists who learn about nonviolence praxis make better strategic choices.[10] AI tools can implement this theory of change at scale. Models trained on civil resistance literature can help organizers learn about the history and strategy of nonviolent action in a user-friendly format that is interactive and responsive to spontaneous questions. And as AI translation models improve, activists and scholars alike will have greater access to source materials, delivering new perspectives and improving our understanding of neglected cases.

AI models could also be used for scenario planning. A key insight from civil resistance literature is that advance preparation helps movements succeed; pre-movement training helps activists craft strategy, maintain nonviolent discipline, and adapt creatively to changing circumstances.[11] Specially trained AI models could help movements prepare by ingesting strategy documents and then posing hypothetical responses from authoritarian opponents. Furthermore, "red teaming" or other scenario exercises could help organizers consider tactical vulnerabilities or alternative action plans in advance, making them more resilient when they do face inevitable hardships or setbacks. (More ominously, authoritarian opponents may already be using—or will soon use—similar tools to prepare more effective countermeasures against democracy movements. Activists would be wise to investigate what their opponents might be learning from these technologies.)

These innovations hold promise. Yet excessive reliance on AI for strategic development also poses real risks. To start, activists must resist the temptation to let AI agents quickly craft strategy documents or campaign roadmaps. This is a mistake for several reasons. Even well-trained AI models may lack knowledge of key context—movements cannot simply rinse and repeat past approaches; instead, they must adapt nonviolent strategies and tactics to current, local circumstances. Moreover, truncated AI synopses of campaign strategy could lead movements to take a reductive or simplistic approach to nonviolent activism. Democracy movements are complex and multifaceted, and it is misleading to paper over nuances for the sake of parsimonious bullet points—even if the average attention span has

dwindled in the digital age. AI models may offer valuable insights for movement organizers to consider, but there is no substitute for careful and locally informed strategic planning.

Even if movements restrict AI usage to purely information gathering or synthesis, they may still encounter problems with AI bias. AI models' text predictions depend heavily on the information corpus used to train them, and that data can be biased, especially if the training sets are not transparent. Moreover, AI models frequently hallucinate data and have proven to be highly sycophantic, doggedly agreeing with and flattering the user irrespective of the query's substantive or ethical content.[12]

These features not only impede strategic planning but also present deep ethical dilemmas. LLM agents present themselves as confident and authoritative voices on almost any topic, including civil resistance. Yet even in the best cases, AI's purported analytic prowess is an illusion. LLMs predict text, and they do so remarkably well. This enables them to provide relevant information in response to user inputs. However, they do not truly "analyze" contextual movement details—they do not care about their users, nor do they understand activists' lived experiences. Civil resistance can be a high-risk activity, and activists who rely on AI models to make consequential decisions may experience mortal danger as a result. Because AI models are not ethical agents in any meaningful sense, technologists have a corresponding and urgent obligation to carefully consider the ethical harms that AI technologies may pose for their users.

No matter the warnings, some activists will likely over-rely on LLM tools for strategic planning. Because AI has clear use cases for democracy movements, it is all the more critical for practitioners to explicitly identify these trade-offs and encourage activists to use these tools responsibly.

Global Variation in AI Adoption

At this time, democracy movements are not adopting AI uniformly. Instead, workshop participants stressed the need to recognize differences in adoption across movements, particularly across country contexts. Variation in emerging technologies is an established insight in digital repression literature, as the world's authoritarian regimes possess varying levels of digital repressive capabilities.[13] Similar variation is likely to manifest on the democracy movement side as well.

Participants identified at least three aspects of variation in AI adoption to consider. The first and most straightforward is differences in **access** to AI technologies. Some movements will have more extensive and reliable access to cutting-edge tools than others. Internet shutdowns are an obvious concern for cloud-enabled platforms, and studies show that autocratic regimes condition the expansion of public technology infrastructure on their ability to control digital spaces effectively.[14] Some movements may also lack the resources for full-service agentic model licenses that will increasingly dominate the AI landscape.

The second is differences in **need**. No two democracy movements are the same—they face different contextual challenges, have different internal organizing structures, mobilize according to different strategies, and pursue a wide range of substantive goals beyond democratic reform. Movements that rely less on online activism, for instance, may find AI adoption less urgent than movements with primarily online audiences.

A final aspect of variation is differences in **trust**. Some activist communities are more likely than others to embrace new technologies, perhaps based on their issue area, past experiences with new technologies, or other community factors that may prove difficult to observe in advance. Whatever the source, AI optimism and/or hesitance will continue to condition the speed of AI adoption across movements, as well as the types of AI tools that activists choose to use.

Workshop participants concluded that sweeping generalizations about how movements will or should incorporate AI over time must be avoided. Tunnel visioning on early tech adopters (particularly

U.S.-based movements) risks creating false expectations for patterns of AI emergence and usage that do not manifest elsewhere. For this reason, participants sought to take a global perspective on AI for democracy movements, stressing the need to synthesize learning across a wide variety of cases.

Models of AI Adoption

A final area of uncertainty lies in the process through which democracy movements are adopting AI technologies. We see a tension between quick, creative innovation on one hand and deliberate, strategic development on the other.

AI's technological dynamism is naturally encouraging a decentralized trial-and-error approach to mass technology adoption—barriers to entry remain low, and "vibe-coding" makes it possible for everyday users to pick up inexpensive (for now) AI tools and explore at will. Participants acknowledged an early hesitance among democracy activists to engage with AI technologies pushed by Big Tech. Nevertheless, organizers are now experimenting with a broad range of AI tools, from fake chatbot candidates in Belarus to protest limits on electoral opposition to AI newscasters reporting on repression in Venezuela amid President Nicolás Maduro's 2024 crackdown.[15]

This ad hoc, creative approach to AI integration has several important benefits. Perhaps the most significant is speed of adoption. Many activists independently prototyping with the latest AI technologies can help democracy movements stay on the cutting edge of a rapidly evolving field, at a time when movements already lag behind their authoritarian opponents in technological know-how. Even informed observers are still learning what AI platforms can and cannot do, and widespread experimentation is necessary to discover how movements can engage these technologies effectively. More broadly, tactical innovation is a known contributor to successful nonviolent campaigns.[16]

However, a trial-and-error creative model of AI adoption that prioritizes rapid, widespread uptake has potential downsides. One is that movements may invest time and resources into superficial or extraneous AI tools that do not truly advance the core strategic dynamics of nonviolent action. Modern democracy movements face numerous challenges. Among others, movement opponents possess increasingly sophisticated repressive capabilities; severe polarization and fragmented information environments make it difficult to connect with broad audiences and form coalitions; and declining public support for democracy and trust in democratic institutions further diminishes these movements' appeal.[17]

If today's democracy movements are to succeed, they must address these challenges head on. For this reason, our workshop was explicitly problem-oriented—we asked participants to identify core challenges for modern democracy movements and then carefully consider how AI could mitigate them. We hoped this would push practitioners to privilege AI applications that speak directly to the central strategic dynamics of civil resistance. Without this deliberate theoretical exercise linking AI to clear democratic theories of change, AI tools may prove to be "solutions in search of problems," with movements embracing splashy new AI capabilities that provide only cosmetic improvements or distract from core strategic priorities.

A rapid and decentralized adoption model also poses problems for knowledge dissemination. The democracy community currently lacks centralized coordinating hubs for AI information, which means that knowledge about AI experimentation is siloed within individual movements or organizers, leading to duplicated and wasted effort. This approach may also be less conducive to careful impact evaluations, making it harder to determine whether AI tools are having their intended effects.

Moving forward, democracy movements must retain their creative vitality while prioritizing forms of AI adoption that are more closely tied to core movement theories of change. Diffuse innovation is excellent for AI productivity tools that are broadly applicable and require little specialization for

movement work. A diffuse approach to digital education could also help increase basic tech and digital security literacy among activists, which is an urgent priority. However, the democracy community must also increase its capacity to produce specialized AI tools. This requires both advanced technical knowledge and a clear sense of what movements really need; a more centralized tech incubator could help to accomplish this. Overall, it is imperative for movements to better coordinate and share learning about existing AI tools to ensure that valuable technological innovations are widely accessible.

Calls to Action

Perhaps the clearest conclusion from the workshop was that AI for democracy movements remains an early work in progress. Succeeding in this space will require effective collaboration between scholars, practitioners, and technologists. To that end, the NVA Lab is organizing ongoing working groups with workshop participants and other interested stakeholders around three lines of effort. We summarize these activities and objectives below.

AI Code of Conduct

First, participants agreed that the world's democracy movements need a clear set of principles and guardrails for responsible AI use. These technologies give rise to fundamental ethical and strategic dilemmas. For instance: What data are AI models based on, and is that data biased? How much agency should AI models have to shape movement strategy? Who owns AI technology, and how much control does that ownership grant? These questions are further complicated by AI's rapid growth, as both strategic and ethical challenges shift as quickly as the technology itself.

Managing these questions requires a clear set of enduring principles that can guide responsible AI use by pro-democracy movements. The **Code of Conduct Working Group** is taking up this task, beginning by synthesizing core components of existing technology codes of conduct in organizational settings. The group is now drafting distinct codes of conduct for end users, funders, and developers on responsible, human-centered uses of AI for democracy movements. Feedback will be solicited from a wide range of stakeholders in the coming weeks, with a target date to release the first draft of the code of conduct in summer 2026.

Research Agenda

Second, participants agreed that the democracy field requires a substantive research agenda on the effects of AI on democracy movements. Both researchers and practitioners need a clearer picture of the overall landscape of AI adoption among democracy movements, enabling key stakeholders to identify gaps in the field and areas for novel research and innovation. Additionally, practitioners would benefit from rigorous impact evaluations of AI applications that demonstrate whether and under what conditions AI tools actually help democracy movements achieve their goals.

Scholars can meet these needs by investing in a new research agenda on AI technologies and democracy movements. The corresponding **Research and Impact Working Group** is beginning this process by producing a review article that will synthesize existing academic scholarship on AI and democracy movements and lay out areas ripe for additional research. The group then intends to develop a rigorous impact evaluation of AI tools identified as especially promising for democracy movement work, to be conducted in fall and winter 2026.

Movement-Tech Infrastructure

Finally, participants agreed that democracy movements urgently required closer and more organic connections to technologists and technology communities. Technologists have much to offer democracy movements, but they are not well integrated with the pro-democracy ecosystem. Many activists are also

deeply concerned about movement capture or cooptation by monopolistic AI tech giants. It is therefore imperative to build bridges between democracy movements and technologists, forging collaborative linkages to ensure that movement AI adoption is effective, ethical, and self-reliant.

The **Movement-Tech Infrastructure Working Group** is approaching this problem on several fronts. On the technologist side, participants are developing ideas for durable, collaborative institutions where technologists can engage with and support activists. The group specifically encourages movement-oriented technologists and investors to build independent tech stacks, platforms, and applications that reflect principles of responsible use and are accessible to democracy activists. On the activist side, the group focuses on essential AI literacy and information sharing among activist communities. Members are considering how to overcome initial (understandable) tech hesitancy among those who are new to AI. Uniting both approaches is the desire to expand pro-democracy norms among frontier technologists and users, fostering stronger alignment around shared values.

Get Involved

As of March 2026, these three working groups have defined core deliverables and begun their work. We expect to convene again in fall 2026 to share the results of our efforts so far and begin to articulate a collaborative, strategic framework for integrating AI for democracy movements.

That said, the success of our efforts will depend on the usefulness of the outputs to communities of practice. Ultimately, we aim to nurture the emerging community at the technology, democracy, and movement nexus. To that end, we are committed to engaging and incorporating more interested stakeholders into these lines of effort as we progress. We are eager to connect with others interested in contributing to these efforts or those already working in this space. We expect to launch a series of webinars to share interim findings and products with stakeholders and will announce those soon.

Whether via engaging with us or other core actors, democracy activists must take AI technologies seriously. There is an immense amount of work to be done, but decentralized innovation is happening everywhere. Done carefully, democracy movements that embrace AI can capitalize on its affordances while mitigating its vulnerabilities. We hope this report has provided inspiration for activists working on the front lines of global democratic resilience, and we look forward to a future where democracy movements are able to harness this emerging technology to their advantage.

Works Cited

Castells, Manuel. *Networks of Outrage and Hope: Social Movements in the Internet Age*. 1st edition. Polity, 2012.

Cebul, Matthew. *Nonviolent Action Against Democratic Erosion: The United States in Comparative Perspective*. Occasional Paper Series. Harvard University, Ash Center for Democratic Governance and Innovation, 2025. https://ash.harvard.edu/resources/nonviolent-action-against-democratic-erosion-the-united-states-in-comparative-perspective/.

Cebul, Matthew, and Jonathan Pinckney. *Digital Authoritarianism and Nonviolent Action: Challenging the Digital Counterrevolution*. Special Report No. 499. U.S. Institute of Peace, 2021.

Cevallos, Albert. "How Autocrats Weaponize AI—And How to Fight Back." *Journal of Democracy*, March 2025. https://www.journalofdemocracy.org/online-exclusive/how-autocrats-weaponize-ai-and-how-to-fight-back/.

Cheng, Myra, Cinoo Lee, Pranav Khadpe, Sunny Yu, Dyllan Han, and Dan Jurafsky. "Sycophantic AI Decreases Prosocial Intentions and Promotes Dependence." *Science* 391, no. 6792 (2026): eaec8352. https://doi.org/10.1126/science.aec8352.

Chenoweth, Erica. "The Future of Nonviolent Resistance." *Journal of Democracy* 31, no 3 (2020): 69–84. https://doi.org/10.1353/jod.2020.0046.

Chenoweth, Erica, ed. *How AI Can Support Democracy Movements: Summary Report of a Research and Practice Workshop*. Occasional Paper Series. Harvard Kennedy School, Ash Center for Democratic Governance and Innovation,

2025. https://ash.harvard.edu/resources/how-ai-can-support-democracy-movements-summary-report-of-a-research-and-practice-workshop/.

Chenoweth, Erica, and Maria J. Stephan. *Why Civil Resistance Works: The Strategic Logic of Nonviolent Conflict*. Columbia University Press, 2011.

Chenoweth, Erica, and Maria J. Stephan. *The Role of External Support in Nonviolent Campaigns: Poisoned Chalice or Holy Grail?* ICNC Monograph Series. International Center on Nonviolent Conflict, 2021.

Chow, Andrew R. and Chantelle Lee. "How AI Is Being Used to Spread Misinformation—and Counter It—During the L.A. Protests." *Time*, June 12, 2025. https://time.com/7293470/ai-los-angeles-protests-misinformation/.

Diamond, Larry. "Liberation Technology." *Journal of Democracy* 21, no. 3 (2010): 69–83. https://doi.org/10.1353/jod.0.0190.

Earl, Jennifer, Thomas V. Maher, and Jennifer Pan. "The Digital Repression of Social Movements, Protest, and Activism: A Synthetic Review." *Science Advances* 8, no. 10 (2022). https://doi.org/10.1126/sciadv.abl8198.

Feldstein, Steven. *The Rise of Digital Repression: How Technology Is Reshaping Power, Politics, and Resistance*. Oxford University Press, 2021.

Ganz, Marshall. "Resources and Resourcefulness: Strategic Capacity in the Unionization of California Agriculture, 1959-1966." *The American Journal of Sociology* 105, no. 4 (2000): 1003–62.

Hill, Kashmir. "The Secretive Company That Might End Privacy as We Know It." Technology. *New York Times*, January 18, 2020. https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html.

Johnson, Thaddeus L., Natasha N. Johnson, Denise McCurdy, and Michael S. Olajide. "Facial Recognition Systems in Policing and Racial Disparities in Arrests." *Government Information Quarterly* 39, no. 4: 101753. https://doi.org/10.1016/j.giq.2022.101753.

Kirley, Beth, Carl Miller, and Fernanda Campagnucci. *Leveraging AI for Democracy: Civic Innovation in the New Digital Playing Field*. National Endowment for Democracy, 2024.

Loucaides, Darren. "How Governments Are Using Facial Recognition to Crack Down on Protesters." *Rest of World*, March 27, 2024. https://restofworld.org/2024/facial-recognition-government-protest-surveillance/.

McAdam, Doug. "Recruitment to High-Risk Activism: The Case of Freedom Summer." *American Journal of Sociology* 92, no. 1 (1986): 64–90. https://doi.org/10.1086/228463.

Pugh, Jeffrey D. "A Catalyst for Action: Training and Education as Networking Platforms for Peace Projects." *Journal of Peacebuilding & Development* 15, no. 1 (2020): 127–32. https://doi.org/10.1177/1542316619875122.

Schock, Kurt. *Unarmed Insurrections: People Power Movements in Nondemocracies*. University of Minnesota Press, 2005.

Sharp, Gene. *Waging Nonviolent Struggle: 20th Century Practice and 21st Century Potential*. Extending Horizons Books, 2005.

Tufekci, Zeynep. *Twitter and Tear Gas: The Power and Fragility of Networked Protest*. Yale University Press, 2017.

Vadukul, Alex. "Coca-Cola's Holiday Ads Trade the 'Real Thing' for Generative A.I." Style. *New York Times*, November 20, 2024. https://www.nytimes.com/2024/11/20/style/coca-cola-holiday-ads-ai.html.

Vinthagen, Stellan. *A Theory of Nonviolent Action: How Civil Resistance Works*. Zed Books, 2015.

Weidmann, Nils B., and Espen Geelmuyden Rød. *The Internet and Political Protest in Autocracies*. Oxford University Press, 2019.

Notes

1. A more comprehensive AI term glossary can be found at https://github.com/LIANJie-Jason/plain-language-ai-glossary.
2. Matthew Cebul and Jonathan Pinckney, *Digital Authoritarianism and Nonviolent Action: Challenging the Digital Counterrevolution*, Special Report No. 499 (U.S. Institute of Peace, 2021); Steven Feldstein, *The Rise of Digital Repression: How Technology Is Reshaping Power, Politics, and Resistance* (Oxford University Press, 2021); Darren Loucaides, "How Governments Are Using Facial Recognition to Crack Down on Protesters," *Rest of World*, March 27, 2024, https://restofworld.org/2024/facial-recognition-government-protest-surveillance/.

3. Erica Chenoweth, ed., *How AI Can Support Democracy Movements: Summary Report of a Research and Practice Workshop*, Occasional Paper Series (Harvard Kennedy School, Ash Center, 2025), https://ash.harvard.edu/resources/how-ai-can-support-democracy-movements-summary-report-of-a-research-and-practice-workshop/.
4. Erica Chenoweth, "The Future of Nonviolent Resistance," *Journal of Democracy* 31, no. 3 (2020): 69–84, https://doi.org/10.1353/jod.2020.0046; Jennifer Earl, Thomas V. Maher, and Jennifer Pan, "The Digital Repression of Social Movements, Protest, and Activism: A Synthetic Review," *Science Advances* 8, no. 10 (2022), https://doi.org/10.1126/sciadv.abl8198.
5. Kashmir Hill, "The Secretive Company That Might End Privacy as We Know It," *New York Times*, January 18, 2020, https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html; Thaddeus L. Johnson et al., "Facial Recognition Systems in Policing and Racial Disparities in Arrests," *Government Information Quarterly* 39, no. 4 (2022): 101753, https://doi.org/10.1016/j.giq.2022.101753.
6. Erica Chenoweth and Maria J. Stephan, *Why Civil Resistance Works: The Strategic Logic of Nonviolent Conflict* (Columbia University Press, 2011); Marshall Ganz, "Resources and Resourcefulness: Strategic Capacity in the Unionization of California Agriculture, 1959-1966," *American Journal of Sociology* 105, no. 4 (2000): 1003–62; Doug McAdam, "Recruitment to High-Risk Activism: The Case of Freedom Summer," *American Journal of Sociology* 92, no. 1 (1986): 64–90, https://doi.org/10.1086/228463.
7. Manuel Castells, Networks of Outrage and Hope: Social Movements in the Internet Age, 1st ed. (Polity, 2012); Larry Diamond, "Liberation Technology," *Journal of Democracy* 21, no. 3 (2010): 69–83, https://doi.org/10.1353/jod.0.0190; Zeynep Tufekci, *Twitter and Tear Gas: The Power and Fragility of Networked Protest* (Yale University Press, 2017).
8. Alex Vadukul, "Coca-Cola's Holiday Ads Trade the 'Real Thing' for Generative A.I.," *New York Times*, November 20, 2024, https://www.nytimes.com/2024/11/20/style/coca-cola-holiday-ads-ai.html.
9. Andrew R. Chow and Chantelle Lee, "How AI Is Being Used to Spread Misinformation—and Counter It—During the L.A. Protests," *Time*, June 12, 2025, https://time.com/7293470/ai-los-angeles-protests-misinformation/.
10. Organizations like ICNC (now CNCR), the Albert Einstein Institute, CANVAS, and others have long sought to educate global activists about nonviolent methods. For evidence that training programs depend as much on capacity and network building as on information transfer, which would suggest that decentralized AI learning cannot fully substitute for in-person training modules, see Jeffrey D. Pugh, "A Catalyst for Action: Training and Education as Networking Platforms for Peace Projects," *Journal of Peacebuilding & Development* 15, no. 1 (2020): 127–32, https://doi.org/10.1177/1542316619875122.
11. Erica Chenoweth and Maria J. Stephan, *The Role of External Support in Nonviolent Campaigns: Poisoned Chalice or Holy Grail?*, ICNC Monograph Series (International Center on Nonviolent Conflict, 2021); Gene Sharp, *Waging Nonviolent Struggle: 20th Century Practice and 21st Century Potential* (Extending Horizons Books, 2005).
12. Myra Cheng et al., "Sycophantic AI Decreases Prosocial Intentions and Promotes Dependence," *Science* 391, no. 6792 (2026): eaec8352, https://doi.org/10.1126/science.aec8352.
13. Feldstein, *Rise of Digital Repression*.
14. Nils B. Weidmann and Espen Geelmuyden Rød, *The Internet and Political Protest in Autocracies* (Oxford University Press, 2019).
15. Albert Cevallos, "How Autocrats Weaponize AI—And How to Fight Back," *Journal of Democracy*, 2025, https://www.journalofdemocracy.org/online-exclusive/how-autocrats-weaponize-ai-and-how-to-fight-back/; Beth Kirley, Carl Miller, and Fernanda Campagnucci, *Leveraging AI for Democracy: Civic Innovation in the New Digital Playing Field* (National Endowment for Democracy, 2024).
16. Kurt Schock, *Unarmed Insurrections: People Power Movements in Nondemocracies* (University of Minnesota Press, 2005); Stellan Vinthagen, *A Theory of Nonviolent Action: How Civil Resistance Works* (Zed Books, 2015).
17. For a discussion about contemporary challenges for nonviolent movements in eroding democracies, see Matthew Cebul, *Nonviolent Action Against Democratic Erosion: The United States in Comparative Perspective*, Occasional Paper Series (Harvard University, Ash Center for Democratic Governance and Innovation, 2025), https://ash.harvard.edu/resources/nonviolent-action-against-democratic-erosion-the-united-states-in-comparative-perspective/.

A PUBLICATION OF THE

Ash Center for Democratic Governance and Innovation
Harvard Kennedy School
79 John F. Kennedy Street
Cambridge, MA 02138

617-495-0557
ash.harvard.edu

HARVARD Kennedy School
ASH CENTER
for Democratic Governance
and Innovation