MARCH 2026

AI and DEMOCRACY
Perspectives from an Emerging Field

By Michelle Shevin, Kelly Born, and Fiona Kirby

Contents

DISCLOSURES:

This report is based on interviews and research conducted in summer and fall 2025. The interviews were commissioned by The David and Lucile Packard Foundation.

The report draws on interviews conducted under Chatham House Rule with 40 individuals (listed on page 8). Quotes are presented anonymously. Each paragraph-style quote reflects the perspective of a single individual unless otherwise noted. Formatting such as bold or italics has been added in some instances to emphasize key themes or to reflect a speaker's tone.

Organizations and individuals mentioned in this report are not necessarily endorsing its findings or conclusions. While some organizations referenced are grantees of the Packard Foundation, inclusion does not imply endorsement of specific viewpoints, strategies, or recommendations.

One of the authors, Michelle Shevin, is a former employee of the Ford Foundation and currently serves on the board of TechTonic Justice. Two of the authors are current employees of the Packard Foundation.

SUGGESTED CITATION:

Shevin, Michelle, Kelly Born, and Fiona Kirby. *AI and Democracy*. February 2026.

COPYRIGHT:

EXECUTIVE SUMMARY

Beginning in summer 2025, Michelle Shevin and Kelly Born embarked on a qualitative research project to understand how funders, researchers, advocates, technologists, and policy experts working at the intersection of artificial intelligence (AI) and democracy are making sense of related issues. This report synthesizes what we heard.

Rather than seeking narrow definitions, this research aimed for systemic understanding. Participants grappled not only with how AI intersects with democracy, but with how both terms have become contested and expansive. "Democracy" was described as encompassing political rights and equality, economic agency and dignity, and institutional accountability and participation. "AI" was described not only as a set of tools, but as an industry, a political economy, a material infrastructure, and a broader structural shift affecting labor, governance, and geopolitics.

Across these conversations, one clear theme emerged: **AI is not simply a new technology layered onto democracy. It is interacting with political and economic systems already under strain.** Institutional trust is low. Economic inequality remains high. Media systems are fragmented. Government capacity is strained.

In this context, AI acts as an **accelerant** – increasing the speed and scale of change across all of these systems simultaneously.

Where institutions are **capable and accountable,** AI could **strengthen administrative capacity or expand public participation.** Where systems are already under strain, AI may deepen existing weaknesses. The question is thus not *whether* AI will affect democracy, but ***how* its development and deployment will redistribute power within and across democratic institutions, and society at large.**

In some areas, AI may intensify longstanding vulnerabilities. In others, it may introduce new capabilities institutions are not yet equipped to manage. If AI primarily automates discrimination or monopoly behavior, stronger enforcement of *existing civil rights or antitrust laws may suffice*. If it enables new forms of autonomous coordination or persuasion at scale, *new oversight tools may be required*. **Distinguishing between what is intensified versus what is truly new may prove essential for calibrating responses.**

The report is organized around seven core thematic areas that emerged from our conversations, listed below. A separate appendix covers organizations, primarily focused on the United States, working at the intersection of AI and democracy.

1. DEMOCRACY IN THE AGE OF AI

Participants highlight the fact that AI tools are **inseparable from the political and economic context** of the societies into which they are being introduced. Several key themes emerge:

» AI tools cannot be analyzed in isolation from problems of **current politics**, **socioeconomic conditions**, and **societal power dynamics**.

» The political economy of AI (including **corporate concentration of power**, **lobbying influence** from AI companies, and **resource control** – including dominance across the AI stack, from chips, to compute, data, and infrastructure) is arguably as consequential as the technology itself.

» Participants continue to diverge re: **what AI might mean for humanity**: some see it as directly competitive with human agency and democratic decision-making; others view it as a tool for augmenting collective intelligence if properly governed.

2. ECONOMY, LABOR, AND MATERIAL CONDITIONS

If not properly designed and regulated, AI could **accelerate existing economic inequalities**, and introduce new risks to labor, markets, and local communities. Participants highlight four major concerns:

» **Labor disruption:** Economic agency is an essential **pre-condition for democratic and political agency**. Widespread AI adoption may mean that workers face significant long-term displacement (though many note that current data sources are dominated by the companies; credible independent data sources are currently not available).

» **Automating inequality:** AI use in public benefits, policing, and other high-stakes systems may **entrench bias**, **obscure decision-making** and, because of the opacity of systems, **lack avenues for redress** even on consequential decisions.

» **Speculative overinvestment:** Data remain mixed regarding 'how big a deal' AI will prove to be, and thus the likelihood and extent of a more fundamental economic realignment. This has raised real concerns of **significant over-investment in and by the frontier labs** and associated market risks.

» **Market concentration: Six to eight firms** dominate access to chips, compute, data, and infrastructure, deepening power imbalances and weakening accountability.

3. CIVIC LIFE AND THE INFORMATION ECOSYSTEM

AI is **transforming the information environment** that underpins democratic trust and participation. Participants express concern about:

» **Declining trust** in institutions and media, alongside **erosion of shared truth** and **manipulation** through deepfakes, hallucinations, and "AI slop."

» Commercial incentives that **prioritize engagement** and growth over all else.

» **Loneliness**, **disconnection,** and the effects of reduced human contact on civic engagement.

At the same time, some see potential for AI to:

» Support **large-scale deliberation** and **participatory processes**.

» Identify **common ground** across divided communities.

» Improve **information accessibility** and interactive methods to **help people break out of information silos**.[1]

---

1   Note: Outside of our interviews, many have noted the potential for AI to improve journalism. While our interviews included few practicing journalists, efforts like the CalPoly and CalMatters partnership are using AI to track legislative hearings, transcribe them, and even source story ideas.

4. GOVERNMENT USE OF AI

Many participants see AI as both a **necessity and a risk** in modernizing government. Several voice cautious optimism that AI can help manage paralyzing levels of **statutory and administrative complexity**, and dramatically improve government **service delivery** and **customer experience**. Others argue that **efficiency gains** from AI may be the *only* hope for government to regain capacity and legitimacy.

However, achieving such pro-governance outcomes would require *a lot*:

» Strong AI **procurement standards** to guide government technology purchasing.
» Closer, more effective (if guarded) partnership between technologists and public servants to enable **continuous testing** and **iterative, high-integrity experimentation**.
» Significant **government transparency** about all AI deployments in order to retain public trust.
» More **visibility into AI training datasets** and / or development of applications that leverage either "RAG" (retrieval augmented generation) models built to reference government data or models custom-trained for specific government uses.
» Clear **lines of accountability** and **options for redress** when harms occur.
» Clear **restrictions on some government use cases** (e.g., police use of facial recognition technologies).

5. GOVERNANCE OF AI

Acknowledging that interviewees did not include representatives of, for example, the venture capital or AI "accelerationist" community, and included only a few participants aligned with the current federal administration, interviewees broadly agreed that **AI governance is necessary**. While they **diverged on exactly *what* guardrails** should look like, there was strong agreement on three dimensions including:

» Robust **privacy protections**.
» **Transparency and accountability** across the AI stack: from data center development and energy usage, to how models are designed and red-teamed, how and which datasets are used, where and how models are deployed, how individual users are engaging with AI, and more.
» Access to justice and **meaningful appeals** mechanisms.

Finally, we heard broad agreement that both governance and deployment decisions, particularly regarding public sector use, must be made with **robust public and community participation**.

Debates persist around:

» Which, if any, **additional regulations** (e.g., liability, antitrust, trade policy, kill switches, or universal basic income) are desirable.
» Where **legislative authority** should reside, with most agreeing that federal action would be strongly preferred, but remains unlikely at present.
» Whether, given the likelihood of ongoing state policy experimentation, state lawmakers have **sufficient access to the inputs necessary** to keep up with the rapidly evolving AI landscape, or **sufficient oversight capacity** to ensure laws are enforced.

6. NATIONAL SECURITY AND SYSTEMIC RISK

AI is reshaping **national security**, **global competition**, and the **balance of power** between democratic and authoritarian governments. Participants raised concerns including:

» Expanding use of AI-powered **surveillance**, particularly in the context of immigration enforcement.
» Military, private, and **asymmetrical uses of AI** and **autonomous weapons**, including use of AI by **nonstate actors** to enact **political violence**.

» Intensifying **geopolitical competition**, including concerns that whichever country attains AI dominance will secure coinciding military and economic dominance, and, related, concerns about supplying AI tools and capabilities to competitive or nondemocratic regimes.

Three priority needs emerged:

» Standardized **risk measurement frameworks** that are shared and translatable across private and government sectors, and across nation-states.
» Improved **global coordination** grounded in democratic values and cooperation rather than arms race competition.
» Clearer policy tools to mitigate national security and foreign interference risks, as experts currently **lack consensus** on which policy levers (e.g., export controls) to use to address risks.

7. IMPLICATIONS FOR PHILANTHROPY

Given this moment of profound democratic instability and accelerated technological and social change, philanthropy requires greater coordination, experimentation, and investment in democratic infrastructure. Participants emphasized the need to:

» Improve **field coordination** by bringing together experts that don't typically work together, building a shared field vocabulary, and developing shared field infrastructure or "public goods" (listservs, etc.).
» Improve **relationships and coordination across funders** that align on the need for regulation (e.g., nearer-term and longer-term risk funders), by focusing on *shared goals and solutions*, rather than on differences in opinion regarding *specific harms*.
» Invest in **measurement capacity,** including efforts to understand harms, to measure security risks, and to understand the economic and social impacts of AI.
» Invest in **foresight capacity**, both efforts to forecast potential futures and efforts to rebuild foresight mechanisms, so that they can be more responsive to the needs of the current moment.
» For those funders investing in **civic tech / positive use cases** for AI: **better coordinate capital across the technological lifecycle** by providing support not just for early-stage ideas but for "getting to scale."

Across these seven areas, **three cross-cutting priorities stand out**:

1. The importance of **building capacity in the institutions that will be absorbing these new technologies;**
2. The **need for independent measurement** as an essential precursor to meaningful monitoring, regulation, and accountability; and
3. The need to **confront the political economy of AI directly** – addressing who owns and controls compute, data, and infrastructure; whose interests are prioritized in policy and governance; and whether existing competition, labor, and corporate governance frameworks are sufficient.

Throughout our conversations, participants emphasized that neither technological fatalism nor technological optimism is particularly helpful. AI's democratic impact will depend less on abstract debates about alleged inevitability, and more on concrete choices about policy design and enforcement, institutional strengthening, coalition-building, and field capacity.

The central question that emerges is not whether AI matters for democracy. It is whether institutions, funders, and civil society actors will adapt their strategies, tactics, and relationships quickly and concretely enough – through measurement, enforcement, capacity-building, and coalition development – to ensure that accelerated technological change does not outpace democratic governance.

Note: These interviews were conducted in the summer and fall of 2025. Since these interviews were conducted, AI systems have continued to improve rapidly – particularly in areas like code generation, tool use, and semi-autonomous "agentic" task execution. These advances reinforce many of the concerns and opportunities identified in this report. As agentic systems move from generating outputs to executing multi-step tasks across digital environments, questions of delegated decision-making, oversight, measurement, privacy, accountability, and institutional capacity become more urgent.

Acknowledgments and List of Research Participants

1. DEMOCRACY IN THE AGE OF AI

What Does "Democracy" Mean Now?

Across 40 interviews, democracy emerged as a contested and evolving concept. Funders and experts grappled with the scope, meaning, and implications of the term. While definitions varied, interviewees' reflections clustered around **three interdependent dimensions of democracy**:

» First, **political equality and rights**: free and fair elections; preservation of civil liberties, from freedom of speech to protection against government overreach; and protections for minority groups.

» Second, **economic agency and material security**: whether individuals have sufficient economic security and dignity to be able to participate meaningfully in civic life.

» Third, **institutional accountability and participation**: the capacity of public institutions to deliver effectively and transparently on their promises, and to respond meaningfully to public input.

Throughout this report, references to "democracy" implicitly engage one or more of these dimensions.

In response to questions about how AI will impact democracy, several participants sought to disentangle whether they should focus their remarks on one or more of an overlapping set of ideas ("liberal democracy," "American democracy," "self-governance," or a "constitutional republic"). A few explicitly sought to address the common misconception that 'working on democracy' implies a core or even narrow focus on elections.

"What do we mean when we talk about democracy? **Democracy is elections**. **Democracy is voting**. Democracy is also having a **society and a populace that is able to participate**, not just removing barriers [to voter] registration – [though] all those things are important – but [also] ensuring that people feel **comfortable and free to speak**."

Some participants come reluctantly to the term itself: "I don't think 'democracy' appears on our website. In fact, we've had a conversation internally about, is it 'democracy' that we're building? Or is it just the ability of people to govern themselves with whatever choice of governance they want? I think even using the word 'democracy'... doesn't sit quite right... it's all about **helping communities solve their problems and accomplish their goals**."

Others agreed on the need for much clearer definitions: "I don't care about a shared understanding of 'AI.' I'm happy to let that be a little floaty, but... what do we actually mean [by 'democracy']? I work with a lot of people who do elections and voting, and when they think democracy, they think [elections and voting]... I think **we need a shared understanding of democracy that sees individual rights as part of a collective right**. I think that's really important, and I think it would be useful if funders can coalesce around something like that."

Regardless of definitions, most if not all agreed about the significant tensions AI poses for liberalism, particularly because "small L" liberalism requires protecting individuality and minority rights, while AI often glosses over individuality in favor of "the average."

"AI can undermine pluralism. There is real concern of homogenization, of AI acting like an averaging machine which could have important cultural and political implications... For a lot of people the norm is just, 'Hey, if we want equality, let's just be colorblind, and treat everybody the same.' That's their notion of baseline fairness – that assimilation is good, and if things are the same for everybody, that's what equality is. But it's a little deeper and richer than that. In a liberal democracy, **allowing people to associate, and have their ideas, their religion, and culture – that itself has value**. AI could be used as a tool to further advance cultural conquest that we've seen in the past. In the same token, AI could be more cost effective than a lot of tools for language access and translation, and other applications could help preserve culture and advance pluralism. I know that English is dominant among language models and that many have shortcomings with regard to underrepresented languages, but compared to everything else we have right now, there's real potential for **language preservation, language translation**, cross-community coalition building, and other uses that advance pluralism. Appreciating the potential of AI to facilitate an inclusive America is essential."

AI's Impact on Democracy – How Big Will It Be?

"The idea of liberal democracy, 'small R' republicanism, etcetera... all of that is actually the result of a quite fragile technological contingency. And it may well break."

We need a shared understanding of democracy that sees individual rights as part of a collective right

"Democracy has to be born anew every generation" wrote American philosopher John Dewey in 1916, speaking to the then-nascent needs of "an industrial democracy."[2] For several of our participants, "AI" is a manifestation of a long-running tradition, in which **emerging technologies prefigure what becomes institutionally possible** and necessary for human societies.

Participants broadly agreed that these are **times of transformation**, without consensus on what changes may emerge, and what they will require.

"The Enlightenment-era [of] nation-states seems to be at its twilight, and it's really hard to know what comes next, but **what seems implausible to me is that nothing changes fundamentally**… When you look at other major technological transitions, it seems to me that every single one of them has preceded an institutional transition, so my expectation is that over the next, say, 20 years, **essentially every institution that predates AI will cease to exist**."

"We're in a moment where **we're going to have to reimagine everything** – we have to reimagine our entire democracy, because what we had before is gone. And it's not coming back… I think we just need a totally **different set of regulatory frameworks**. And they have to be global."

"[Many including me are] making a bet… that we have to simply reform the institutions as they exist, without a value judgment on whether that's the right thing or not… [My view is] that the incremental reforms that we have been fighting for since the entire time I've been doing this work are now very clearly just not even f—ing worth it. There's no point. **We should be going for transformational change – a shift to a fundamentally different operating model**."

"**Our operating model of government has been left in the industrial era**. We didn't quite update it for the internet era. We sort of threw some websites on the front end of a very industrial-era model, and limped along. So now we kind of have to **leapfrog government into the AI era**, and that's very awkward to do, given that there's just so much we don't know about what will be demanded."

Several experts found it hard to imagine a world where human agency competes meaningfully with machines and financial interests.

"Democracy is… a social contract based on a balance of power enabled between the masses and the [state]… one of the things that's **[hard to imagine are] reasonably good futures** in a world where humans are centered in decision making in any meaningful way, **and not just AI systems and market incentives [and] very small numbers of humans**."

Rather than understanding "democracy" and "AI" as disparate spheres intersecting each other in discrete ways, many participants emphasized their deep entanglement. Some participants even framed the concepts as **directly competitive philosophies for decision making**:

"So **one [democracy] is about channeling agency and the other [AI] is about proxying it** – both vying for control

"So **one [democracy] is about channeling agency and the other [AI] is about proxying it** – both vying for control of power and systems that define reality. The intersection, then, is more of a head-on collision at the end of the spectrum of human self-determination."

---

2   Dewey, John. *The Public and Its Problems*. Facsim. ed. Athens (Ohio): Swallow press-Ohio university press, 1991.

of power and systems that define reality. The intersection, then, is more of a head-on collision at the end of the spectrum of human self-determination."

What is "AI," Anyway?

Participants often used the term "AI" to describe different things within the same conversation. At times, they were referring to **specific technological tools embedded in decision-making systems** – hiring software, benefits eligibility systems, content moderation tools, or regulatory review assistants. In other cases, they were describing **the AI industry itself** – a concentrated set of six to eight firms that control data, compute, material infrastructure, and significant political influence.

And at a broader level, some were pointing **to AI as a structural and cultural shift affecting labor markets, geopolitics, information ecosystems, and institutional design.** Because of this, concerns about AI's implications for democracy extend far beyond questions of immediate deployment risk, into issues of corporate power, resource allocation, and political participation, continuing long-standing democratic debates regarding **influence**, **accountability**, and **structural inequality**.

The most coherent definitions and concerns we heard about "AI" emphasized power dynamics, for example: "'AI' is a marketing term that touches on multiple different practices, tools, and approaches to **data-intensive analysis, predictions, and surveillance that serve industry and power**."

For many interviewees, AI is not, strictly speaking, a technology, nor merely a collection of associated technological capabilities in which computers are used to approximate, augment, or automate tasks that historically require human intelligence. Here, like "democracy," "AI" contains multitudes: it is a tech stack, a set of material infrastructures, a powerful industry, an economic and social force, a story about technology, and more.

"[By 'AI,' early scholars] meant an interdisciplinary subfield where you apply technical tools to simulate intelligence. [Today,] you can say 'artificial intelligence' and be describing **an ecosystem of models, algorithmic techniques, energy infrastructure, thought processes, companies**, and so on. So when somebody says 'the American economy is resembling a giant bid on artificial intelligence,' they don't mean specific analytical techniques, they mean that [eco]system.'"

Regardless of definitions, this document spans references to **three primary, practical applications.** First, **AI to automate or assist in decision-making** – in hiring, benefits eligibility, policing, lending, content moderation, and regulatory review. Second, **AI to rank and predict** – generating risk scores, recommendations, and classifications that shape who receives resources, visibility, or scrutiny. Third, **AI (increasingly) as an actor** – through "agentic" systems that can execute tasks, interact autonomously, or coordinate activity at scale. **In each case, AI does not simply generate information; it influences real-world outcomes.**

These **uses matter because they alter how power is exercised.** AI systems may replace or mediate human judgment, compress timelines for decision-making, and shift authority away from frontline actors or democratic participants toward profoundly opaque, centralized systems.

But **definitional ambiguity matters less than practical impact.** Whether one defines AI as a tool, an industry, or a structural shift, its **democratic implications hinge on how it redistributes decision-making authority, accelerates economic restructuring, and concentrates or diffuses institutional power.** The technology itself may be new, but the underlying democratic questions – who decides, who benefits, who bears risk, and who can challenge outcomes – remain familiar.

To that end, many participants explicitly referenced AI's "**political economy**," calling attention to the **market forces, economic systems, and power dynamics** that surround AI as a system: "I wish that more philanthropy was focused on the political economy of AI, and the way that it operates in a **system of concentrated corporate power** that is [now] working even more closely hand in hand with government austerity or surveillance imperatives in an increasingly authoritarian context in the US. And globally."

As one expert noted: "One of the prevalent debates... is **whether a rising tide will lift all ships, or only [those of] the people who already own ships**. If AI follows the pattern of every technology that precedes it, **those with power, access, and resources will learn to leverage AI in decisions** – whether they're directly involved or not – to maintain the status quo that benefits them and out-compete those without those advantages. This trend isn't new – it's called The Matthew Effect in the Bible, which recognizes that **without extraordinary intervention, advantages accumulate to the already advantaged.** Or more simply, the rich get richer – and there's nothing in AI's architecture, ownership, or politics that suggest such an extraordinary effort... [The problem here is] less about 'bias,' and more about the ability to impose systems with known errors anyway."

Concerns about the political economy of AI – about concentrated corporate power, control over infrastructure, and the risks that presents – appeared *everywhere* in our conversations, across nearly every thematic area discussed below.

The growing relevance of AI to society as a whole may be **attracting new and growing numbers of stakeholders** to a conversation that has for many years been somewhat more narrowly focused on tech policy.

> Concerns about the political economy of AI – about concentrated corporate power, control over infrastructure, and the risks that presents – appeared everywhere in our conversations, across nearly every thematic area discussed below.

"There's something interesting happening in the data center conversation, where we're **getting very uncommon constituencies to finally care about tech policy**. You would imagine that would have happened in surveillance land, where it was like, 'Oh my gosh, aren't we all afraid that we're all being watched by our phones and our TVs and our data is leaking everywhere?' But [we often heard] 'It's too late anyway, we're always being watched.' It's a collective action problem. But, if there is a real, localized threat that people can perceive... if suddenly there's a nuclear plant being built [in my neighborhood] because of [AI companies]... it has really allowed an opportunity for AI policy to become everybody's problem. A lot of **the environmental groups** have been on [AI policy] for a while because of the water and energy questions. But increasingly, there's **new labor considerations** that are becoming a part of this. I was just talking to a bunch of architects and urban planners about their role in [AI dynamics], and... [we are **finally seeing AI] as a reorganization of our economy, and our sectors, and our relationship to each other**; [appreciating what AI] is actually doing to our society writ large."

Some participants noted that the immediacy of AI's **material infrastructure concerns** sometimes appear to be **eclipsing conversations around the sociotechnical implications** of AI, including foundational questions about data privacy and due process.

"There's obviously a lot of discussion about the collection of information for these AI models, but it feels like that discussion's almost over. People were talking about it, and now we've sort of moved on to data centers and how much energy they use... It's sort of like, 'Oh, that's done.' But I don't think [it] can ever be done, right? Because there's obviously **ongoing data collection**, and I think it's really, really important that we think about AI as we're thinking about **data privacy**."

Others underscored that in some ways, we have come full circle from a conversation about the potential for open-source, democracy-enabling, human-rights respecting AI back to conversations about "traditional" democratic problems related to lobbying, land use, and power.

"In the early days of discussions around AI tools, machine learning, neural networks, there was a **vision that was more cooperatively owned**. We were thinking about open sourcing some of these things, there was real energy around building together and thinking more creatively. But now that the **strategy around AI is so corporate heavy**, and is so much about industrial policy, energy demands, shifting water resources… I think **the questions about democracy are becoming much more traditional democratic questions**. 'A big AI company is coming to my town, and they have lobbied all of city council, they're providing all these narratives about what AI is supposed to do for the future, and it has really dissolved my ability to be a participatory decision maker in my city council processes, and my decisions about land use, and zoning…' This AI race [has] really shifted the scope of AI and democracy questions from, 'Isn't it bad that mis- and disinformation is going to destroy our elections?' – which is still an issue for me, but… because the scope of AI is becoming ever more material, I think we're coming back to traditional democratic problems **around resource sharing, budgets, and local power**."

Nearly all of our participants referenced some version of this critical point: **AI tools cannot be separated from the politics, socioeconomic conditions, and power dynamics of the societies they are being introduced into.**

"I sometimes **wonder if centering AI specifically, and technology more broadly… ends up displacing the focus that should be on underlying issues and structural dynamics**… Technology provides a way for us to reflect and work on **complex social problems that ultimately [are] amplified, or are complicated by the use and development of technology**. So, one of the things that has become really pressing in the work that I do is the **role of corporate and billionaire influence**. [While] capture is not a new issue… [it's important to look at] the **ways in which technology is interacting with other phenomena, [including] socioeconomic phenomena, to create the problems that we see**."

A recurring tension in these conversations concerned whether AI represents a fundamentally new democratic challenge, or whether it primarily intensifies problems we already understand. **Many of the risks identified are not new: corporate capture, labor exploitation, biased decision-making, and weakened journalistic systems remain longstanding democratic vulnerabilities. AI threatens to accelerate and scale these existing pressures.**

Yet some capabilities appear meaningfully different. Real-time **personalized persuasion at massive scale, autonomous agents operating with limited human oversight, and the potential for rapid, widespread economic displacement could compress changes that once unfolded over decades into just a few years**. The speed, coordination, and scale enabled by AI systems may strain institutions that were not designed to operate at machine speed.

Some participants emphasized that gaining clarity about what is genuinely new – versus what is an intensified version of an old challenge – can help determine **when we simply need *stronger enforcement of existing laws*, when *entirely new legal frameworks* may be needed, and when the real need is to *build institutional capacity*** – in government, journalism, or civil society – to withstand and adapt to these changes.

Across nearly every thematic area, concerns about concentrated corporate power and resource control surfaced as structural drivers of risk.

2. ECONOMY, LABOR, AND MATERIAL CONDITIONS

Four main concerns were raised about the potential economic implications of AI:

1. Labor market disruption, including risks of unemployment and the erosion of worker's economic and, thus, political agency.
2. Second, **"Automating inequality"** as a result of government or corporate deployment of AI systems that are presumed to be inherently biased because they were trained on a corpus of (biased) data.
3. Third, growing risks of an **economic bubble and potential subsequent market crash** driven by large investments of speculative capital.
4. Fourth, the **concentration of wealth and power** in the owners of a handful of dominant firms as they secure market control and (presumably) push to weaken regulatory protections and extract maximum value from consumers, including vulnerable populations.

Labor and Unemployment Effects

Participants discussed the potentially significant impacts AI may have on employment and hiring, the subsequent **implications widespread unemployment would have for democratic participation,** and the **lack of objective data** about likely and actual labor market impacts. Interestingly, these labor market impacts of AI were seldom independently introduced by participants. When introduced, many interviewees expressed some alarm.

Several noted a continuing pattern of encroachment, in which layers of algorithmic management and worker surveillance gradually erode agency and autonomy at work. Particularly, how the **loss of worker agency** under AI-driven management could contribute to a deeply stratified society, which would in turn have

destabilizing consequences for democracy. They also noted that the labor debate has expanded to **include knowledge professions**.

A few participants sought nuance and pragmatism regarding a conversation that too often prioritizes *automation* of human labor over *augmentation*. One participant argued for increased attention to "human machine teaming," including research to understand and calibrate human-machine trust.

Overall, participants appreciated that **democracy requires empowered constituents**, and **if workers are disempowered, they will not be able to engage.**

Multiple participants expressed frustration regarding the relative **dearth of rigorous data** regarding AI's labor implications as compared to a surplus of corporate mythologizing. Participants commonly named AI companies and major consulting shops as notorious producers of potentially biased analyses and forecasts.

"Anthropic is, in the guise of, 'we're doing everything good for the world,' essentially making a play to capture the measurement conversation around AI's impact on the economy. On one hand, you could be like, 'that's so great, they're helping to foster this.' On the other, if you read between the lines, **they are looking at the data coming from their platform. And they want other researchers to look at data from their platform**. And then based on that, we'll figure out how the economy's going related to AI... I think there's something there about seeding the terms of debate."

Many participants spoke to the astronomical investments being made in the buildout of AI companies, and infrastructure.

Others noted that narratives portraying AI as inevitable and superior to human labor are taking hold at a time of weak labor protections and ongoing cuts to public support. Together, these dynamics **increase job insecurity** and the potential need for **a strong safety net** to address potentially **widespread economic insecurity**.

"Venture capital appears to be creating inflated funds in order to drive up the cost of AI when AI is not a proven technology. So the first thing I think about is the **potential of an economic bubble, a crash, and a bank bailout that feels imminent to me in the next three to five years...** And now **we have no safety net**, because [austerity measures and individuals have] gutted the entire federal [apparatus to] provide any kind of safety net to people who would be directly harmed."

Exacerbating Inequality through Deployment of AI in Public and Private Sectors

Participants discussed how **AI might embed or amplify bias,** particularly when deployed for *decision-making*. AI is landing amid efforts to realize a pluralistic and multiracial democracy within a **multigenerational struggle for civil rights and racial equality**, alongside growing **political, educational, geographic, and religious fractures and increasing economic inequality.**

"The **material precarity of people's lives undermines democracy**... We know poor people vote less. We know poor people generally participate less in civic meetings, etcetera. So that's not a new phenomenon. But... **AI intensifies people's material precarity** in ways that are just an order of magnitude more severe than what we've seen. I think that really **threatens to permanently disenfranchise poor people from the political process**, and I mean 'disenfranchise' both in terms of voting, and in terms of other forms of civic engagement and participation that are also crucial to democracy."

Other participants cautioned against allowing public momentum around newer AI problems to drive research and advocacy priorities, when **old problems remain unsolved**.

"To me, **the [algorithmic] bias discussion is in its infancy with regard to legal solutions**. We also have an attempted dismantling of disparate impact[3], which is one of the more effective tools in identifying AI tools that produce unjustified discriminatory outcomes – certainly more effective than tools that only look for discriminatory intent.... **We have to examine the AI bias issue not just in the context of evolving technology, but also in the context of the changing nature of law with regard to discrimination**."

The opaque nature of AI systems **limits visibility** into decision-making processes, creating a misleading appearance of objectivity and rendering opportunities for redress more difficult.

"I take a generally broad understanding of AI, because to the people who are subjected to it, a lot of times, the technical sophistication doesn't matter. The problems are oftentimes the same. **You don't know it's being used**. If you do know it's being used, **you don't know how it works**. You can't fight it, practically – it's hidden. It has the **allure of objectivity**... I focus mostly on AI based decision-making, the point at which it is changing somebody's life actively. And the closer to that decision-making point we get, the more concerned I am."

Many participants named socioeconomic inequality as a key reason AI affects populations unevenly. Several mentioned *Automating Inequality*, a 2018 book by Virginia Eubanks[4] which detailed efforts to automate parts of public systems through integrating administrative data held by government agencies, as foundational to their understanding of AI and democracy intersections. Several participants also amplified recent work looking specifically at how AI systems impact low-income Americans.

One expert suggested that some current concerns may be unfounded, and are instead rooted in *past* failures in government deployment that were the result of misapplications or older technologies, rather than newer processes like retrieval augmented generation (RAG)[5] that allow greater control over the sources referenced by LLMs.

"I hear much, much less [these days] about... bias. I think as people have realized that this isn't about turning over things like eligibility decisions, for instance, to an automated system – in fact, I think those fears [about bias in government technology] were driven more by this **incredibly poorly done automation that occurred in what I consider sort of a pre-AI era.** And I maintain that as just badly done [decisions about when to use AI]. If the rules governing eligibility for a program are deterministic, **you not only don't want AI, you don't need it**. You just need good deterministic software to be compliant with the rules."

Americans in need of public assistance find that eligibility, service delivery, coverage, and support are increasingly algorithmically intermediated in ways that **too often privilege financial interests** over duty of care, or that automate existing biases within systems while **reducing transparency and contestability**. Ensuing failures are **rarely met with redress** as institutions charged with oversight, governance, and accountability seem to err toward prioritizing national technological supremacy over public care.

---

3   Disparate impact analysis identifies disparities in outcomes, assesses whether they are justified by legitimate objectives, and intervenes when they are not. Even when race, gender, and other protected classifications are excluded as variables, AI systems frequently rely on proxies – education, occupation, geography, social networks – that replicate disparities. Because AI systems typically operate without explicit discriminatory intent, eliminating disparate impact analysis reduces the ability to detect and address discriminatory outcomes produced by increasingly automated decision-making.

4   Eubanks, Virginia. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. First Picador edition. New York: Picador St. Martin's Press, 2019.

5   "RAG" or retrieval augmented generation is an AI approach that improves an LLM's answers by first retrieving relevant information from external sources (such as documents, databases, or trusted websites) and then using that material as context for the model's response, which can increase accuracy and help reduce hallucinations and some forms of bias by grounding outputs in identifiable evidence.

One participant noted: "I feel really nervous about… procurement [of unproven AI technologies] that are being used within states to streamline everything from benefits assessment, to handing out benefits to working class and poor people… **[I worry about] the kinds of potential mistakes that could be caused**, and the real world harms [that have] already been demonstrated. But now I think there's even more incentive for states to adopt AI, and [this] AI technology is terrible."

Participants emphasized that AI differentially impacts low-income Americans in multiple ways across sectors and systems, commonly described as some version of a 'perfect storm.' Here, a confluence of incentives and feedback loops have created a landscape of **bureaucratic frustration** and **undignified abuses** of people without ready access to justice.

"**People living in poverty were essentially the frontier of experimentation with automation**. [Including] precarious workers like delivery workers, rideshare drivers – they were [some of] the first people on which algorithmic management was tested on and refined on. And… the ways in which these abuses unfolded, essentially made human rights law, but also labor laws and benefits protections, outmoded, and, in some senses irrelevant, because **there was no way for democratic institutions** to really **understand how to deal with these abuses**, or there wasn't really any interest in dealing with these abuses, because the political interest was very much in this idea… [focused on] innovation and efficiency and moving forward."

Market Concentration and Political Power

Others expressed concerns that only six to eight companies have the financial resources to play in this game – even traditional venture capital firms are priced out of investing in frontier models and are instead investing at the application level.

This concentration of economic power in the hands of an increasingly small group of elite tech company owners risks its own set of democratic distortions. Powerful tech companies now spend massive lobbying dollars and exert concerted influence on public sector processes (Meta alone spent over $24M on lobbying in 2024[6]).

In addition to traditional concerns about monopoly power, participants emphasized that **AI, when used for decision-making, risks further undermining democratic rights and governance due** *both* **to society's inability to measure or monitor whose interests are prioritized (and why),** *and* **citizens' inability to influence these decisions**. Emphasizing that "the ability to influence decisions, even partially, across a wide range of domains is… the goal of politics," one participant argued that "when AI is intermediated in decision making, there's no way for users to measure / monitor whose interests are prioritized and why or how." Therefore, **AI adoption can become an "anti-participatory source of political power itself**."

Participants also observed ways in which political dynamics constrain and shape corporate priorities.

"[AI companies] react to political pressure. There was a lot of **political pressure, especially in 2024, to look like they were doing something in the election space**, and if anything, **right now they're feeling the reverse**: worried that they're going to be called part of a 'censorship industrial complex.' [Now AI companies seem] more focused on things like wellbeing and mental health, and [are] not talking as much about democracy or elections. **Democracy's in retreat around the world, [and] I don't think that's a coincidence**… [I'm very worried that] the information environment [for the next ten years] is getting shaped now, and the incentives for these private companies is their own interests, [including] not being on the wrong side of where political power is."

---

6   https://www.opensecrets.org/orgs/meta/summary?id=D000033563

3. CIVIC LIFE AND THE INFORMATION ECOSYSTEM

Participants described a range of possible futures for AI's role in democratic participation that spanned both promise and risk. They raised concerns about **declining trust and media credibility**, while also considering how AI might enable more **inclusive forms of civic engagement**. Many cautioned, however, that without careful design and governance, AI could **intensify manipulation and isolation** while **reducing participation**, weakening the social cohesion and institutional legitimacy essential to democracy.

Media, Trust, and Truth

"AI has raised the question, 'How can you trust anything you see?'"

Impacts of AI on the information ecosystem that undergirds American democracy were some of the topics most frequently raised by participants, **including acute concerns about deepfakes, hallucinations, and convincing deceptions**, alongside a general atmosphere of doubt and distrust.

"The destruction of the concept of truth is so fundamental to the harms [of AI] that I don't think you can ignore it… the introduction of AI, whether or not it's used in journalism… doesn't matter, because it's polluted the environment such that **no one believes anything anymore**."

Alluding to what is sometimes called 'the liar's dividend,' participants described a tide of "AI slopaganda" that **deteriorates the capacity for shared reality, and therefore one of the fundamental preconditions of democracy – access to shared information and facts.**

"There's now so much noise out there, and at least a certain swath of the population has been trained to be skeptical, which is great, but there's this thing that can happen where… you're just like, 'Well, now I just can't believe anything… I'm just going to be a nihilist and not believe anything.' And that's also a bad outcome."

AI deployment is also happening amid declining trust in institutions. Many participants pointed to a feedback loop of tech ownership in a consolidating media landscape as detrimental to democracy.

"Having an **adversarial free press is absolutely crucial to democracy**. One of the challenges is that we still think that corporate media could be that, and it really isn't. These are companies, or billionaires, with a lot of interests in front of the government other than their journalism, and we... are seeing that they're not going to fight [to hold the government accountable]. That means there has to be a **new cohort of media that arises out of this period of destruction."**

Participants shared acute concerns regarding a **tide of 'AI slop'** and persistently untrue information flowing through and out of AI mediated information platforms.

A few interviewees described AI companies **abdicating responsibility** for addressing obvious problems with election disinformation on their platforms in the lead-up to the 2024 election.[7] One noted: "There should be some sort of **election benchmarking**. Somebody should be doing monthly testing on things that are important to our democracy, not just elections, but also facts about how our government [works]."

The increasing popularity of generative AI tools suggests further shifts in how information is being consumed and integrated. One researcher found that what kept people engaged with AI chat agents "wasn't anything about legitimacy or trustworthiness – the veracity of the content. The thing that was bringing people back was that **they had agency**. **When they used ChatGPT, it was responsive to them**. They... could steer it."

The question then becomes: if AI is going to continue to impact an already challenging information environment, what can be done about it?

At least one participant regrets how the framing of "disinformation" implies a free speech issue when the **concern is actually incentives**. Focusing on whether a particular piece of information is true ignores tech companies' profit motive to accelerate negative information dynamics. Another reported a complete reframe, sharing, "A lot of focus has been on misinformation and an attempt to squash it out, and my current focus is on the opposite: **How do you promote the good stuff**?"

Civic Participation and Collective Intelligence

Participants voiced cautious curiosity about how AI might strengthen democratic participation through helping to **elevate citizen voices**, cultivate **productive discourse**, and show people **areas of commonality** rather than difference.

Several participants reported curiosity about the potential of AI to reinvigorate participation processes. One emphasized research regarding "how you [could] use AI and technology to help **elevate citizen voice**... [including] sensemaking... and mechanisms for consensus and deliberation." Some see potential for **innovation in polling**; others see promise for LLMs to **distill political viewpoints** in the midst of less structured public engagement processes. We heard many mentions of vTaiwan[8] and Audrey Tang, the groundbreaking Taiwanese activist and digital affairs minister, and about efforts in computational "deliberative democracy: inviting communities of all backgrounds to engage in meaningful, informed discussion about... [a wide range of topics]."[9]

---

7   Palta, Rina, Julia Angwin, and Alondra Nelson. "How We Tested Leading AI Models Performance on Election Queries." Proof News, February 27, 2024. https://www.proofnews.org/how-we-tested-leading-ai-models-performance-on-election-queries.

8   vTaiwan: a decentralized process that uses in-person and online interactions to bring together citizens and government officials in Taiwan to deliberate on national issues.

9   Siu, Alice. "Let the Public Govern AI." Tech Policy Press, September 5, 2025. https://techpolicy.press/let-the-public-govern-ai.

Several expressed enthusiasm for how **AI could enable broader participation that might help to invigorate the democratic fabric of communities**, from public comment periods to city hall deliberations, helping to cultivate more productive democratic discourse.

"There is an opportunity for AI to reflect back to us the collective, and our place in it... in a really dynamic, iterative, nuanced way: **Here's what all of America thinks about this topic, and here's where you fit in the map of other people**. And [maybe] you can have an iterative conversation with the chatbot about [e.g.] what other Americans said about equality... what [those who disagree with you] actually think, and why. That starts to get at this challenge of siloing, the feeling of disconnection, and this narrative that, 'it's okay for me to hang out in my information silo because we're so divided anyway that all I would hear if I turned on CNN is just disagreement and screaming conflict.' **We... need to find new channels to reflect back to people how much social cohesion and common ground there [actually] is."**

Others were even more optimistic:

"Right now, social media algorithms reinforce... biases and make [people] more confident in wrong things [that] they previously believed. **Can we train AI systems to cause people to change their minds towards more true beliefs against their biases?**... There's a bunch of stuff that could be done on AI for consensus building... The first thing is AI for truth tracking, getting on the same page about what's real. Second thing is more stuff like scaling up vTaiwan, [an open process for deliberation on national issues]. People have different preferences... some people really prioritize feeling safe, some people really prioritize being inclusive... Can we get AI systems that can talk to everyone at once, and come up with a proposal, and then explain that proposal to everyone, and get everyone's buy-in into that proposal? Can we pilot it at the local level?"

At least one participant had ambitious ideas for direct, rather than representative, democracy.

"Maybe if we lived in a democracy, we could **direct vote via text message** once a day on an issue or something... I think we've just been digitizing and automating our old system of democracy... if we actually want

a new system of governance that is tech native… **we need to re-envision governance**."

Another noted that "there are ways in which AI [could enable] real democracies… The **'pure democracy,' where every person has a say, that was impossible without AI. But it might be possible**, with an AI system."

Finally, several participants voiced optimism about the **potential for AI to remove obstacles that have been holding democracy movements** back from true success, some framing this as conditional on pro-democracy actors being willing to start using AI.

One expert described relevant research: "[I had been] trying to understand why so many democracy movements have been getting beat so badly over the last decade. One of the things that has really struck me **[is movements'] incapacity to deal with technology issues**. And really, the technology advantage that governments have over them… And what I noticed is **an extreme reluctance on behalf of a lot of movements, especially in [the] democratic context, to integrate AI in any way into their operations**; they almost uniformly talk about it as… 'avoid, avoid.' And they have their critiques of the tech companies… [but] I realized that **there were instances in which I could imagine [AI] freeing up time for a lot of organizers**."

Another noted that: "[Advocates] want to build community power, and because some might assume that the interests of big tech are against the interests of communities, their first instinct might not be 'let's think about how technology can be used to organize and scale our impact.' I really want to push back on that and say, '[the] water is warm, come on in here, become fluent with these tools, and participate in these policy debates.'… **We need community power building groups in the policy and governance discussions, and we need them to think strategically about AI adoption to improve the efficacy of their work**."

That said, several offered words of caution, noting that AI's democratic impact depends on whether it **inhibits or supports human connection** and voice. One participant cautioned: "When I think of democracies, I think: **what is the glue that brings people together**? And my current working hypothesis is that **if you put AI between the people, that glue becomes weaker… we've got to be very careful not to put the AI between the people**."

AI infrastructure and environmental impacts

Adjacent to core concerns about democracy, many participants mentioned the environmental impacts of AI. **About 20% of interviewees referenced data centers as a key potential mobilizing force**, as tangible impacts of AI are already being felt in communities via infrastructure build out, electricity pricing, and more.

"I see a lot of work clustering around the materialities of AI, [which] has to do with labor, but also... **infrastructural thinking and theory**. We are talking about infrastructures of electricity, of internet, of underwater wires, and data centers... and the upstream impacts of AI on a lot of populations, in terms of labor, but also in terms of **environmental costs**, in terms of **displacement of populations** – because of pollution, because of the noise [from data centers] and the water they take away from people. This field of study also includes the **mining of materials** that are needed to actually power these technologies, which also has to do with **labor and the impact on specific communities** that, not randomly, are generally located in the global South."

Participants commonly emphasized the importance of **providing adequate opportunities for community self-determination** given the known impacts of data center projects on water supply, energy prices, and other aspects of lived environments at the local level. According to MediaJustice, in 2025, $200 billion worth of data center projects were being built across the American South alone.[10]

"We know that AI takes up a lot of data, and also a lot of power. They're building all of these data centers, at this huge environmental cost, that really is of a piece with the ways in which climate change is unfolding. I think it is incumbent on us to find those connections."

Well beyond framing climate as a potential entry point for local democratic engagement, many participants expressed deep concern about AI's environmental impacts: "I think that the whole question of the climate impacts of AI is underserved, understudied, and underreported. I think there's consensus generally that [AI models] are having impacts. But I do not think there is consensus about how we tackle this problem."

---

10   Dulani, Jai, and Samuel Rosado Zaidi. "The People Say No: Resisting Data Centers in the South." MediaJustice, September 9, 2025. https://mediajustice.org/resource/the-people-say-no-report/.

Risks AI Poses to Civic Engagement

While participants commonly raised the possibility of using AI tools to foster new forms of democratic engagement and participation, an even larger number were concerned about **four critical risks: thin engagement, manipulation**, **fake personas,** and **a crisis of connection and belonging**.

Thin Engagement & Gradual Disempowerment: While some participants were excited about the possibilities for more deliberative democracy, as noted above, others discussed the risks of "thin engagement" wherein people might engage in democracy only or primarily via online mechanisms, rather than in person, with real humans.

Risks of "gradual disempowerment" were also mentioned multiple times: "There's a really interesting paper on 'gradual disempowerment,'[11] which I think gets at the cumulative effects of 'thin participation' over time, [where] people having their needs met increasingly by an AI therapist, and the AI social services, and the private companies providing X, Y, and Z – that they are disempowering themselves from the democratic process, and even from [the] social fabric."

Manipulation: For other participants, concerns abound that AI will yield more powerful tools of **persuasion** and **manipulation**, deployed by AI companies and other interests.

"When I talk about manipulation, I'm not talking about deception, such as 'Democrats vote on Wednesdays, Republicans vote on Tuesdays.' I mean manipulation that collects data from you, **generates responses that appeal to you, and covertly nudges you in a particular direction.** Studies show that social media microtargeting is limited in its effectiveness. But the ability for AI to **engage with users in real time, and persuade them, and push them in a particular direction** in a very persuasive way is a real threat. We need more data, research, and solutions to address behavioral manipulation."

Agents, Bots and Fake Personas: While some participants noted excitement about "digital twins" (in which data-driven profiles of people can help model, and improve understanding of, real people), others noted that robust democratic processes ultimately rely on **engagement with real humans**. The anticipated rise of AI agents only heightens these concerns.

"Synthetic personas and synthetic voting with AI is one of the fastest ways to disenfranchise people... I think there's a real risk, and we see so many of these tools being built out in the private sector, where you've got **businesses having legitimate reasons to say, let's build a synthetic persona of you know, 'Gen Z** who are into this type of shoe'... you can see where that goes off the rails very quickly in a democratic context."

Here, participants point to the need not just for deep consideration and sociotechnical research, but also technical work to safeguard processes that would be rendered illegitimate by thin participation: "We need standardization at a federal level about what **proof of humanity** looks like... but I haven't seen anyone actually go apply this in, say, a deliberation."

Connection & Belonging: Despite several participants' expressed enthusiasm about participatory innovation and epistemic expansion, others see the proliferation of AI technologies forcing a deeper reckoning with how people relate to each other: "We've known about the loneliness epidemic as a third order effect of the car. **AI is pouring fuel on the fire**. Community bonds are necessary in order to compromise and govern."

---

11   Kulveit, Jan, Raymond Douglas, Nora Ammann, Deger Turan, David Krueger, and David Duvenaud. "Gradual Disempowerment: Systemic Existential Risks from Incremental AI Development." arXiv, January 29, 2025. https://doi.org/10.48550/arXiv.2501.16946.
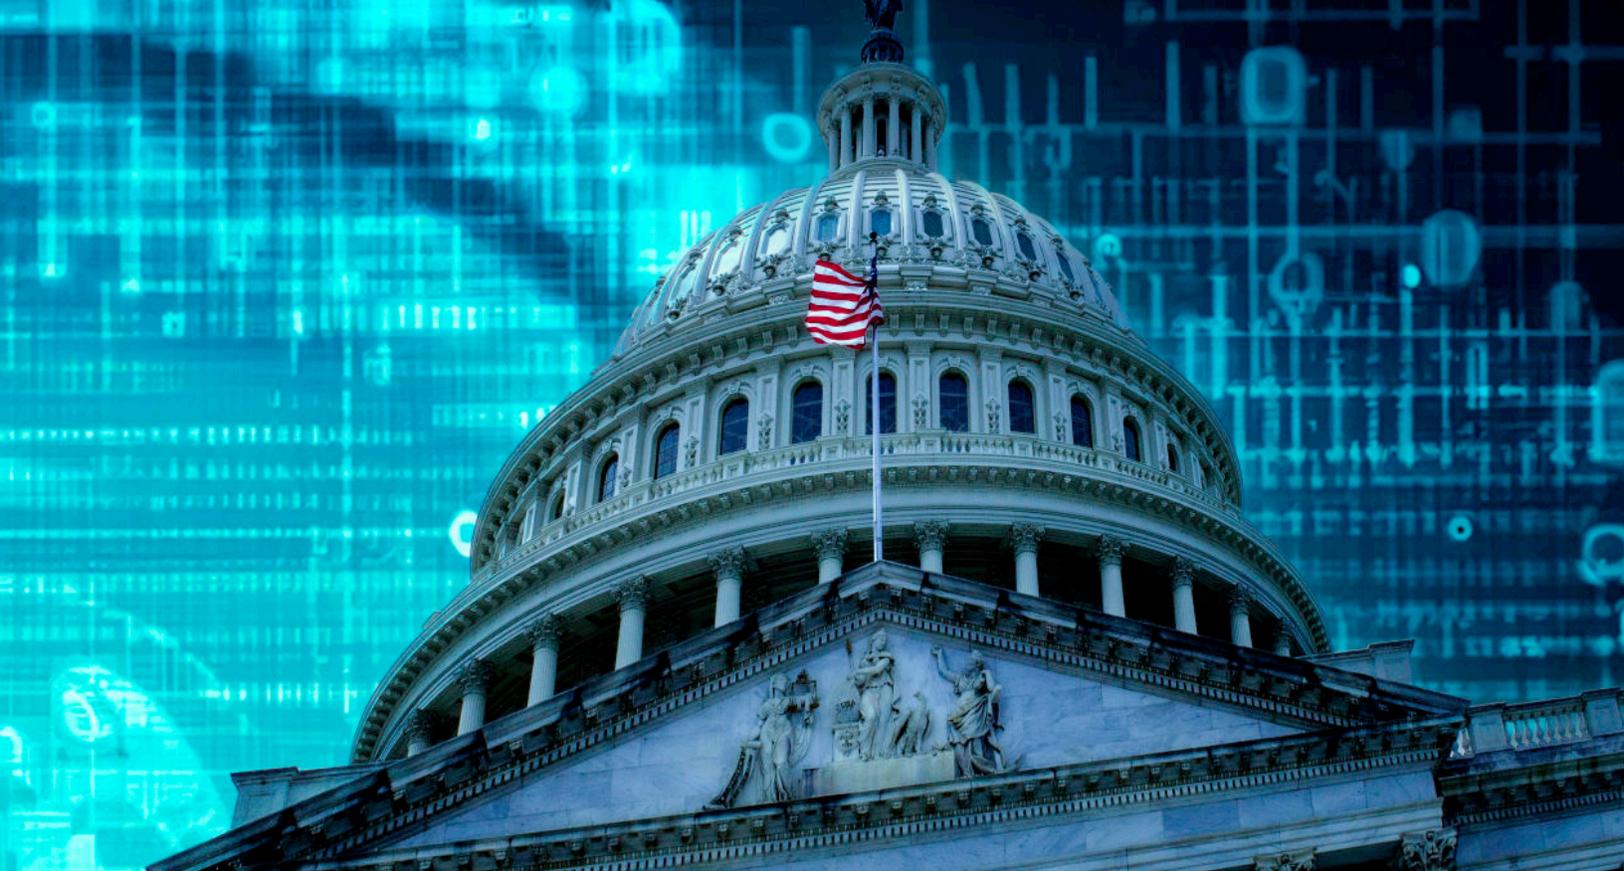
4. USE OF AI IN GOVERNMENT

Unlike in most other areas, multiple respondents were cautiously optimistic about the ability of AI, if deployed responsibly, to **improve the functioning of complex government processes**. Some went so far as to suggest that **AI might be a missing link to helping improve customer service and functionality in government**. Yet others expressed concerns that AI might increase bias (as noted above) while **reducing transparency and accountability,** or be built *on top* of institutional dysfunction rather than reducing their complexity, thus creating even more opaque systems. Participants highlighted the importance of **close – and careful – collaboration** between technologists and government practitioners.

Can AI Improve Government at a Critical Time for Democracy?

Mirroring common excitement about the potential for AI systems to reinvigorate processes for public engagement, some participants are optimistic about **genuine opportunities for renewal and reinvention of democratic institutions**. Several participants are leaning in to support high integrity experimentation with AI tools at the state and local level.

"I think one of the dangers with AI is that you can be out there just building cool things for cool things' sake, [rather than asking] 'What are the problems? What do we think are the most potentially transformative applications?' Let's create safe containers for that kind of experimentation for government, because most [of] government doesn't have the capacity or the risk tolerance to experiment... So I think about [our work] as this living lab."

Other participants argued that the practice of government has **become sufficiently complex** so as to *require* the use of large language models (LLMs), natural language processing, and other AI capabilities to unwind and address.

"Government is going to need to be able to act in bold ways, given the potential pace and scope of [societal] change. We're not really prepared to do that. [At] the top level, **AI is going to impact the bar that government has to meet.** But as a tool, not only will [AI] change the nature of the government workforce, but it is also **a tool for dealing with the incredible complexity that's grown up around our civil service systems**... models like STARRA, the Statutory and Regulatory Review Assistant out of Stanford... there are a handful of these things, some of them open source and some of them not, that are designed to deal specifically with statutory and regulatory and policy language. The only hope you have of taming this complexity... comes from the use of LLMs. You really can't do without them... we have hit this inflection [point]: **these [government] systems don't work right, at the same time that there are starting to be tools that actually... make things possible** that were not possible before."

Others note that trying to hamstring government use of AI may backfire. In her forthcoming book "RECONSTITUTION: A Blueprint for Governing in an Age of Acceleration," Marci Harris describes "**a strange tension in government modernization efforts. In some cases, well-intentioned advocates want data and technology to work but not too well.** Some actually encourage government databases that can't talk to each other, incompatible agency systems, and processing delays **because they believe that these inefficiencies protect people from government overreach**. For a while, they were right. But **that time is over... The collapse is happening everywhere at once**. Modern systems can connect databases instantaneously and enforce rules at scale. Immigration enforcement increasingly uses data from multiple agencies. Facial recognition technology allows 20 federal agencies to access driver's license databases. Predictive algorithms determine benefits eligibility..."[12]

In the recent past, advocates and policymakers introduced some guardrails to government use of technology, largely in response to large-scale failures of predictive analytics, face recognition, and automated decision making tools. Many participants described these efforts as hard won, if insufficient. But others noted how **frustrated civil servants have been by these "protective inefficiencies"** (as Marci Harris terms them) and regulatory interventions. As one expert describes:

"**What I was really seeing in the federal government was not safe use of AI – just no use of AI**, while we jumped through a whole bunch of well-meaning, but badly crafted hoops... When people talk about the harms that could come from misuses of AI in government, to me, they're really not fundamentally different from the harms that have come from the misuse of regular old programming in government. [Take the] 2013 MIDAS[13] (Michigan Integrated Data Automated System) debacle... unemployment insurance [automated fraud accusations] that bankrupted people. The reason software bankrupts people is because we do it badly, and we don't test and learn. **"We don't fundamentally put all of our energy into testing. And that's true of AI, too**... one thing that's fundamentally different with AI is **you have to continually test. If you want safe AI, you need to actually enable far quicker [and ongoing] experimentation and testing. Not this, 'In 5 years, maybe we'll feel comfortable trying this thing out.'"**

---

12   Harris, Marci. "RECONSTITUTION, Chapter 8," 2025 Draft Manuscript.

13   Michigan's Michigan Integrated Data Automated System (MiDAS), launched in 2013 to detect unemployment-insurance fraud, automatically generated large volumes of fraud determinations with little meaningful human review. Subsequent reviews and litigation found the system's fraud determinations were wrong the vast majority of the time (Michigan's Auditor General reviewed a large sample and found an error rate of roughly 93%), leading to tens of thousands of people being falsely accused between 2013–2015. Those flagged often faced aggressive collection actions and severe penalties – wage garnishment, intercepted tax refunds, liens, and damaged credit – before they had a realistic opportunity to challenge the accusation, resulting in significant financial harm for many (including reports of bankruptcies). Michigan later reversed many determinations and issued refunds, and the state agreed to a $20 million class-action settlement to compensate affected individuals for harms tied to the automated process and lack of due process. Cahoo v. SAS Analytics Inc., No. 18-1296 (6th Cir. 2019) https://law.justia.com/cases/federal/appellate-courts/ca6/18-1296/18-1296-2019-01-03.html

Several experts emphasized both the risks of **failing to deploy AI, which could cause governments to fall further behind,** and the importance of evaluating AI not against a standard of zero bias or error, but **in comparison to the inevitable bias, error, and inefficiencies present in existing** *human* **systems.**

How to Deploy AI in Government, Responsibly

Participants noted that responsible deployment of AI in government requires addressing process complexity, building talent and capacity within institutions, partnerships between technologists and government practitioners, and the use of government procurement power.

One noted the **need to ensure that AI systems are applied toward** *structural transformation of how government operates***, rather than just addressing surface level efficiencies.** And to ensure that government does so transparently.

"I think that a lot of harm is coming to people because of these archaeological layers, essentially, of law and policy that have made our systems ungovernable… One thing… [I mostly] hear from industry and the uninitiated is: 'Great. You have something that's complex. AI can navigate that complexity for you.' That's actually a terrible idea. **If you don't use the AI to reduce the complexity, you end up with systems where the AIs are just talking to each other, and nobody actually knows what's going on in the back end…** you [are left with] a fight between black boxes. It's actually terrifying [that we might get] to the point where, really, nobody knows what the truth is."

> "I think that a lot of harm is coming to people because of these archaeological layers, essentially, of law and policy that have made our systems ungovernable… One thing… [I mostly] hear from industry and the uninitiated is: 'Great. You have something that's complex. AI can navigate that complexity for you.' That's actually a terrible idea. **If you don't use the AI to reduce the complexity, you end up with systems where the AIs are just talking to each other, and nobody actually knows what's going on in the back end…** you [are left with] a fight between black boxes. It's actually terrifying [that we might get] to the point where, really, nobody knows what the truth is."

To ensure AI deployment actually benefits the public, **greater partnership is needed between tech experts, who understand the tech, and government leaders, who understand what government problems need to be solved**:

"[We need to let government] leaders lead on **what problems need to be solved**. They [don't always] have the expertise to know the right applications, necessarily, of emerging tech… How do we help them get clarity on the problem? How do we educate and push builders to go after those problems?… If left in a vacuum, the builders will just go to the easiest thing, or the thing they understand the most, the thing that seems the most interesting… We need to listen better to what the choke points are**, and then we need to actively aim the [smartest] builders at those problems**."

Given ongoing government spending on technology, **governmental procurement practices have the potential to influence how technology is developed, and what guardrails and safeguards are in place**. Engaging public sector leaders as buyers of technology could be an important leverage point for ensuring AI Is deployed in ways that benefit government and society.

"The thing I find that excites [public sector] leaders is this idea that they could help define what comes next: can they define the parameters and the direction that vendors are going in?... But it's a world of tension right now, because how do you embrace and experiment with tech while also knowing there are a lot of downsides... My mind tends to go towards: **how do we help people [in government] be better buyers?** Right now, you have 50 states all negotiating their own agreements. Where is the leverage in that?... We're not going to have a public model [anytime soon], so it's all still being built on the backs of a handful of tech companies."

Finally, some experts noted that to *govern* AI thoughtfully, government must be *using* AI. These proponents questioned the prudence of waiting to adopt or deploy new technologies until regulatory frameworks are in place (i.e., to ensure technologies were safe and reliable for deployment). One expert shared:

"I think a couple of years ago there would have been much more of a pushback on the use of these technologies. Now I feel as if people are saying, 'You can't manage these technologies if you don't understand them... if you're not using them.' And when you start using them, you realize that there's some good things about them, and there's some not so good things, and then there's some really not-great-at-all [things]... [So use of the tools is] something I am hearing more appreciation for these days, particularly when we think about governments. **How can they set regulation around [AI] if they don't really understand and aren't using the technologies in some way?**"
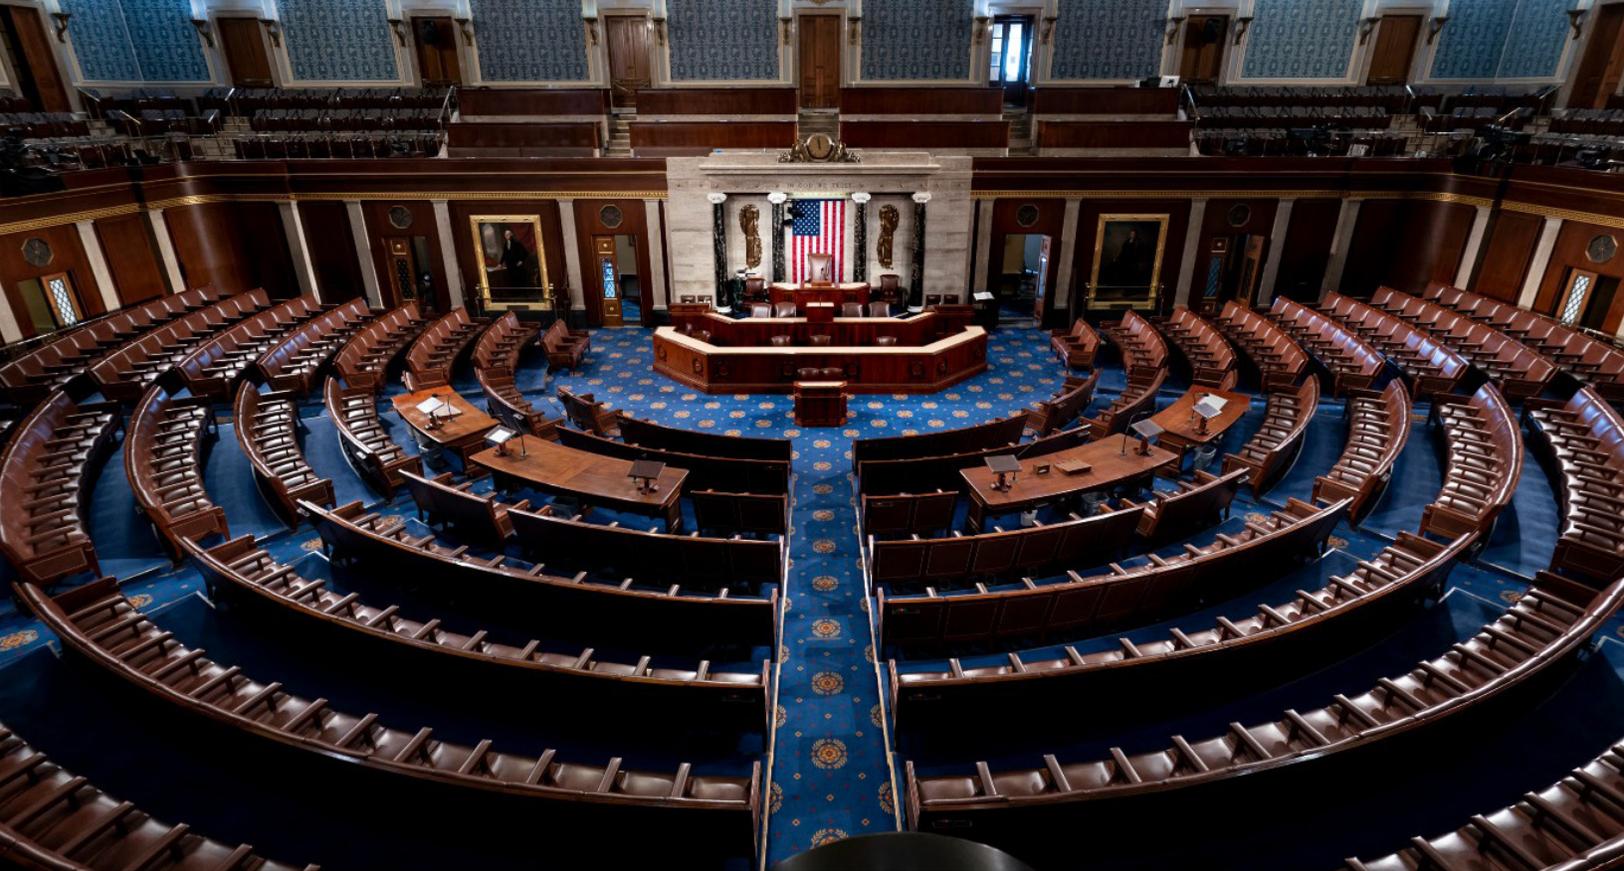
5. GOVERNING AI

Having explored how government is deploying AI internally, we now turn to how AI itself is *governed* across public and private sectors.

Governance of AI can happen at multiple levels, including at every stage of its development and deployment, and by multiple stakeholders. **Legislative bodies** play a central role, but others do as well. This includes **civil servants** who make decisions about technology procurement, service delivery, and policy enforcement. It also includes **businesses across all sectors of society – from finance to healthcare to education – deploying AI in pursuit of their missions.** It includes **duty-bearing professionals** and the **professional associations** (such as bar associations in the legal field) that uphold the legitimacy of their expertise.[14] And of course **AI companies themselves**, who make AI governance decisions regarding overall corporate governance, and at every stage of product development.

Mapping Policy Gaps and Priorities

While we sought to include a broad range of perspectives on AI and democracy – including individuals with prior professional experience in private sector technology companies – our interviewees did not include VCs or "AI accelerationists" and only included a few participants aligned with the current federal administration. Nevertheless, it was notable that ***nobody* in our sample argued for *no* regulation of AI**. Instead, we heard important nuances about timing, the necessary coalitions required to build power, gaps in the evidence base, practical concerns about enforceability, and questions regarding the intended outcomes of interventions.

---

14 Gansky, Ben, and Sean Martin McDonald. "The Governance Gap: Preparing Professions for Digital Transformation." CIGI, November 27, 2023. https://www.cigionline.org/publications/the-governance-gap-preparing-professions-for-digital-transformation/.

One expert shared: "I think **there is consensus, even amongst [AI] companies, that they are moving too fast,** that the systems are not high quality yet. That there's real risk [whether AI will] sustainably drive positive and beneficial outcomes... Of course... regulation takes time... and in some ways it should take time. But... certainly in the US [there is consensus that] things like a **federal privacy law,** and **better online harms protection**, at least at the federal level, are really important."

Across participants, there was general agreement around the need for baseline protections – particularly **privacy**, **transparency**, **accountability**, **use case restrictions**, and **rights to appeal** in cases of "consequential decisions" (such as employment, housing, credit / access to finance, incarceration / judicial action, and benefit eligibility). Some argued that better enforcement of laws already on the books may be sufficient, while others believe new policies are necessary, particularly to address gaps in privacy protections, corporate accountability, and the governance of rapidly advancing or "super-intelligent" systems.

Despite agreeing on its necessity, participants have **wide-ranging views regarding where and how governance of AI should occur.** Regulation could, for example, aim at:

» **Frontier AI companies** themselves (e.g., Google, OpenAI, Anthropic, XAI)
» **Individual models** (such as OpenAI's GPT 5.2 or Anthropic's Sonnet 4.5)
» **Deployers** (e.g., loan officers, hospitals) in specific impact areas (housing; healthcare; employment; etc.) where there are sometimes existing regulations that can be adapted, and sometimes needs for new rules

Notably, participants across the political spectrum urged caution regarding direct liability at the foundational model layer, while largely seeming to agree that liability at the frontier company and deployer levels are both important to pursue (albeit in different contexts and to different degrees).[15] Others nodded to the need for clearer regulation of chips and underlying infrastructure.

While participants expressed a range of policy priorities, several emerged as cross-cutting – what one funder described as a set of "table stakes for democracy":

» **Privacy** protections that assume and assure the existence of "a sphere that is outside the view of the government."
» **Transparency and explainability** regarding the use of AI, how AI systems work, the data used to train AI models, and where and how it is being deployed for consequential decisions regarding, for example, benefits, incarceration, immigration, and more.
» **Use case restrictions,** including for example profiling, police use of facial recognition technologies, social scoring systems, automated eligibility systems, and other potentially harmful AI applications.
» **Access to redress** mechanisms that allow people to appeal consequential decisions made about them by AI, an essential component of due process.

Accountability and liability were also mentioned by a majority of interviewees, though, as noted above, uncertainty remains about whether it should accrue at the level of the company, deployer, or foundation model. While many participants argued that companies must bear liability for harms, biases, or resource extraction associated with AI systems (e.g., from data centers to copyright infringements), a few participants

---

15 Participants who urged caution regarding model-layer liability often pointed to challenges of: 1) attribution (i.e., whether the problem lay in model design or in how the deployer chose to integrate it into their systems); 2) the general-purpose nature of foundational systems (i.e., general purpose tools might find thousands of downstream applications, many of which the original developer did not design for and cannot control), and 3) concerns that poorly designed liability regimes could entrench large incumbents rather than reduce harm.

believed strongly that existing tort and civil law are sufficient. (Note: While not mentioned by our interviewees, others argue for "safe harbor" – to shield companies from liability for several years in order to avoid stifling competition until the risks and capabilities of AI are better understood).

Several participants emphasized that **governance must happen in partnership with impacted communities**. This might include professionals working in industries impacted by AI, beneficiaries of public systems being transformed by AI, and broad constituencies (e.g., people of color, women) that are often opaque to (or poorly represented among) builders of AI technology.

Beyond pursuit of new legislation, many in civil society are continuing to focus on **implementation of existing policies, strategic litigation, and creative use of existing enforcement authority.** State attorneys' general, in particular, have been applying existing consumer protection, civil rights, and corporate governance statutes to AI-related harms – even in the absence of AI-specific statutes – at times establishing important precedents in influential jurisdictions such as California. Existing legal frameworks – including disparate impact (re: discrimination), antitrust (re: corporate concentration), tax policy (re: wealth distribution), copyright (re: intellectual property), consumer product liability (re: harms from consumer technology) – already protect Americans from many harms upstream and downstream of AI systems. Of course, each of these areas represents a contested policy debate.

However, one participant, while agreeing that the work to adapt existing institutions and regulations to AI is worthwhile, also warned about cases like speeding or jaywalking, where **relevant laws were written *assuming* imperfect enforcement, not 'a surveillance state.'** "Because the laws were written [with] a set of expectations about what technology was capable of [that no longer holds], merely ratcheting up the compliance and enforceability of existing laws that were written and calibrated for imperfect enforcement could become very totalitarian."

Many participants again referenced the political economy of AI, arguing that **philanthropy often avoids confronting the deeper structural forces at play**.

"I think **philanthropy has incentivized the field to take piecemeal solutions**, and to not piss off the tech sector, and to align ourselves with... neoliberal policies that have enabled the tech sector to flourish... I [think] it is the [result of] a ... tech based industrial policy [that]... gives [tech CEOs] free rein, even though we're whittling away at liberal democracy, we're going to continue to enable these people who believe they are gods to continue to build their technology unfettered... I don't really see much of the movement, or philanthropy, understanding that **the reason that we have this problem is because of an erosion of liberal democracy; that the tech industry has never wanted any kind of regulatory frameworks placed on them**."

Such arguments point towards heightened antitrust enforcement.

Others argued that, **rather than a simple focus on which policy ideas might be preferred in *principle*, the field should pragmatically wait for policy windows to emerge and then do its best to work within them.**

"Most policy decisions are made when windows of possibility open up. [Those policies] have to be fairly pragmatic... [it] suggests to me that the right way to think about these policy

> "Because the laws were written [with] a set of expectations about what technology was capable of [that no longer holds], merely ratcheting up the compliance and enforceability of existing laws that were written and calibrated for imperfect enforcement could become very totalitarian."

questions is often [to focus on] **what society we want to achieve, and what are pragmatic steps [to get us there]?** Most of the time, an 'AI Bill of Rights' type of framework doesn't add a lot of value. There are always going to be exceptions... [but] a better question is... [if] we want a vibrant and successful democracy, and we can see the direction AI systems are going [then], let's back up: **What are the core problems in democracy that can either be made worse or better by AI?** [From there we can] begin to lay a framework, and very rarely will that map onto some grand discussion of, [for example,] the need for more privacy. Sometimes it might, but... I'd want to know **to what end**."

Where and How Can Policy Happen?

Participants emphasized not just identifying areas of policy agreement, but building the coalitions and platforms that ensure that necessary policies get passed. Multiple participants believed that the field already has much of the necessary analysis, and the next step is coalition-building: "There are a lot of good [places where] we could get to consensus. A lot of good research and work has been done, and I think it is just about **building the coalition** to make that happen."

Participants discussed two interconnected dynamics: the need to creatively examine *where policymaking and governance are possible*, and the need to *build new coalitions,* including potential roles for the private sector, and non-traditional tech players.

Where Policymaking and Governance Are Possible

Amid a complicated federal policy ecosystem, some participants are looking to the **state and local level to foster innovation** in terms of both policy enforcement and cross-sector partnerships. [*Note: these interviews were conducted before the December 2025 AI Executive Order on AI, seeking to limit AI policymaking at the state level*].

"There are a lot of funders who are thinking... it's not just about the silver bullet of **Washington DC policy, all eggs shouldn't be in that** single basket."

In a similar vein, one funder highlighted the value of leveraging **state and local governments as "laboratories of experimentation"** from which the field can derive insights to better understand harm mitigation, rather than pausing adoption entirely due to theorized risks. Others recommend **watching state attorneys' general (AG) offices for innovation in enforcement of AI regulation**. In one case, AGs from California and Delaware used OpenAI's nonprofit status to apply a particular standard of liability for product harms.[16] Still others are closely observing the implementation and enforcement of emerging state AI regulations, including in California and Colorado where contested privacy, transparency, and accountability requirements are still being defined, with heavy input from industry lobbyists.[17]

Building New Coalitions

Many participants emphasized the need to build stronger relationships between **researchers, local communities, public leaders, and advocacy organizations** to capture moments of opportunity that can result in stronger governance ecosystems. For example, community organizers in Minnesota moved from resisting a proposed predictive model in the county education system, to being invited by local government to co-create a Transforming Systems Together Board to embed community voice in public decision-making

---

16   McDonald, Sean Martin. "Letter from Attorneys General to OpenAI," LinkedIn, September 6, 2025. https://www.linkedin.com/posts/sean-martin-mcdonald-04b1b221_letter-from-attorneys-general-to-openai-activity-7370129858846433280-ld5W/.

17   Oduro, Serena. "In Delaying Its AI Law, Colorado Shows Tech Lobby's Power In State Politics." TechPolicy.Press, August 29, 2025. https://techpolicy.press/in-delaying-its-ai-law-colorado-shows-tech-lobbys-power-in-state-politics.

regarding AI.[18] At the cross-regional level, a 2025 campaign from AI Now secured 70 coalitional signatories, including many local advocacy groups.[19]

Several expressed a renewed enthusiasm and humility for tackling specific problems in partnership with communities they have formed relationships with. Many emphasized the need to broaden the coalition still further: "[We need to involve] all of the organizations with political influence and reach. Where's the AARP on AI? How would you get to them and make them believe that this is an issue they need to pay attention to?"

Many discussed the need – at a time when federal policy is unlikely and state policy is fragmented – to **engage more directly with (some) elements of the tech community itself:**

"I've had conversations with people who work at Meta… [who note that] before… movement groups came with a sledgehammer, and not enough of a willingness to [coordinate with] the people who could be allies within those institutions.' [These are not] the Mark Zuckerbergs of the world… **the people working within those institutions who are trying to change [things]**. [The Meta source indicated that] we are in a new moment, however, with the authoritarian breakthrough coinciding with **people who work inside of tech companies [who fear] unfettered AI [and] are defecting.** There's information to gather… [but it] requires a lot of careful consideration and talking to people who have organized tech workers in the previous 15 years."

Others **argue for a still greater private sector role in the coalition**, especially from groups that can legitimately speak for consumers. Such experts argue that neither industry nor civil society can assume they represent all public preferences.

"**I really do think there's a private sector role, and I get that people don't trust the incentives**… and there's good reason for a lot of that. There's an important role for traditional consumer groups to educate and help figure out: **how do we think about what is best for consumers?** Because I worry that some of the civil society groups, especially on the civil liberties and privacy side… think and believe that they are speaking on behalf of consumers, and I just don't think that's true… By way of example, a lot of consumers, whether or not they have real choice… are going to use services that are not privacy protective because they're easy and convenient. That is a reasonable choice for a consumer to make, and as much as you don't like it, that's the way a lot of people are going to decide it. And I get that you want to be the defender, and change things for the better [but that] does not [necessarily] mean that you actually are doing what [consumers] want. You are not their agent."

Others take it one step further, arguing that **not all private sector actors are the same: the frontier models face very different incentives than other private sector users / deployers of AI** – e.g., the small and medium sized business who are highly incentivized to ensure the trustworthiness of any AI technologies they deploy. (Hence, the growing AI assurance technology marketplace, which some project will be valued in excess of $275B by 2030[20]).

At the same time, most participants noted **limits to the breadth of possible coalitions**: "I honestly don't know if it's possible in this new world to have a multi stakeholder [initiative]. I think we might be at war, and I think the [AI] companies have picked their side."

---

18 Pfefferkorn, Marika, Aasim Shabazz, and Michelle Shevin. "Co-Powering an Emergent Horizon." Stanford Social Innovation Review (SSIR), July 14, 2022. https://ssir.org/articles/entry/co_powering_an_emergent_horizon.

19 AI Now. "People's AI Action Plan." Accessed August 15, 2025. https://peoplesaiaction.com.

20 https://aign.global/ai-governance-insights/patrick-upmann/ai-assurance-technology-in-the-ai-governance-context-development-until-2030

Constraints on Governance

Political Economy of AI

Many participants emphasized that governance debates cannot be separated from the ownership structure, market incentives, and lobbying power shaping the AI ecosystem. Addressing the political economy of AI begins with a basic question: **should AI systems be required to meet the legal and professional standards of the domains in which they are deployed**?

If an AI system is used in healthcare, lending, employment, or public benefits, it should be subject to the same substantive requirements that govern those fields – including civil rights law, due process protections, and professional standards of care. Beyond that, structural issues of market concentration, labor protections, and corporate governance also shape how AI systems are built, deployed, and regulated. Without confronting both domain-level accountability and broader concentration of power, regulatory outcomes will be shaped primarily by the interests of dominant firms rather than democratic oversight.

"Both democracy and AI are systems designed to manage and realize interests – and both systems need significantly more, more accessible, and more reactive scaffolding and mechanisms in order for them to be fit for purpose… At present, it's not even clear whether there's consensus that AI production and deployment should have to acknowledge the specific requirements of its use case."

The Black Box: AI Itself as a Governance Constraint

Several participants pointed to a deeper structural risk: good governance depends on feedback. Laws evolve through setting precedents, enforcement, learning from that enforcement, and reforming legislation to address any shortfalls. **AI systems, particularly when opaque or decentralized, can disrupt that feedback loop. If harms are difficult to trace, if decision-making is illegible, or if systems change faster than institutions can observe and respond, governance loses its evolutionary capacity.** This creates not just regulatory lag, but a risk that democratic systems cannot update themselves in response to emerging conditions.

> AI systems, particularly when opaque or decentralized, can disrupt that feedback loop. If harms are difficult to trace, if decision-making is illegible, or if systems change faster than institutions can observe and respond, governance loses its evolutionary capacity.

Government Regulation vs. Consumer Choice

While all interviewees agreed on the need for some form of regulation, many also cited the **limitations of regulatory approaches and the need for other forms of public resistance and adaptation.**

Some participants noted **that fundamental attitudes toward privacy and surveillance have already shifted. People have become accustomed to their private information (photos, videos, addresses, etc.) being public**, and are skeptical that anything other than personal choice can change it.

"**It's just not clear to me if you can actually pass a law [protecting privacy**, but]… We've seen people shift off of large platforms onto Signal and group chats and things like that, **I think people will vote with their feet**… or will [pro-privacy approaches] develop in an arms race fashion, like the Hong Kongers with their umbrellas [popularly deployed to prevent surveillance]?"

"I don't know if we can build a new tolerance for the wisdom of not having the government know everything about you or… **everybody else knowing everything about you**… **it's so deeply internalized in the next generation**… **There would need to be a broader normative reorientation for any [privacy] litigation to stick…** What would a new conversation about privacy – that actually registers with and makes sense to young people – actually look like?"

Another expert advocated approaching regulation from an **economic justice rather than a technology lens.** When asked about possible areas of consensus, this expert professed: "I think there's consensus that **we need to find ways to govern that aren't about direct policy-making about 'AI'**… we're in a closing space… I think there is consensus that we really need to… build AI governance into a broader economic justice… platform."

We also **heard *a lot* about practical difficulties involved with implementation and enforcement** of regulatory interventions: "I've been working with regulators and other practitioners to help develop frameworks for what **meaningful enforcement** might look like in this space… **there's a skill and expertise gap between regulators and the producers of these systems."**

6. NATIONAL SECURITY AND SYSTEMIC RISK

AI threatens to shift the balance of power not just within democracies, but between democratic and authoritarian states.

While few raised national security or geopolitical concerns unprompted, when these issues did arise interviewees expressed serious concerns on three fronts: **surveillance; military deployments (e.g., autonomous weapons), especially by nonstate actors; and intensifying geopolitical rivalry.**

Surveillance

Many participants noted concerns about increased **surveillance capacity, including current applications for immigration enforcement**.

"AI as designed today... its model is a task-replacing tool, or **a surveillance tool that enables [the] individuals... [controlling the data]. So the authoritarian play is advantaged by these kinds of tools**."

Participants emphasized that populations already under surveillance, or whose movement through space is tightly controlled, face disproportionate risks from increasing AI deployments.

"We are interested in government uses of AI for national security... pretty broadly defined... [everything] from **surveillance by the government,** its linkage to **surveillance by private actors**, so basically the whole data broker issue; to the use of **algorithmic tools for making decisions** and recommendations relating to things like whether someone is regarded as a national security threat, whether they're placed on a watch list, whether they are flagged for immigration enforcement... that whole spectrum of consequences."

AI and Asymmetric Power

Some participants expressed particular concerns about unfettered access to powerful capabilities by **nonstate actors, non-democratic regimes, and coalitions able to accrue asymmetric power**.

Emphasizing a long list of safety issues AI labs are effectively distracted by, one participant pointed to the possibility that progress could be made on important problems while bigger ones go unaddressed: "[Highly capable AI] creates really, really bad dynamics [in terms of] societal power structures; the likely... worlds we end up in are not going to be good... because of the way that power and AI interact. And in particular, **the kinds of power accumulation [resulting from] very decentralized AI, where essentially [you have] broad access to weapons of mass destruction.**"

Participants emphasized agentic capabilities adding to already existing asymmetries that empower private actors.

"I think [a democracy where AI agents are proliferating] looks like a world where Waco-style, Ruby Ridge-style events are happening all the time... I think it's very plausible that within two years, we're going to have **off-the-shelf, autonomous drone technology that you could put into a consumer drone,** anything with a camera, and **orchestrate drone swarms**... I think we need to start fortifying the nation-state relative to these private actors."

To mitigate risks, participants emphasized the need for standardized **risk measurement frameworks** and improved **global coordination on regulation**.

International Coordination, Competition, and Conflict

Given the potential for extreme change from several vectors, participants emphasized a **need for greater global coordination** and cooperation: "We see a lot of opportunity ... [for global AI governance to be] guided by democratic and rights-based values as opposed to arms race scenarios."

However, interviewees did not necessarily agree on which specific policies would best address national security risks posed by advanced AI capabilities. For example, one participant described export controls as "going through a period of unraveling consensus."

Participants emphasized an increasingly complicated geopolitical landscape, inextricably intertwined with technology production and infrastructure: "**What does it mean to have [AI sovereignty]**: sovereign databases, sovereign models, infrastructure that's provided by [domestic] providers as opposed to having to rely on US providers?... [this has even greater] implications in the case of many non-democratic states."

Across these domains, participants emphasized that **AI increases the scale and accessibility of capabilities that were once limited to states**. This shift may strain existing legal frameworks designed for slower-moving or more centralized threats.

7. PATHWAYS FOR PHILANTHROPY

Participants consistently framed **philanthropy's task as strengthening the ecosystem around AI and democracy by closing coordination gaps**, **building capacity**, and **investing in the infrastructure** needed for durable, adaptive response.

Across nearly 40 conversations, there was perhaps **nothing we heard more consistently** than some version of '**we need more opportunities to make sense of what is going on** and what needs to happen in relationship with each other.'

When asked about the roles that philanthropic funders could play, several areas of agreement emerged: **strengthening field coordination**, **investing in measurement and foresight, supporting institutional adaptation**, **bridging divides between near-term and long-term risk funders**, and **aligning on a central set of common priorities**.

Defining (and Coordinating) "The Field"

In what one expert referred to as "**a thousand voices shouting in different directions**," many participants respond to questions about "the field" with tongue-in-cheek overwhelm or confusion. "**What field?**" and **"What do you mean by 'the field'?"** were two characteristic responses.

"We [funders] don't all agree. We don't agree on what the problem is; [**we don't] even have a definition for the problem...** [Technologists and tech funders are not connected to pro-democracy actors] It's tech over here, siloed... we [can't] confront the problems that tech is wreaking on global democracies or economies if we continue to silo it. **[Tech] has to be something that's far more integrated into other areas**."

Another funder shared, "There's so much noise in the space, and so much happening, and so many groups and so much attention, and a lot of it is not crisply defined. **What change one is looking for, and how we'll measure success, is not defined.** [It's unclear] where to even start... it's like whack-a-mole, and if you come up for air in six months, it's a different landscape."

One expert reported, "The democracy piece of it is still getting figured out... from my perspective, democracy in the funding sphere for too long has been considered really narrowly, and it's really been about voting rights and voting technology, and secure elections, and all of that is really important. But I think expanding what is considered democracy, and the different institutions that make democracy possible, and the ways that they intersect with technology... **isn't a coherent ecosystem yet.**"

Yet another funder reported, "It's very scattershot at the moment. On the academic side, **we haven't seen coalescence around definitions and methods that allow academics to have one conversation about a topic as opposed to a bunch of different topics**. One of the great outcomes of cross-funder field-building work on 'public interest technology' is that now that field has enough of a shared vocabulary that computational social scientists can speak to policy experts, and to economists... We're not there yet with AI. It's going to be thorny. We'll devote resources to convening organizations for that vocabulary to emerge."

Not all funders shared this opinion. "There is definitely a field now," affirmed one technology funder who has been working on "ecosystem development" for many years through a public interest frame. "**We all use different terms to define this space right now... [but] we're siblings.**" Summarizing the results of a recently commissioned independent evaluation, this funder shared that people who were skeptical "even three years ago... are now true believers."

In terms of **bringing in more funders**, one philanthropic leader suggested that the technical nature of the subject can intimidate potential funders: "Part of the problem we face in philanthropy is that **a lot of people don't feel like they're tech experts** – and they're not, it's true, but [they think] therefore they don't have a perspective... to bring into the conversation. I spend a ton of time telling people, '**you don't actually need to understand how neural networks work to be able to be a smart grantmaker** around these issues."

A Missing Middle: Bridging Views About AI Risks

"The framework that I'm sick of, but that... I think it is sort of on the way out, is [that] either AI is going to solve everything, or AI is going to kill us. That [is] not only a false dichotomy, but actually, those points are actually on the same end [of the spectrum], because they both assume that AI is really effective."

Despite the clear urgency of the moment, and widely expressed need for clearer coherence and sense-making, **the field remains clearly divided between funders with very different understandings of risk and potential harms**.

For several years, "catastrophic" or "existential" risks (the latter is sometimes shortened to 'X risks') from AI systems have been positioned as distinct from a set of already-occurring harms. Many participants referenced an ensuing division between funders and advocates, often framed as **a "safety vs. ethics" debate**. However, an emerging surface area (sometimes mentioned by participants in terms of an emerging theory of "gradual disempowerment"[21]) offers a potential bridge between the two camps, highlighting **shared priorities regarding corporate accountability, improved measurement and testing of AI systems,** and **pragmatic approaches to reducing risks and redressing harms**.

---

21 Kulveit, Jan, Raymond Douglas, Nora Ammann, Deger Turan, David Krueger, and David Duvenaud. "Gradual Disempowerment: Systemic Existential Risks from Incremental AI Development." arXiv, January 29, 2025. https://doi.org/10.48550/arXiv.2501.16946.

Perhaps because our interviewee sample skewed heavily towards nearer-term risk funders, only a small minority (<10%) of our interviewees professed to prioritize concerns about "superintelligent" and "runaway" AI systems. Instead, most interviewees focused on concerns about how much agency, autonomy, and decision-making that governments, corporations, and individuals are ceding to AI systems that are decidedly *not yet* superintelligent. Several participants also emphasized that, for communities already experiencing harms and marginalization as a result of AI deployment, near-term concerns often *are* existential.

"For people who experience these tools as biased or harmful to their communities, the distinction between 'existential risk' and more immediate harms can feel abstract. If AI is eliminating jobs, reinforcing bias, and marginalizing the people you care about, **that reality may not seem meaningfully different from a scenario in which AI takes over entirely**. From that perspective, the fear of artificial general intelligence erasing all human agency is **not fundamentally distinct** from a nearer-term concern: that a narrow, relatively homogeneous group will use AI to concentrate power and diminish the agency of most other Americans."

Ultimately, participants noted that **people who believe AI will rise up to take over may share many priorities in common with those concerned about the risks of gradually putting AI in charge.**

A smattering of participants sought to bridge the 'catastrophe is coming' and 'catastrophe is already here' framing, noting instead **a persistently "missing middle"** of discourse and inquiry: "It's funny, **you solve one thing, and then you still end up with the conclusion that AI changes everything.** It's like: AI might take over the world. Even if we solve that, AI might enable authoritarians. Even if we solve that, AI might enable terrorists. Even if we solve that, there will be, like, 85% plus unemployment, like nothing we've ever seen, right? And what is our tax base then?

> "It's funny, **you solve one thing, and then you still end up with the conclusion that AI changes everything**

What are the underpinnings of democratic participation then? What about meaning? I feel like there's such a long to-do list for society, and… a challenge you might find is that [there are] people like me, who really take the crazy disruption in the near term seriously… And then there's other people who might be inclined to think about democracy and labor markets, etcetera, but mostly they still don't believe that AI's going to be that big a deal. So there's a bit of a missing middle, where **all the people who think AI is going to be a huge deal are working on existential risks**. **And then other people who are working on ethics and fairness in the near term mostly think AI isn't very good**."

Another wished for **"all of philanthropy to be more AGI-pilled,"** pointing to the "residual skepticism of some large foundations" that upholds a rhythm in which "advocacy works in these five-year chunks that… all [need to] be recalibrated."

Yet many participants underscored that **our ability to harness the benefits and mitigate the harms of these technologies will hinge on our ability to find middle ground between these two extremes**: "The next two years is going to be a tightrope… **how do we create safe containers for experimentation without swinging too far in either direction**, where you either get stuck being the funder or the organization that is purely accelerationist, or the one that's, like, 'Put the brakes on. This is dangerous.'… **For the future of our democracy and government service delivery, it doesn't help us to have a bunch of people either on one side or the other."**

Convening and Collaboration in (Careful) Pursuit of Coherence

A majority of participants noted the need for **more in-person convening.** Many described a need for in-person spaces to share expertise. However, some noted concerns about industry capture, accessibility, and risks of groupthink / pressure for coherence.

"This community often mistakes convening for influence and expertise for agency, but it needs more of all of them."

One participant referenced the constraints placed on convening spaces by growing industry presence: "There were spaces that were really good. They did a good job at bringing people together and really putting out issues that were, in a way, even radical… We need radical ideas in this space, because otherwise we [are] just going to fall victim [to] the hype… there were spaces like that… But sadly, a lot of those spaces have been **co-opted by industry**, [and] serve purposes of ethics washing… [at one recent conference] it was very difficult to [even] have a conversation about difficult topics."

Others described a desire for places in which they could **engage in productive discussions, rather than panels,** or for places **bringing together people who do not typically work together** but have expertise in relevant areas: "I really am keen on convening places for NGOs, academics, policy makers, [and] democracy advocates to be meeting. That's not happening enough, [but] **cross-pollination is important**. [In] the election space or the democracy space in general, there's a long history of relationships that have been built for working together. I'm not sure that really exists in this emerging AI space."

A few participants emphasized that **convenings should prioritize accessibility considerations**. One executive director cautioned that "Funders love convenings and in-person events, but if I want to build scalable infrastructure, I can't be jetting around to events all the time," while another shared the reminder, "Online spaces… are still of great help for a lot of people with accessibility issues, or people who cannot just leave their countries for people with visa issues, etcetera."

At the same time, several participants cautioned funders against forcing the field towards alignment: "**Funders should not expect the field to cohere**. We should **create conditions to allow grantee organizations to exercise self-determination**, and make flexible resources available for them."

Another noted that "**All of us should be in pursuit of a pluralist, diverse, multiracial democracy… I want funders to stop trying to cohere**… I think we have over-prioritized collaboration and co-funding in ways that have actually undermined and reduced the amount of dollars and resources going into good hands. It has encouraged funders to decide together in ways that result in groupthink… I think **this is a time where we need every possible diverse point of view, every possible perspective, angle, choice**… [No one has] solutions in this moment. I think that… [philanthropy's] job is to [do grantmaking] with integrity and honesty and values. It is not… our responsibility to have answers or to drive strategies that are incubated in beautiful conference rooms in San Francisco or New York. I believe our responsibility in this moment is to steward resources in the service of creating the conditions that allow for the solutions to emerge from the ground."

Funding Technical Solutions

Interviewees were conflicted about whether philanthropy should play a role in building tech solutions. Multiple participants spoke about the **need for philanthropy to invest in building essential technology tools, thus demonstrating what public interest approaches look like in practice.** Conversely, **many others cautioned against such funding,** reflecting on many past "tech for good" investments that proved unsustainable.

These skeptics instead advocated that philanthropy should be **preparing struggling *institutions* (from hospitals to schools to newsrooms and governments) to be able to effectively absorb emerging tools and maintain existing technologies**.

"One thing I've loved seeing is that a lot of funders are now shifting towards funding tooling. Even a few years ago, funders did not really understand [it and would say things like] 'oh, we fund convenings and papers... **how do we get to the world where we're making solutions?**"

Another participant spoke about the catalytic role philanthropy could play in supporting globally-informed technology development: "All over the world... everybody is dealing with the same stuff... We should **get a bunch of people together from around the world [together] with people who want to do tool experimentation, and have a hackathon that's informed by the latest activist needs from around the world, as opposed to just in one country**... Activists from other places ... know what's coming; they lived through this five years ago. They can tell you where it's going."

However, others caution against philanthropy funding technical solutions that may be unsustainable or lack market demand. Some argue instead for **devoting resources to ensure that the people and institutions that will be deploying these technologies are stronger**, more resilient, and capable of absorbing them.

"There was this moment 7 to 10 years ago where so many organizations were creating apps to solve problems. And I think that **funders should be really careful about thinking that... a bunch of investments around tool development or small model development, etcetera, [will solve the problem]**. And then that stuff is not sustainable, or it becomes [a] resource drain... There is a whole history of philanthropic investment around technology, and I don't know that it's actually been written up very well, but as someone who's been doing grantmaking for 16 years now, I've seen waves of this 'tech is going to solve things' [mentality], and it turns out a lot of organizations get built up and given money, and then they either crash or things don't work out... [That said,] supporting new... technology that's desirable [to] communities that we should be answerable to... feels very important, too."

Participants emphasized that the temptation to demonstrate AI capabilities for particular use cases should point us to a fundamental need for attention to underlying infrastructures.

"If we are actually concerned about AI and labor, AI and healthcare, AI and [whatever else,] I don't feel it is worth our time to be sending any more money to the technical side of this... Lots of people will do that. Even if we believe that AI in healthcare is really going to be the next big frontier, then I would hope we give as much energy to preparing hospital systems to engage in that meaningfully, training up doctors to make sure that we understand the implications of this... **These [AI deployments are taking place] in a world that is full of public infrastructure that is actively crumbling**. I think [any] AI solution... if it does not have infrastructure to meet it, it's not going to be able to do the things we hope it will. My hope [for funders] is that if we really meaningfully care about AI in the world, **we'd have to spend as much effort in the world as on the AI question**."

Another highlighted that "**Institutions using AI not only need to have a system for understanding [these technologies] as relational, political, and moral tooling** – they have to develop... a system for recognizing their fitness for that purpose and either prompting their improvement and evolution; or replacing them with tooling that is fit-for-purpose."

Building Measurement Capacity

Across domains – from labor markets to election integrity to national security – participants described a **striking gap in independent measurement tools and capacity**. Technical benchmarks developed by AI companies have become the dominant way systems are evaluated, yet there is comparatively **little infrastructure for measuring broader economic, social, or democratic impacts. This gap has real consequences. If companies define the metrics, they shape what counts as risk. And if they shape what counts as risk, they influence public understanding, and regulatory priorities.** Strengthening independent measurement capacity is central to effective governance of these technologies.

"In a nutshell, we've ceded the measurement conversation to technical benchmarks at tech companies. We have no measurement systems for the effects on the economy or broader social impacts. We've had a whole set of papers [arguing that] 'We need social metrics.' Everybody kind of agrees on that… **whoever controls the measurement story controls how we understand this in the future."**

Others note the need for measurement systems to undergird policy development and enforcement: "No one's talking about evaluations… **what is your model being measured against, and how is it being measured, and who is doing the measurement. [This] is… the best way to control the outcomes of models**… I truly could not tell you how to understand if a social media algorithm is going to subvert democratic processes. And that was my job at [previous employer]. Nobody knows how to do that… **Regulation would have to go hand-in-hand with the operationalization of [measurement systems]."**

> "No one's talking about evaluations… **what is your model being measured against, and how is it being measured, and who is doing the measurement. [This] is… the best way to control the outcomes of models**… I truly could not tell you how to understand if a social media algorithm is going to subvert democratic processes. And that was my job at [previous employer]. Nobody knows how to do that… **Regulation would have to go hand-in-hand with the operationalization of [measurement systems]."**

Some noted how a **lack of shared measurement standards** and infrastructure inhibits risk mitigation: "I look at the EU AI Act risk levels, the NIST AI Risk Management Framework risk levels, and then industry also has their own, like Anthropic and OpenAI [will say, for example,] 'Oh, we've moved into risk level 3, for bio weapons'… there's maybe the beginning of some consensus [regarding] what these different tiers [mean]? [But] how should we [structure] model scorecards? **How should we measure each new model against some of these risks**?" There is a clear need for greater standardization.

Another noted: "There is a real opportunity to bring much more rigor and **empiricism around measuring security concerns of the systems**, especially as commercial systems are adopted into national security and defense contexts."

Bringing in Different Forms of Capital

Several participants discussed the **need to bring in other forms of capital that might be better positioned to drive market change:** "I'm concerned that we don't have enough go-to-market solutions we can trust…

[We need] more mission-oriented impact investors. We need [AI alternatives that]... have [things like good] governance [and] data privacy [built-in]."

Another interviewee described the value in working with mission-driven investors: "We do not do enough work with the responsible investment ecosystem... **Changing the actual structure of how products get built**... is such an underfunded, untapped opportunity. [There is a] funding desert for companies that want to responsibly grow and scale. **We do a really good job at the incubation stage, funding new ideas at like $100,000... But we do a terrible job resourcing those to scale patiently**... And [unfortunately] getting philanthropic funding in that context might not actually be helpful, because it's a concessionary signal, right... there's only so much philanthropy can do for those ideas in a way that wouldn't discredit them."

Others emphasized **the need for funders to invest not just in prototyping tech projects, but in seeing them to scale, ensuring new pro-democracy tech tools are piloted, iterated on, and then deployed widely**: "Philanthropic funding can be catalytic [in this space]... my dream... is to think about how philanthropy [can consider]... the pipeline. It's a lot easier to fund a pilot than it is to help a startup win a government contract after [it has] done pilots in a bunch of different states... **Let's look at a lifecycle here**. You have a bunch of accelerators out there, and maybe some of those accelerators are doing things for [the public sector]. Then you have [a fund] who can be taking some product ideas and applications out of accelerators and helping provide funding to pilot them in 1 to 3 states. And then, [what is] **the set of funders who are going to take the things that are successful and help them [scale]** from 3 states to 10? And then where's the investment capital to help a company get from being able to do a thing in 10 states [to] being able to compete against the Deloitte's for a contract? I think the field is... ready for that level of maturity and coordination. I'm not saying that [good benefit and service delivery, or improving government effectiveness overall] is the key to strengthening our democracy. I think that's probably a stretch... but I do think [it's important]. So, **how can investments be coordinated in a way that add up to a greater whole?**"

Getting Ahead of the Future

Many funders spoke to the need to be better prepared for the future that is coming: by imagining futures that deviate from the dominant tech company narrative; by preparing for agentic AI; and by developing better foresight capacities.

Expanding the Imagination

A small but creative cohort suggested more fundamentally rethinking the societal role of AI, and how it's being developed.

"'To AI or not to AI,' is the wrong question. The question is actually about **how we are distributing power, and who gets that power in the future**... AI is increasingly [a] Western invention of [about] five companies. I think if we're having this conversation about 'to AI or not to AI,' we're not spending enough time trying to really meaningfully engage in what's happening to the balance of global capitalism and neocolonialism. 'To AI or not to AI' is [akin to asking,] 'To give it to Microsoft, or to give it to OpenAI, or to give it to Google, or... whatever,' Another question could be, **'Do we believe AI is worthwhile in the world?' And if so, do we want to develop other potential avenues for [its development]?** Is there another universe in which it does not depend on a hyperscale data center that Elon [Musk] is going to build in the middle of [wherever]? Maybe it could be an open source [model developed] in Chile... I think we can expand our political imagination enough to [ask,] **'What is AI trying to get us to?'** Is the thing that we care about the most the speed, or the scale, or the ability to analyze a bunch of information? Because if we care about that, there are lots and lots and lots of other avenues there, and we don't need to take [seriously] these arguments that OpenAI is the only answer to that set of questions."

One expert wants to see more organizations working on the "transformative tech trilemma," which invites us to try to advance "[technical] progress, safety, and participation" in tandem. Instead, those goals can tend to operate in silos, with most work on 'progress' happening in the private sector – and, problematically, leaving the private sector to define what 'progress' means – and most work on participation happening outside of it.

Others pushed back on the concept of AI inevitability and argued for a more comprehensive reckoning with the risks of AI: "I'm hearing especially from philanthropy and civil society... there's pressure coming from above to say, 'this is inevitable.' So we have to make the best of it. We have to find and fund the best, so that we can be first to use this very powerful tool instead of saying, '**How about no?' How about we learn from the past?** [Look at] Ed Tech: [the story is] decades and decades long, and... promise after promise, and billions of dollars spent that were directly given as basically welfare to companies for procurement for [educational technology]. And we've seen the line of children's education go straight down... I want to see more of that [kind of question] and **less of the 'How are we going to use AI for democracy?'... [The question should be:] How are we going to have democracy, period?**... We are all grappling with: What are philanthropic interventions that can change the material conditions of communities being devastated by imposition of AI into their lives without consent?"

At the same time, others urged funders to seriously engage scenarios in which AI capabilities advance far more quickly and dramatically than many institutions are prepared for: "I just want everybody to **grapple with the possibility that we actually have superhuman AI in 5 years, 10 years, 15 years**. When I talk to most people who are working on something else, and I ask them to grapple with that possibility, they tend to fight the hypothetical. They tend to imagine, you know, an AI that's like ChatGPT, but a bit better, and they call that 'superhuman AI,' and I want to be like, **'No, no, no, no, imagine it is better than you at everything you do from a computer. And imagine that, on top of that, it has millions and millions of drones that it can access.** Just really live in a world where this is taking off like a rocket ship, it's going somewhere big.' That doesn't mean you believe in any particular argument for AI takeover, or authoritarian takeover, or bio-misuse, etcetera... **If I were giving advice to a funder that works in a different space, like labor markets or democracy more broadly, I would be like, 'Host a convening of your grantees, and bring in some people who are really living in this headspace,** that AI is going to change everything, and have everybody sit together and think about the implications... of that.'"

Getting Ahead of Agents

Many of the concerns raised earlier in this report – from manipulation and gradual disempowerment to national security and institutional capacity – converge most clearly in discussions of increasingly autonomous or "agentic" AI systems.

Several participants suggested that parts of philanthropy may be **underestimating the pace and implications of AI capability gains.**

"**We need a better hierarchy of harms and risks when it comes to agentic AI**... a mental model that is not fear-based but is empirically grounded. That helps one have a more balanced understanding of the benefits and harms, because the current narrative is skewed by the existential risk community [or on the other hand, just focused on current harms]. We need something in between."

Another framed the challenge in terms of democratic self-governance: "Liberal democracy presupposes that people are able to make their own decisions about important things, that they have a right to both get information from others and to delegate the decision to ... others in certain ways... **How do we make sure [an] agent isn't unduly influencing** you because a Russian bot did an adversarial attack on your AI agent, and now you're

whooped by some Russian information operation and you don't even know it. Let alone [by] Amazon or whatever… **What happens when you have a billion agents interacting with each other and self-modifying in response to the interactions they're having with other people's agents, to get our instructions implemented?** That's an ecosystem… We're technically not far along in solving that problem. That keeps me up [at night.]"

Fundamentally, participants expressed concerns that agentic capabilities are beyond current governing institutions' capacity to manage, leaving a gap that currently only the private sector is poised to fill.

"There's going to be many bona fide risks from AI… I think over time, we're going to have worlds full of very cheap to run superhuman agents, and everyone is going to have those in their pocket… how does the governance end up getting decentralized in lieu of [the] national government being able to grapple with all that?"

Another participant emphasized: "'People have been building systems that govern delegated agency for our whole existence… While the nuances vary a little, they almost all focus on: **narrow and specific-authority; explicitly articulating the interests being represented** (and often the methods used); **and accessible, independent, expert oversight**. While none of those requirements are socially exceptional, technically implementing any of them into AI tools… would require fundamental changes."

Building Foresight Capacity

While not unique to AI, several participants advocated for funders to invest more deeply in foresight capacity: "**We need to do better at predicting the range of scenarios that could happen**, and working backwards from there. I think foresight and forecasting can be a big investment area."

Others warned that traditional forecasting models are poorly calibrated to the speed and uncertainty of AI-driven change: "All of the foresight people [who] funders rely on to shape strategies… are behind the curve… **It's not about predicting the future, it's about preparing and being thoughtful**… we need to build better foresight mechanisms. We need to rebuild the field of philanthropy to be much more live, much more current."

CONCLUSION: GOVERNING AI AT THE SPEED OF DEMOCRACY

Across these interviews, one conclusion is difficult to avoid: **AI is not entering a healthy democratic system. It is interacting with institutions already under strain** – marked by economic inequality, declining trust, fragmented media systems, strained government capacity, and concentrated private power.

In some areas, **AI may intensify *longstanding vulnerabilities***. Corporate concentration of power, economic precarity, biased decision-making, and weakened journalism are not new problems. But AI risks accelerating and scaling them, embedding them into black-box systems and reducing avenues for redress. In these cases, **stronger enforcement of existing civil rights, labor, consumer protection, and antitrust laws** – alongside greater regulatory capacity – may be more urgent than creating entirely new frameworks.

Yet in other domains, **AI may introduce *new capabilities* that overwhelm the assumptions built into current governance systems**. Real-time personalized persuasion and manipulation at scale, increasingly autonomous agents, and rapid labor displacement could compress big societal changes into shorter timeframes than democratic institutions have been built to manage. Such **new capabilities could require new regulatory tools or oversight mechanisms, or new institutional designs.**

In yet other cases (and if deployed responsibly) AI could ostensibly help improve democratic processes – increasing government capacity, or improving pathways for citizen participation and deliberation.

Across all seven sections of this report, three structural gaps repeatedly surfaced:

1. **Institutional Capacity**: Many public agencies, journalists, labor organizations, and civil society groups **lack the technical expertise and staffing necessary** to use, monitor, evaluate, and govern

AI systems effectively. Pro-democracy actors are struggling to keep up with the pace of change, to use these tools where impactful, and to ensure that existing laws are enforced and any necessary laws are passed.

2. **Independent Measurement:** AI companies largely define the technical benchmarks used to evaluate system performance. **Independent infrastructure for measuring labor impacts, information harms, discriminatory effects, and security risks remain underdeveloped.** Credible third-party measurement is an essential precursor to meaningful understanding, governance, and enforcement.

3. **Political Economy and Market Concentration:** A small number of firms control all of the necessary compute, data, and infrastructure, and wield significant lobbying resources. This ownership structure and concentrated market power is driving both technological development and regulatory outcomes. **Governance cannot focus solely on technical safeguards while ignoring concentrated economic power, or guardrails will be nearly impossible to enforce**.

Several specific recommendations emerge for civil society and philanthropy:

First, funders must invest in the capacity for meaningful collective action. Democracies must be able to articulate their legitimate public interests, adapt those interests as conditions change, and enforce them in practice. In many cases that means strengthening government agencies, regulators, journalists, labor organizations, and community groups to be able to use, evaluate, and oversee AI systems effectively. Procurement reform and building sociotechnical talent pipelines, or providing enforcement support to ensure existing laws are applied, may prove as consequential as pursuing stand-alone policy.

Second, philanthropy must invest in independent measurement infrastructure. AI companies dominate the measurement infrastructure. The field needs independent ways to assess labor impacts, election integrity risks, algorithmic bias, national security concerns, and more. Standardized risk frameworks and benchmarks cannot be left solely to industry.

Third, addressing the political economy of AI must account for its ownership, its influence, and its fitness for purpose. Whether via research, advocacy, or movement building, more must be done regarding competition policy, labor protections, and corporate governance. Otherwise, market concentration will drive (or inhibit) regulatory outcomes.

Fourth, the field must prepare for faster cycles of change. Civil society is used to a slower pace of change. We will need to become more adaptive. Investments in foresight, scenario planning, building coalitional relationships across sectors, and rapid-response capacity are essential to reduce the lag between technological shifts, philanthropic action, and democratic adaptation. As some participants argued, philanthropy needs to learn to 'move fast and build things.'

Fifth, we must continue to strengthen democratic resilience beyond AI. Labor rights, local and independent journalism, civic participation, professional governance, and government modernization are not unrelated to AI governance. They are foundational to it. Strengthening them is itself an AI governance strategy.

AI will not determine democracy's future. But it will both amplify what already exists, and create as yet unseen new problems. The question is whether democratic systems can adapt quickly and thoughtfully enough to ensure these rapidly accelerating systemic changes don't undermine our capacity for self-governance.

Looking at the role of philanthropy, perhaps the positive contribution required of us is that we begin to deprioritize static strategy setting, linear problem solving, and set-it-and-forget-it policy making. Instead, those of us engaged in the work of governance may now find ourselves facing the need for **ongoing adaptation to shifting conditions and increasing uncertainty**. In such a world, more focus is needed on weaving coalitions that enable ongoing adaptation, supporting experimentation, and cultivating broader ecosystems of agency and action.

And, importantly, doing so together. As one interviewee said: "I would love to see the funding ecosystem cohere, because this is a space of such clear market failures that philanthropy... can make a really massive difference now."

REFERENCES

AI Now. "People's AI Action Plan." Accessed August 15, 2025. https://peoplesaiaction.com.

Dewey, John. *The Public and Its Problems*. Facsim. ed. Athens (Ohio): Swallow press-Ohio university press, 1991.

Dulani, Jai, and Samuel Rosado Zaidi. "The People Say No: Resisting Data Centers in the South." MediaJustice, September 9, 2025. https://mediajustice.org/resource/the-people-say-no-report/.

Eubanks, Virginia. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. First Picador edition. New York: Picador St. Martin's Press, 2019.

Gansky, Ben, and Sean Martin McDonald. "The Governance Gap: Preparing Professions for Digital Transformation." CIGI, November 27, 2023. https://www.cigionline.org/publications/the-governance-gap-preparing-professions-for-digital-transformation/.

Harris, Marci. "RECONSTITUTION, Chapter 8," 2025 Draft Manuscript.

Kulveit, Jan, Raymond Douglas, Nora Ammann, Deger Turan, David Krueger, and David Duvenaud. "Gradual Disempowerment: Systemic Existential Risks from Incremental AI Development." arXiv, January 29, 2025. https://doi.org/10.48550/arXiv.2501.16946.

McDonald, Sean Martin. "Letter from Attorneys General to OpenAI," *LinkedIn*, September 6, 2025. https://www.linkedin.com/posts/sean-martin-mcdonald-04b1b221_letter-from-attorneys-general-to-openai-activity-7370129858846433280-ld5W.

Oduro, Serena. "In Delaying Its AI Law, Colorado Shows Tech Lobby's Power In State Politics." TechPolicy. Press, August 29, 2025. https://techpolicy.press/in-delaying-its-ai-law-colorado-shows-tech-lobbys-power-in-state-politics.

Palta, Rina, Julia Angwin, and Alondra Nelson. "How We Tested Leading AI Models Performance on Election Queries." Proof News, February 27, 2024. https://www.proofnews.org/how-we-tested-leading-ai-models-performance-on-election-queries.

Pfefferkorn, Marika, Aasim Shabazz, and Michelle Shevin. "Co-Powering an Emergent Horizon." *Stanford Social Innovation Review (SSIR)*, July 14, 2022. https://ssir.org/articles/entry/co_powering_an_emergent_horizon.

Siu, Alice. "Let the Public Govern AI." Tech Policy Press, September 5, 2025. https://techpolicy.press/let-the-public-govern-ai.

APPENDIX

The following list includes more than 130 organizations, primarily focused on the United States. This is not an exhaustive list – groups included were recommended by interview participants and others close to this research.

We are open to making this a living document. If you believe an organization should be added or the work described below has changed, please email fkirby@packard.org.

***Note: Categories are imperfect – organizations listed in one category often also work in others. This list is illustrative and not intended to represent the entirety of the field.**

Contents

Auditing & Evaluation

Organizations that assess and evaluate AI systems to ensure they support the public interest.

» **Collective Intelligence Project:** The Collective Intelligence Project (CIP) is an incubator for new governance models for transformative technology. CIP focuses on the research and development of collective intelligence capabilities: decision-making technologies, processes, and institutions that expand a group's capacity to construct and cooperate towards shared goals. They apply these capabilities to transformative technology: technological advances with a high likelihood of significantly altering society.

» **Humane Intelligence:** Humane Intelligence is building a community of practice around algorithmic evaluations through its programming platform for model evaluators and individuals seeking to learn more about model evaluations. By creating this community and practice space, they aim to professionalize the practice of algorithmic auditing and evaluations.

» **METR:** Model Evaluation & Threat Research (METR)'s mission is to develop scientific methods to assess catastrophic risks stemming from AI systems' autonomous capabilities and enable good decision-making about their development. To that end, they evaluate frontier AI models to help companies and wider society understand AI capabilities and what risks they pose. METR also prototypes governance approaches which use AI systems' measured or forecasted capabilities to determine when better risk mitigations are needed for further scaling.

» **OCEAN:** ORCAA Collaborative Expert Assistance Network: OCEAN uses algorithmic auditing expertise to support those harmed by AI and make AI systems more accountable. They do this work by providing expert analysis in lawsuits on behalf of people harmed by AI, educating to build capacity in the legal system for investigating AI and algorithmic systems, and providing AI policy assistance at the state and local level.

Civic Participation[22]

Organizations that use AI to promote collaboration, deliberation, civic engagement, and pluralism.

» **All Tech is Human:** All Tech Is Human has brought together individuals and organizations across civil society, government, and industry. They aim to grow the responsible tech field by promoting knowledge-sharing and collaboration among multiple stakeholders to co-create a better tech future.

» **Deliberative Democracy Lab at Stanford:** The Deliberative Democracy Lab (formerly the Center for Deliberative Democracy), housed within the Center on Democracy, Development and the Rule of Law at Stanford University, is devoted to research about democracy and public opinion obtained through Deliberative Polling®.

» **Digital Public:** Digital Public aims to help communities, governments, non-profits, and technology companies engaged in digital transformation to design contextual, legal, digital infrastructure for participation and governance.

» **The GovLab:** The Governance Lab is an action research center whose mission is to improve people's lives by transforming how institutions govern. Its work focuses on designing, deploying, testing

---

22   There are also a number of organizations focused on grassroots movement-building with particular constituencies, such as faith, veterans, climate activists, or parents and youth (eg. Common Sense Media, Mothers Against Media Addiction), which may not be fully covered here.

and learning from innovative methods of using data and technology -- including AI -- to help public institutions better understand and respond to the people they serve.

» **Harvard Ash Center / Nonviolent Action Lab:** The Nonviolent Action Lab is an innovation hub for research on advancing democracy worldwide through civil resistance. The Lab produces and disseminates up-to-date knowledge on nonviolent action, how it works, global trends in success and failure, trends in political violence and state repression, and analysis of these trends.

» **New_ Public:** New_ Public strives to ensure that online communities foster connection, facilitate pluralism, and encourage civic engagement. They frame their work as digital urban planning to imagine a thriving social infrastructure on the internet.

» **Google Jigsaw:** Jigsaw is an incubator within Google that builds and invests in technologies to give people greater agency in the world around them.

Civil Rights

Organizations using research and advocacy to advance equity in the design, governance, and use of AI technologies.

» **Algorithmic Justice League:** The Algorithmic Justice League (AJL) combines art and research to illuminate the social implications and harms of artificial intelligence. AJL's mission is to raise public awareness about the impacts of AI, equip advocates with resources to bolster campaigns, build the voice and choice of the most impacted communities, and galvanize researchers, policymakers, and industry practitioners to prevent AI harms.

» **Center for Critical Internet Inquiry (C2i2) at UCLA:** The Center for Critical Internet Inquiry (C2i2) is a critical internet studies community committed to reimagining technology, championing social justice, and strengthening human rights through research, culture, and public policy. C2i2 initiatives include the Minderoo Initiative on Technology and Power, the Critical Internet Studies Program, Digital Civil Rights, and the Data, Power, and Racial Justice Initiative.

» **Fight for the Future:** Fight for the Future is a queer-woman led team of activists, strategists, and technologists who have helped organize some of the largest online protest actions in history, including campaigns against SOPA / PIPA and for net neutrality. Through high-visibility campaigns they defend human and civil rights online, advance an accessible, just, and censorship-free Internet, and demand that tech regulations do not come at the expense of the vulnerable.

» **IDA B. WELLS Just Data Lab at Princeton:** The IDA B. WELLS Just Data Lab brings together students, educators, activists, and artists to develop a critical and creative approach to data conception, production, and circulation. Their aim is to rethink and retool data for justice.

» **Leadership Conference on Civil and Human Rights' Center for Civil Rights and Technology:** The Leadership Conference on Civil and Human Rights is a coalition of more than 240 national organizations fighting to protect, defend, and expand the rights of every person in the United States. Born out of the civil rights movement, they have been a convener, thought leader, and advocate for public policies that work to make the U.S. more fair and more free.

» **Upturn:** Upturn challenges the material harms of AI, algorithms, and data on people's lives. Their research and advocacy interrogates oppressive uses of technology through a historical and political lens. In particular, they focus on the real-world consequences of AI systems on housing, labor, healthcare, financial security, and other important aspects of people's lives. They partner with local and national

advocates, organizers, litigators, academics, and policymakers to address these immediate harms, while supporting movements building toward a just economic system that meets people's needs.

» **Tech Policy Lab:** (*See Generalist / Multi-Issue section*)

Effective Government

Organizations that seek to improve government effectiveness through the responsible adoption of AI and other technologies.

» **Future Caucus:** Future Caucus's mission is to activate young leaders to bridge the partisan divide and transform American politics. They work directly with young policymakers on both a federal and state level throughout the U.S., providing them with the resources and support to develop and pass policy solutions and forge partnerships on the issues impacting Americans.

» **GovAI Coalition:** The GovAI coalition is a group of 3,000 members from over 900 local, state, and federal agencies and over 200 industry and civil society partners united to promote responsible and purposeful AI in the public sector. Their mission is to enable public sector AI to serve the public good through resource sharing, standards, and advocacy.

» **InnovateUS:** InnovateUS provides no-cost, live and at-your-own pace learning on data, digital, innovation and AI skills for public professionals. They work with government leaders to design, develop and deploy skill-based learning. Their mission is to improve lives by enabling public service professionals to deliver more effective and equitable policies and services.

» **MITRE:** MITRE believes AI is vital for advancing the public good, enhancing decision-making, and modernizing government services. They support efforts across the federal government to plan and implement AI solutions, including through their Federal AI Sandbox, a tool developed with NVIDIA to train LLMs for government tasks.

» **Partnership for Public Service AI Center for Government:** The Partnership for Public Service AI Center for Government is committed to fostering AI leadership, talent and literacy across federal, state and local governments by developing AI leaders, strengthening AI talent capacity, increasing AI literacy and community learning and celebrating agencies using AI effectively.

» **POPVOX Foundation:** POPVOX Foundation works to help democratic institutions keep pace with a rapidly changing world. Through publications, events, prototypes, and technical assistance, the organization helps public servants and elected officials better serve their constituents and make better policy. This includes regular AI trainings for legislative staff in the US and globally and the development of AI-enabled tools to improve legislative capacity.

» **Stanford RegLab:** RegLab partners with government agencies to design and evaluate programs, policies, and technologies that modernize government. They are an interdisciplinary team of engineers, data scientists, social scientists, and legal experts building an evidence base and high-impact demonstration projects for better government. Their home at Stanford allows them to leverage advances in AI to build a bridge between government agencies and frontier technology.

» **Tech Talent Project:** The Tech Talent Project is dedicated to increasing the ability of federal, state, and local government agencies to recruit modern technical leaders in order to achieve critical economic, policy, and human outcomes. It was founded in 2017 by former technology leaders from the Obama and Trump White Houses.

» **TechCongress:** TechCongress works toward building a critical mass of technology expertise in government. Since 2015, their fellowships have placed nearly 150 technologists in Democratic and Republican congressional offices to advise on technology issues, from writing the rules on AI to delivering SNAP benefits. They are piloting a state legislative technology fellowship in Connecticut, which includes managing the biweekly State AI Policy Forum with over 250 state lawmakers from 48 states to align AI policy frameworks. They aim to scale this model to 50 states by 2031.

» **Recoding America Fund:** The Recoding America Fund (RAF) is a six-year philanthropic initiative aimed at reforming governments at both the federal and state levels. With partners, it is supporting efforts to reform the U.S.'s civil service systems; change how governments invest in, develop, and oversee digital infrastructure; reform government procedures; and demonstrate a new operating model that reimagines the relationship between the legislative and executive branches to improve outcomes.

» **(Data & Society)** *Public Tech Leadership Collaborative (see Generalist / Multi-Issue section)*

» **Tech Policy Lab** *(see Generalist / Multi-Issue section)*

» **U.S. Digital Response** *(see Generalist / Multi-Issue section)*

Elections & Structural Democracy

Organizations that analyze and address the impacts of AI on elections.

» **ASU Mechanics of Democracy Lab:** The Mechanics of Democracy Laboratory (MODL) at Arizona State University focuses on improving the quality, transparency and resilience of elections across the United States. MODL provides trainings for election officials, conducts research into improving systems and processes and engages in public education to counter misinformation. MODL's AI + Elections Clinic addresses the opportunities and risks of emerging technologies in voting and election management.

» **(Brennan Center for Justice) AI and Elections:** The Brennan Center sees AI as a powerful tool with the potential to both help and hurt democracy. As they work to identify the dangers and opportunities AI poses, they are detailing steps the public and private sectors should take to address the effects of AI on elections and democratic governance.

» **Center for Constitutional Democracy at University of California, College of the Law, San Francisco:** The Center for Constitutional Democracy brings together scholars, practitioners, and community leaders for innovative work across legal strategy to defend and advance rights, justice, and democracy. Priorities include anticipating and proactively addressing threats to constitutional rights and democratic institutions in an era of rapid AI and other technological change.

» **California Initiative for Technology and Democracy (CITED):** The California Initiative for Technology and Democracy seeks state-level solutions to the threats that disinformation, AI, deepfakes, and other emerging technologies pose to democracy and elections. CITED leads a legislative portfolio every year, provides public leadership and strategic communications, hosts community events to raise the visibility of these issues and educate the public, engages experts nationally and in other states, and works in coalition with civil rights, civic engagement, GOTV, and powerbuilding partners throughout California.

» **Multiracial Democracy Project at GW Law School:** The Multiracial Democracy Project (the "Project") is an initiative of the George Washington University Law School. Through research, analysis, convenings, and public education, the Project produces solutions that facilitate the U.S.'s transition to a well-functioning

democracy that truly represents an increasingly diverse population. The Project serves as a bridge between academic, civil rights, democracy reform, technology, and policy communities.

» **Public Citizen :***(see Generalist / Multi-Issue section)*

Journalism & Media

» Using AI to improve the practice of journalism

  » **Gigafact Parser:** Gigafact Parser is a civic platform for news professionals to discover what politicians are saying. AI transcribes and analyzes audio and video recordings of public comments by public officials to surface their claims and talking points.

» Journalism to hold AI companies accountable

  » **Coalition for Independent Tech Research:** The Coalition for Independent Technology Research works to advance, defend, and sustain the right to study the impact of technology on society. They are a coalition of academics, journalists, civil society researchers, and community scientists committed to advocating for and organizing in defense of research that is ethical, transparent, and privacy-preserving. Coalition members are, and will remain, independent of the technology industry they study.

  » **Indiana University's Observatory on Social Media:** The Observatory on Social Media is an interdisciplinary research center at Indiana University. OSoMe unites data scientists and journalists in studying the role of media and technology in society, and building tools to analyze and counter manipulation on social media.

  » **The Maybe:** The Maybe works to dismantle dominant narratives and elevate new voices to expand what's seen as possible. They build strategies, initiatives, and audiences to make that real—from deep research and storytelling to connection and roadmap development, they help thinkers, leaders, and organizers to shape what happens next.

  » **MediaJustice:** MediaJustice builds power to challenge how corporations and governments use media and technology to shape our collective future. They connect movements fighting corporate control, analyze how tech enables harmful policies, and amplify community-led alternatives to surveillance and extraction.

  » **Proof News:** Proof is a newsroom convening journalists to hold AI and Big Tech accountable. This year, the team is focused exclusively on AI's effects on work, reporting stories about what is happening to workers right now. Proof conducts investigations, builds public datasets, collaborates with other newsrooms, researchers and influencers, and promotes best practices in evidence gathering. It is committed to telling stories on social and video platforms to reach new audiences.

  » **Tech Policy Press:** Tech Policy Press publishes opinion, analysis, and news. At a time of great challenge to democracies globally, they seek to advance a pro-democracy movement in tech and tech policy.

  » **Tech Transparency Project (Campaign for Accountability):** TTP is a research initiative of Campaign for Accountability (CfA). This year the Tech Transparency Project is focused on issues that resonate

across the political spectrum including tech's impact on children, national security, veterans, and older adults.

» **WITNESS:** WITNESS helps people use video and technology to protect and defend human rights. They identify critical situations and teach those affected by them the basics of video production, safe and ethical filming techniques, and advocacy strategies.

» Training journalists to better cover AI

» **Pulitzer AI Accountability Network:** The Artificial Intelligence Accountability Network supports and brings together journalists reporting on AI and with AI globally. Working with journalists and newsrooms that represent the diversity of the communities affected by AI technologies, the Network seeks to address the knowledge imbalance on artificial intelligence that exists in the journalism industry and to create a multidisciplinary and collaborative ecosystem that enables journalists to report on this fast-evolving topic with skill, nuance, and impact.

» **Tarbell Center for AI Journalism:** The Tarbell Center for AI Journalism supports journalism that helps society navigate the development and deployment of advanced AI. They provide funding and training to strengthen AI reporting at major news outlets. Their fellowship is a one-year program for journalists, providing 3 months of full-time training on AI and a fully-funded placement at outlets such as Bloomberg, The Guardian, and TIME. Their grants program provides funding directly to newsrooms to build dedicated AI reporting capacity — supporting investigative teams, new verticals, and in-depth coverage of AI's societal implications.

Labor

Organizations analyzing and seeking to mitigate the impacts of AI on workers and labor markets.

» **Budget Lab at Yale:** The Budget Lab is a research group that focuses on tracking the impact of AI on the labor market and helping people to understand who is being impacted, how quickly, and what will affect that. They are also starting a workstream on the fiscal impact of AI, focused not on the debates about the growth side, but on how AI could affect both spending and revenue, and how to adjust our tax code in response.

» **Center for Labor and a Just Economy:** The Center for Labor and a Just Economy (CLJE) at Harvard Law School is a hub of collaborative research, policy, and strategies to empower working people to build an equitable economy and democracy. Their program on tech and labor is engaged in long-term participatory research with working class women on the impact of Artificial Intelligence on both paid and unpaid labor, as well as researching potential disruptions by cryptocurrency and other speculative industries on working class well-being.

» **CoWorker.org:** CoWorker.org is a laboratory for workers to experiment with power-building strategies and win meaningful changes in the 21st century economy.

» **Distributed AI Research Institute (DAIR):** DAIR is an interdisciplinary and globally distributed AI research institute rooted in the belief that AI is not inevitable, its harms are preventable, and when its production and deployment include diverse perspectives and deliberate processes it can be beneficial. Their research reflects their lived experiences and centers their communities.

» **Jobs for the Future:** Jobs for the Future transforms education and workforce systems to drive success for all, including populations who face barriers to economic advancement.

» **MIT Initiative on the Digital Economy:** The MIT Initiative on the Digital Economy (IDE) is a team of thought leaders and researchers examining how people and businesses work, interact, and will ultimately prosper in a time of rapid digital transformation.

» **National Domestic Workers Alliance:** The National Domestic Workers Alliance works for the respect, recognition, and rights of domestic workers throughout the U.S. They organize domestic workers and develop leaders, develop policy solutions, work to shift key narratives, use technology to find new ways to make work better for domestic workers, and mobilize voters.

» **Rideshare Drivers United:** Founded in 2018 by a group of LA drivers, Rideshare Drivers United is a driver-led organization. Their values are deeply rooted in worker democracy, grassroots action, and labor rights. They've grown from 400 members in Fall 2018 to over 20,000 California driver members today. They have built three California-based chapters, in LA, San Diego, and the Bay Area.

» **Stanford Digital Economy Lab:** They pursue a deeper understanding of the digital economy and its impact on the future of work and society. They focus on ensuring that as machines become more capable, humans become more empowered, not displaced. To that end, they conduct research that helps companies, policymakers, and communities navigate technological change. Their work spans four key areas: the economics of transformative AI, digital platforms and society, new measures of the economy, and AI agents.

» **TechEquity:** TechEquity focuses on the intersection of AI and economic equity, with a particular focus on labor / workforce, housing, affordability, and consumer protection. It uses a three-part theory of change that includes building public support through strategic communications and original research, advocating for legislation and regulation in California and other states, and doing implementation / enforcement work at California administrative agencies. It leads the public interest coalition on AI policy in California with 40+ members and sees California as a way to leverage nationwide change.

» **UC Berkeley Labor Center:** The UC Berkeley Labor Center's Technology and Work Program provides worker organizations and policymakers with the research, policy analysis, and training they need in order to respond to rapid technological change in the workplace and ensure that AI and other digital technologies benefit rather than harm workers. They focus on low-wage industries and the workers of color, women, and immigrants who are often on the frontlines of experimentation with emerging technologies.

» **The Worker Agency:** The Worker Agency's mission is to win campaigns that change people's lives for the better. They work to build campaigns that drive people-first solutions to today's problems, and collaborate with progressive organizations, social justice, and labor groups to amplify the impacts of their work.

» Brookings **AI & Work** initiative (see Generalist / Multi-Issue section)

» **(Burnes Center for Social Change) Labor and Economic Justice Initiative:** (see Generalist / Multi-Issue section)

» **(Data & Society) Labor Futures Program:** (see Generalist / Multi-Issue section)

Material Infrastructure & Environment

Organizations that examine and respond to the environmental, infrastructure, and community impacts of AI systems.

» **Rooted Futures Lab:** Rooted Futures Lab is dedicated to centering environmental justice in technology. Their mission is to address the systemic neglect and exclusion of environmental justice in technology development and application. They envision a world where technology supports human and planetary well-being for the global majority.

» **Open Power AI Initiative (Electric Power Research Institute):** The Electric Power Research Institute (EPRI) provides a stable, science-based foundation for how electricity systems are planned, governed, and operated around the world. As artificial intelligence increasingly shapes energy infrastructure investment and system planning, these tools will have real consequences for communities, affordability, and access to reliable power. Open Power AI was established to help ensure that these decisions are grounded in transparent analysis, public-interest accountability, and attention to real-world impacts.

» **(Data & Society) Trustworthy Infrastructures Program:** (see Generalist / Multi-Issue section)

» **Tech Policy Lab:** (see Generalist / Multi-Issue section)

National Security and International Relations

Organizations that analyze and seek to mitigate risk related to AI, national security, and geopolitics.

» **Americans for Responsible Innovation:** Americans for Responsible Innovation works to help policymakers develop a governance framework for rapidly advancing technologies that protects the public from harm while continuing to foster innovation. Founded by Brad Carson (former Congressman, Under Secretary of the Army, Acting Under Secretary of Defense), their work focuses on guardrails, national security, and emerging risks. ARI has been heavily engaged in supporting Aspen's Congressional delegations on AI. ARI has a sister organization, Center for Responsible Innovation, which focuses on promoting responsible innovation, developing actionable and politically feasible policy ideas, and educating policymakers on AI.

» **Brennan Center for Justice, Liberty and National Security (LNS) Program:** The Brennan Center for Justice is a law and policy institute that works to defend and revitalize systems of democracy and justice in the United States. The LNS program analyzes AI's impact on national security policy and practice, and the ensuing risks to civil liberties, civil rights, and democratic values. It develops recommendations on how legislative and regulatory frameworks and case law should evolve to meet these challenges.

» **Center for a New American Security (CNAS):** The Center for a New American Security (CNAS) develops national security and defense policies. CNAS engages policymakers, experts, and the public with research, ideas, and analysis to shape and elevate the national security debate.

» **Carnegie Endowment for International Peace:** Scholars generate ideas and analysis to help inform countries, institutions, and leaders as they take on global problems. They partner with governments, industry, academia, and civil society to address the governance, international security, and democracy challenges from AI. By confronting both the short-term and medium-term challenges, they aim to mitigate the most urgent risks of AI while laying the groundwork for addressing its broader and subtler effects.

» **The Citizen Lab at U Toronto:** The Citizen Lab is an interdisciplinary laboratory based at the Munk School of Global Affairs & Public Policy, University of Toronto, focusing on research, development, and

high-level strategic policy and legal engagement at the intersection of information and communication technologies, human rights, and global security.

» **Transnational Institute:** The Transnational Institute (TNI) is an international research and advocacy institute committed to building a just, democratic, and sustainable planet. For over 50 years, TNI has served as a nexus between social movements, scholars, and policymakers. Their work includes analyses of the political economy of AI and strategies to bring technology under popular control.

Organizing & Community Power

Organizations that support community-led responses to AI deployment and governance.

» **Citizens and Technology Lab (CAT Lab) at Cornell:** Works with communities to study the effects of technology on society and test ideas for changing digital spaces to better serve the public interest. They envision a world where digital power is guided by evidence and accountable to the public. They seek to enable anyone to engage critically with the tech tools and platforms they use. Working hand-in-hand with diverse communities and organizations around the world, they identify issues of shared concern related to digital discourse, digital rights and consumer protection.

» **Detroit Community Technology Project:** Rooted in the Detroit Digital Justice Principles, Detroit Community Technology Project's (DCTP) mission is to use and create technology rooted in community needs that strengthens neighbors' connection to each other, and to the planet.

» **Encode:** Encode helps policymakers and the public navigate AI. They work to support humanity in this time of unprecedented transition through legislative advocacy and public education.

» **Kairos:** Kairos runs campaigns intended to allow users to see themselves not just as participants in the digital realm, but as co-creators and constituents with the ability to move platforms and ultimately win stronger regulation of the internet. All of their campaigns are designed to ensure that technology and the internet are in service of Black and brown liberation and not used to uphold systemic injustice.

» **Science, Technology, and Public Policy (STPP) at UMich:** University of Michigan's Science, Technology, and Public Policy (STPP) Program is dedicated to advancing the public interest, and specifically social equity and justice, in the development and governance of science and technology through education, research, and community and policy engagement. It produces briefs and reports for civil society and governments on a range of topics, including an AI Handbook for Local Governments, how communities can resist data centers, and the social, equity, and environmental implications of LLMs.

» **Surveillance Resistance Lab (part of the Collaborative Research Center for Resilience):** The Surveillance Resistance Lab publishes resources to scale up the ability of communities, workers, and advocates to build power towards a vibrant democracy.

» **TechTonic Justice:** TechTonic Justice is here to build long-term power among the people and communities that AI leaves behind. They strengthen local justice movements by working alongside community partners and frontline advocates to ensure those at the ground-level are able to effectively respond to complex AI harms and organize against their use – both today and in the future.

» **Twin Cities Innovation Alliance:** Twin Cities Innovation Alliance (TCIA) is a coalition of stakeholders representing a cross sector of public, private and community organizations, corporations and institutions. Their mission is to build and develop a critical mass of diverse, highly engaged residents, policy makers, and entrepreneurs, made up of minorities and people of color traditionally identified as the end users and consumers of innovation and design, transforming them into purveyors and beneficiaries.

Political Economy

Organizations that analyze the risks and work to mitigate the impacts of AI with a competitiveness / antitrust and / or economic justice perspective.

- » **Ada Lovelace Institute:** The Ada Lovelace Institute is an independent research institute with a mission to ensure that data and AI work for people and society. They believe that a world where data and AI work for people and society is a world in which the opportunities, benefits and privileges generated by data and AI are justly and equitably distributed and experienced. They recognize the power asymmetries that exist in ethical and legal debates around the development of data-driven technologies, and will represent people in those conversations. They focus not on the types of technologies they want to build, but on the types of societies they want to build.

- » **AI Now:** Founded in 2017, the AI Now Institute produces diagnosis and policy research on artificial intelligence. AI Now develops policy strategies to redirect away from the current trajectory: unbridled commercial surveillance, consolidation of power in very few companies, and a lack of public accountability.

- » **Open Markets Institute:** The Open Markets Institute works to help people relearn how to use competition policy to build stronger democracies, more just and equitable societies, more innovative and sustainable economies, and a more secure and peaceful world. By combining policy, legal, and market structure expertise with sophisticated communications and outreach efforts, Open Markets seeks not only to hold today's monopolies accountable for abuse of power, but to rebuild an economic system where progress is easier to achieve, because power is far more widely and equitably distributed.

- » Political Economy of Tech Working Group: Led by **danah boyd**, a network of scholars are conducting research and coordinating experts to interrogate the political economy in which AI is operating. The goal of this work is to identify novel levers of change to address potential societal-level harms from the mass deployment of AI in a wide array of sectors and communities.

- » **Public Citizen:** (see Generalist / Multi-Issue section)

Privacy, Due Process & Access to Justice

Organizations defending privacy and civil liberties in the digital age.

- » **Center for Information Technology Policy (CITP) at Princeton:** Princeton University's Center for Information Technology Policy (CITP) is a nexus of expertise in technology, engineering, public policy, and the social sciences. Their researchers work to better understand and improve the relationship between technology and society, with a particular focus on AI, Data Science and Society, Privacy and Security, and Digital Infrastructure and Platforms.

- » **Electronic Frontier Foundation:** The Electronic Frontier Foundation (EFF) defends civil liberties in the digital world. Across a broad portfolio of work, EFF focuses on how AI systems are reshaping democratic institutions, labor, education, culture, and public accountability. Their work helps policymakers, advocates, and courts understand the real-world consequences of AI adoption before it becomes entrenched infrastructure.

- » **Electronic Privacy Information Center:** EPIC has been an independent watchdog fighting to defend human rights in the digital age for more than three decades. They take on Big Tech and government overreach through advocacy, litigation, and research. They believe the mass extraction and misuse of

personal data by government and corporate powers destroys privacy and undermines democratic governance. They provide expert guidance to policymakers at the state and federal level on proposals to ensure that AI systems serve the public interest and support democratic institutions.

» **Center on Privacy and Technology at Georgetown Law:** The Privacy Center conducts research and advocacy to expose and oppose digital era government and corporate surveillance. The Center defends privacy not only as an individual right, but as a public good which is necessary for democratic self-governance. Through research, advocacy, community partnerships, litigation, and education, the Privacy Center supports global movements against the use of data to fuel systems of economic and political control.

» **Just Futures Law:** Just Futures Law seeks to transform how litigation and legal support serves communities and builds movement power. They are experienced legal strategists who are part of a larger organizing vision led by impacted people.

» **National Association of Criminal Defense Lawyers (NACDL):** NACDL's mission is to serve as a leader, alongside diverse coalitions, in identifying and reforming flaws and inequities in the criminal legal system, and redressing systemic racism, and ensuring that its members and others in the criminal defense bar are fully equipped to serve all accused persons at the highest level.

» **Tech Policy Lab:** (See Generalist / Multi-Issue section)

Safety & Systemic Risk

Organizations that research and seek to mitigate large-scale and systemic risks from AI systems.

» **Center for AI Safety:** The Center for AI Safety is a San Francisco-based research and field-building organization. They believe that artificial intelligence has the potential to profoundly benefit the world, provided that it can be developed and used safely. Their mission is to reduce societal-scale risks from artificial intelligence.

» **Center for Secure Emerging Technologies (CSET) at Georgetown:** A policy research organization within Georgetown University's Walsh School of Foreign Service, CSET provides decision-makers with data-driven analysis on the security implications of emerging technologies. CSET is currently focusing on the effects of progress in artificial intelligence, advanced computing, and biotechnology. They seek to prepare a new generation of decision-makers to address the challenges and opportunities of emerging technologies.

» **Future of Life Institute:** The Future of Life Institute's mission is to steer transformative technologies away from extreme, large-scale risks and towards benefiting life. They support efforts to develop institutions and visions capable of managing the risks of new technologies.

» **Golden Gate Institute for AI:** The Golden Gate Institute for AI is a think tank dedicated to addressing the most critical challenges of the transition to a world with powerful AI. They support leaders as they navigate this high-stakes period by bridging disciplines, convening experts, and publishing accessible analysis of AI's toughest questions.

Talent and Technical Assistance

Organizations that build technical capacity and provide implementation support for responsible AI adoption.

- » **Center for Civic Futures:** Center for Civic Futures provides spaces for government leaders to learn, test, and act. Through pilots, partnerships, and shared insight, they help decision-makers move ideas from exploration to evidence, and evidence into practice. Center for Civic Futures includes the AI Readiness Project, a national research and advisory institute focused on technology policy and best practices in state and local government; and the Public Benefit Innovation Fund, which deploys dollars focused on real world experimentation identified by the needs and priorities of leaders.

- » **Coding it Forward:** Coding it Forward connects AI-native talent to the social sector. Since 2017, Coding it Forward has connected over 700 young people with 80+ local, state, and federal government offices nationwide to provide the technical capacity a modern democracy requires. Their greater movement has reached over 30,000 early-career technologists dedicated to ensuring that the most powerful tools are governed and deployed by a representative workforce committed to the public good.

- » **Foundation for Civic AI:** Recent advances in AI technology have created tremendous new opportunities for non-profit organizations to leverage data to support their mission and increase their impact. Foundation for Civic AI provides its grantees access to this new AI infrastructure at no cost.

- » **Horizon Institute for Public Service:** The Horizon Institute supports the US government in building the technical expertise it needs to keep pace with emerging technologies like AI by helping technology experts move into policy and public service careers. Horizon's programs span the career pipeline— from early exploration to placement—connecting talent with policy opportunities inside and outside government. Its programs include a fellowship that places experts in the executive branch, congressional offices, and think tanks; online resources that reach 100,000+ people per year; a suite of in-person events and workshops; and the AI Policy Leadership Network.

- » **Human Rights Data Analysis Group (HRDAG):** Provides scientific analysis of human rights data for their partners seeking justice and accountability in the United States and around the world.

- » **(Burnes Center for Social Change):** AI for Impact (See Generalist / Multi-Issue section)

- » **U.S. Digital Response:** (See Generalist / Multi-Issue section)

Generalist / Multi-Issue

Organizations working across multiple issue areas to advance technology policy and democratic governance broadly, including governance of AI.

- » **AI & Democracy Foundation:** They are a governance innovation and infrastructure organization focused on ensuring that democracy can keep pace with AI. They believe that without rapid action, AI advances will lead to immense power concentration and chaotic destabilization. They are working toward high-quality deliberative democratic processes for decision-making on and with AI that can provide an alternative path forward and help mitigate these risks — enabling us to maintain democracy in a world increasingly dominated by AI.

- » **AI Futures Project:** The AI Futures Project is a research group forecasting the future of AI. Their work includes AI 2027, a scenario that represents a best guess on what the future of AI might look like.

» **AI Objectives Institute:** The AI Objectives Institute (AOI) is a research lab that conducts economics research, advises governments, and builds tools for policymakers to ensure AI and economic systems serve genuine human objectives, enabled by broad public input and scalable cooperation.

» **Allen Lab for Democracy Renovation, Harvard Ash Center:** The Allen Lab for Democracy Renovation works to ensure that public policy, political institutions, and the technologies that support them, including artificial intelligence, are designed and judged by how well they strengthen constitutional democracy—expanding freedom and political equality, building fully inclusive institutions, and widening avenues for participation and connection, all rooted in the conditions people need to flourish.

» **American Economic Liberties Project:** The American Economic Liberties Project launched in February 2020 to help translate the intellectual victories of the anti-monopoly movement into momentum towards concrete, wide-ranging policy changes that begin to address today's crisis of concentrated economic power. They work with a network of allies to call on government to use policy tools to challenge monopolies' dominance over markets and society.

» **Aspen Digital:** A program of the Aspen Institute, Aspen Digital convenes thinkers and doers to anticipate how the world is changing in the age of AI – including in security, governance, media, workplace, creativity and more. Initiatives include cross sectoral gatherings that drive change and thought leadership that advances the public interest.

» **Brookings:** Brookings is a think tank focused on research, analysis, and education on public policy issues. The **Brookings AI & Work initiative** has three mutually reinforcing elements: a storytelling series that centers workers' experiences and perspectives; innovative quantitative and qualitative analysis to understand the risks and opportunities of AI and work; and cross-sector convenings to develop policy solutions and best practices that ensure equitable benefits from AI advancements and mitigate potential harms. The **Brookings Artificial Intelligence and Emerging Technology Initiative** seeks to shape the future of AI policy – advancing research, collaboration, and institutional innovation. Major areas of focus include: AI governance and its impact on democracy; economic transformation from AI; AI impact on global development and geopolitics; developing standards for maximizing benefits and minimizing harms; and understanding frontiers of AI innovation and future capabilities.

» **Burnes Center for Social Change at Northeastern:** The Burnes Center works to develop innovative, participatory, and equitable approaches to solving public problems using new technology. The **AI for Impact** program develops and deploys AI products for use in government agencies and civic organizations. Each tool addresses a specific operational challenge. Built by full-time AI for Impact fellows, these products offer a practical model for human- and community-centered AI development inside public institutions. Through its Labor & Economic Justice initiative, the Burnes Center is using AI tools to help states eliminate unnecessary college degree requirements in state employment and working with several national unions to launch an AI training program for worker organizers. The Labor Initiative has extensively studied and discussed AI's effects on workers and work power through its independent media channel **Power At Work.**

» **The Center for an Informed Public:** The University of Washington's Center for an Informed Public researches the impacts of AI on elections, including deepfakes, bots, and social media algorithms. CIP has also hosted multiple convenings with journalists on how to cover mis- and disinformation, including around AI. For example, they are hosting Kashmir Hill (New York Times) and Jeff Horwitz (Reuters) this April to discuss the harms of chatbots.

» **Center for Democracy and Technology (CDT):** The Center for Democracy & Technology (CDT) fights to advance civil rights and civil liberties in the digital age. They shape technology policy, governance, and design with a focus on equity and democratic values.

» **Change.org:** Best known as a petition platform, Change.org is also incubating a deliberative democracy platform.

» **Common Good AI:** Common Good AI's mission is to foster inclusive civic engagement by transforming how communities find common ground and solve problems together using new digital tools. The organization partners with non-profits, foundations, academic institutions, and other organizations to design end-to-end engagements using new deliberative tools for activities including local agenda setting, coalition building, deliberation, dialogue, or decision-making.

» **Data & Society:** Data & Society studies the social implications of data, automation, and AI, producing research to ground public debate about emerging technology. The **Labor Futures Program** uses empirical research and targeted engagement with stakeholders and decision-makers to create opportunities and levers for workers to shape the technologies that impact their everyday lives. The **Trustworthy Infrastructures Program** team researches community-driven responses to technology's entrance into the most intimate parts of people's working, material, personal, and public lives. The **Public Tech Leadership Collaborative** is a peer learning community anchored by ten research institutions and including scholars, practitioners, and state, local, and federal government leaders who cultivate trust-based relationships to ensure that government use and procurement of data and technology serve the public interest.

» **Epoch AI:** Epoch AI is a multidisciplinary research institute investigating the future of artificial intelligence. They examine the driving forces behind AI and forecast its economic and societal impact.

» **(Foundation for American Innovation) Artificial Intelligence Policy:** The Foundation for American Innovation champions the technology, talent, and ideas essential to American prosperity, security, and flourishing. They note that while AI promises to reshape American life, this will demand the evolution of existing economic and political systems, requiring collaboration between industry and government.

» **GovAI Hub:** Social Finance, in partnership with HumanServices.ai, launched the AI Learning and Innovation Hub in 2025 to help government agencies test, refine, and scale AI solutions. Agencies participate in cohorts of 4-6 organizations over a six-month period, collaborating and learning from one another along the way.

» **Institute for Advanced Study's Science, Technology, and Social Values Lab:** Founded in 2023 by Harold F. Linder Professor Alondra Nelson, the Science, Technology, and Social Values (ST&SV) Lab explores emerging scientific and technological phenomena (including artificial intelligence, genomics, public health, and space) and their intersections with the frustration and fulfillment of civil, political, and social rights.

» **Knight First Amendment Institute at Columbia:** The Knight First Amendment Institute defends the freedoms of speech and the press in the digital age through strategic litigation, research, policy advocacy, and public education. The Institute aims to promote a system of free expression that is open and inclusive, that broadens and elevates public discourse, and that fosters creativity, accountability, and effective self-government.

» **Metagov:** Metagov cultivates tools, practices, and communities that enable self-governance in the digital age. They work toward a governance layer for the internet that is empowering, creative,

interconnected, and accountable through a consortium of projects in the U.S. and Scotland on digital governance of and with AI.

» **National Conference on State Legislatures (NCSL):** The National Conference of State Legislatures, created by state legislators and legislative staff in 1975, serves America's 50 states, commonwealths, territories and the District of Columbia. Its mission is to advance the effectiveness, independence and integrity of legislatures and to foster interstate cooperation and facilitate the exchange of information among legislatures. Every state legislator and staffer is a member of the organization and has complete access to the latest in bipartisan policy research, training resources and technical assistance tailored specifically to their needs.

» **Partnership on AI (PAI):** Partnership on AI (PAI) is a partnership of organizations across academia, civil society, tech industry, and corporate working together to ensure that AI technology benefits, rather than harms, people and communities. For nearly a decade, PAI has developed tools, resources and recommendations to educate and inform on what safe, trustworthy, and inclusive AI looks like.

» **Pol.is:** Polis is an open-source platform that helps cities, states, and countries to find common ground on complex issues. Launched in 2012, it's become part of the national democratic infrastructure in Taiwan, the UK, and Finland, collecting and analyzing viewpoints from thousands of participants to identify points of consensus.

» **Public Citizen:** Public Citizen is an advocacy organization that works to defend democracy, resist corporate power, and ensure that government works for the public interest. Public Citizen's work covers both political economy and the impact of AI on elections. They have a lobbying arm, policy expertise, and an in-house litigation team. They also co-chair a large AI coalition, the AI Big Tent Collective, with hundreds of public interest organizations involved. Over the last three years, they helped move **many state legislative pieces** limiting usage of deep fakes in elections.

» **RAND:** RAND is a research organization that provides leaders with the information they need to make evidence-based decisions. It studies the potential opportunities and risks of artificial intelligence, including ways to strengthen the AI workforce, countering AI-driven social media manipulation by U.S. adversaries, AI governance and how to make future AI systems safer, potential military applications of AI, how AI is being used in the classroom, and identifying occupations at risk of being replaced by AI.

» **RebootDemocracy.ai:** Rebooting Democracy in the Age of AI features original Insights and news of the week on AI, Governance and Democracy in the United States and globally. It is also the home of the **Observatory of Public Sector AI.**

» **Remesh:** Remesh is an AI platform for conversational research. The platform has been used to facilitate large-scale deliberation, covered in **Democratic Policy Development using Collective Dialogues and AI:**, a project which combined AI-enabled collective dialogues that make deliberation democratically viable at scale with bridging-based ranking for automated consensus discovery.

» **Stanford Law School AI Initiative:** The Stanford Law School AI Initiative brings together faculty, students, centers, labs, and external collaborators to inform how AI is developed, governed, and used responsibly in law and public life. They act as a hub to bring together the school's AI-driven centers, labs, and clinics. This includes the Legal Innovation through Frontier Technology Lab (liftlab), the Deborah L. Rhode Center on the Legal Profession, and the Regulation, Evaluation, and Governance Lab (RegLab).

» **Tech Policy Lab (University of Washington):** The Tech Policy Lab is an interdisciplinary collaboration at the University of Washington that aims to enhance technology policy through research, education, and thought leadership. Their work covers multiple areas, including introducing diverse voices into AI policy-

making, designing AI to be more responsive to cultural context, and early research into AI and bias. In the context of effective governance, the Lab has trained a generation of technologists and lawyers who work across industry, government, and academia. In the context of material infrastructure and the environment, the Lab is leading a project around the forces that hide AI's material and environmental impacts.

» **U.S. Digital Response:** USDR builds governments' digital capacity to enable service delivery that builds trust, and provides a safe space for state and local governments to responsibly adopt generative AI tools in content creation, task automation, translation, and analysis. Their technical assistance also includes support for hiring key AI roles into government, facilitating swift AI adoption and program management.

» **Vanderbilt Policy Accelerator:** The Vanderbilt Policy Accelerator for Political Economy and Regulation (VPA) focuses on a small number of topics in political economy and regulation. In its work on AI, they focus on networks, platforms, and utilities law, industrial policy, and public options that address the challenges and opportunities of AI.

Regrantors and Pooled Funds

Organizations providing funding to support the public interest orientation of AI implementation and regulation.

» **Berggruen Institute:** Since its founding, the Berggruen Institute has forged new conceptual frameworks to meet the challenges and harness the opportunities of the arriving future. As an independent think and action tank, they have the autonomy to reach beyond academic disciplines to bridge social divides, partisan dispositions, and cultural boundaries.

» **(Global Fund for a New Economy) (GFNE) Political Economy of AI Initiative:** GFNE's Political Economy of AI initiatives examines how artificial intelligence is reshaping global power, governance, and economic structures, with a focus on ensuring that AI development benefits the public interest rather than concentrating control in the hands of a few. Through policy research, coalition-building, and strategic convenings, GFNE works to advance democratic oversight, equitable access, and global cooperation in the governance of AI systems.

» **Humanity AI:** a $500 million five-year US-focused initiative dedicated to making sure people have a stake in the future of artificial intelligence (AI). As part of Humanity AI, ten foundations supporting the arts, labor and work, democracy, education, and security are driving new investments toward establishing a people-driven future where AI delivers for humanity, strengthens communities, and enhances human creativity.

» **Mozilla Foundation:** Mozilla Foundation aims to build a better technology future, working alongside developers, innovators, and advocates. They work to combat misinformation, AI bias, and platforms' impact on global elections. New initiatives include the Technology Fund, the Data Futures Lab, and an expansion of the responsible AI Fellowship.

» **Reset.Tech:** Reset Tech's mission is to guard against digital threats to security, safety, and fundamental rights. They seek to "reset" the connection between media and democracy to restore the promise of technology that works for people and free expression. They work to hold the biggest tech companies accountable to the values of democratic societies by promoting new ideas to change laws, industry standards, and consumer attitudes.

» **Current AI:** (see European Efforts section).

European Efforts

This list is incomplete because it is out of scope for this research.

» **Current AI:** Current AI is a global partnership to accelerate public-interest AI. Backed by over €400 million, they bring together governments, philanthropy, researchers, and industry to build shared AI infrastructure AI: open tools, trusted data, and people-first applications. They also provide implementation support, access to datasets, accountability frameworks, and compute.

» **European Artificial Intelligence & Society Fund:** The European AI & Society Fund is a pooled fund that mobilizes the public interest community in Europe to fight for artificial intelligence that serves people, society, and planet. They award grants, build the field, fundraise, and forge collective strategies.

» **Go Vocal:** Go Vocal's mission is to build stronger democracies by making public decision-making more inclusive, participatory, and responsive. They are a community engagement platform to help governments collect public input, streamline workflows, and implement AI.

» **Make.org:** Make.org is a neutral and independent civic tech organization whose mission is to engage people in the positive transformation of society. Their teams build democratic methods and infrastructure that enable citizen participation in public decision-making, increase democratic resilience, identify societal consensus, and help keep the public informed.