# Artificial Intelligence: The Real Risks, the Fear Campaign, and the Race America Cannot Afford to Lose

*An evidence-based examination of legitimate AI risks, organized narrative and policy campaigns, foundation-funded civil-society infrastructure, media amplification, the data-center second front, and the strategic consequences for American AI leadership*


## Abstract

Artificial intelligence is being presented to Americans through two radically different stories. In one, AI is a foundational technology with the potential to accelerate science, medicine, defense, education, productivity, and economic growth. In the other, AI is portrayed primarily as a threat: a job destroyer, surveillance system, environmental disaster, misinformation engine, monopoly machine, danger to children, and potentially an existential risk to humanity.

Some of those dangers are real. AI can be wrong. It can be abused. It can displace particular kinds of work. It can amplify fraud, deepfakes, cyberattacks, surveillance, and manipulation. Powerful companies can use it badly, and governments can use it badly. Serious technology therefore requires serious safeguards.

But the public information environment is not simply a spontaneous collection of independent warnings. A documented advocacy and philanthropic infrastructure is actively shaping the AI debate. In July 2025, more than 90 technology, labor, environmental, consumer, civil-rights, and activist organizations launched the People's AI Action Plan specifically as a counter-weight to the Trump administration's AI strategy. Its signatories included Food & Water Watch, Citizens Action Coalition, and MediaJustice. Those same three organizations later appeared on the coordinating committee of the 500-plus-organization Stop Data Centers Coalition. AI Now Institute, one of the organizers of the People's AI campaign, separately built a policy toolkit expressly designed to stop, slow, and restrict AI data-center construction. [5-8]

The funding infrastructure is equally visible. In 2023 Open Society Foundations joined nine other philanthropies in a coordinated initiative representing more than $200 million in AI-related funding. The foundations committed to align investments, scale civil-society efforts, support AI-governance advocacy, and advance international AI rules and norms. Open Society has separately funded AI Now regulatory research. By 2026, the overlapping Humanity AI philanthropic network said its participating funders would drive more than $500 million in new investments over five years, while directing grants to AI policy groups, public-conversation projects, and journalism. [11-16]

The strongest documentary conclusion is therefore not that every critic of AI is dishonest, nor that every foundation or NGO shares a secret command structure. It is that a large, coordinated, well-funded ecosystem exists to shape public attitudes, political rules, media coverage, and infrastructure decisions around AI. The same ecosystem increasingly treats the technology and the physical infrastructure required to run it as parts of one political fight.

That matters because the United States government now explicitly describes AI as a global race in which leadership will determine economic, security, technological, and international-standard-setting advantages. The question is not whether AI should be unregulated. The question is whether Americans are being given a proportionate picture of its risks and benefits—or whether legitimate dangers are being bundled, amplified, and organized into a fear environment that can slow the very technology the country is simultaneously trying to lead. [1]


## 1. Begin With What Is True

A serious investigation should begin by conceding the strongest legitimate concerns rather than dismissing them.

Artificial intelligence can produce false answers. It can generate convincing fabricated images, audio, and video. It can be used for fraud, manipulation, automated surveillance, cyber operations, discriminatory decision systems, or irresponsible medical and financial advice. Some occupations and tasks will be disrupted. Large AI companies can acquire enormous economic and political power. AI systems also require large amounts of computing infrastructure and electricity.

Those are real governance problems. They justify evaluation, accountability, cybersecurity, privacy protection, transparency in high-stakes uses, and rules against clearly harmful conduct.

But 'AI has risks' is a very different proposition from 'AI is fundamentally dangerous,' 'AI is being done to you,' 'AI inevitably destroys jobs,' 'AI infrastructure is extractive,' or 'the only responsible response is to slow, stop, ban, or politically contain the technology.' The distance between those propositions is where narrative manipulation can occur.

The same distinction drove the earlier data-center investigation: a real problem can be converted into a much larger political conclusion by changing 'can happen' into 'will happen,' 'some systems' into 'AI,' and 'requires safeguards' into 'must be stopped.'

## 2. The AI Fear Environment Is Real—and Growing

Public concern about AI has increased sharply. Pew Research Center reported in August 2026 that 52 percent of Americans were more concerned than excited about the increased use of AI in daily life, up from 37 percent in 2021. Earlier Pew work found a striking gap between the public and people with AI expertise: 56 percent of surveyed experts expected AI to have a positive impact on the United States over the next twenty years, compared with only 17 percent of the general public. [2-3]

That gap does not prove propaganda. It does show that the public is encountering AI through an unusually fear-heavy information environment compared with the people most familiar with the field.

The fear campaign also has multiple strands. One is the existential-risk movement. In March 2023 the Future of Life Institute published an open letter calling for at least a six-month pause on training systems more powerful than GPT-4 and urged governments to impose a moratorium if labs would not pause voluntarily. The letter eventually drew tens of thousands of signatures and generated worldwide headlines. [4]

A second strand is the civil-society and anti-corporate AI movement. It focuses less on extinction and more on jobs, monopoly, surveillance, discrimination, environmental justice, immigration, policing, children, labor, and corporate power. That is the strand most directly connected to the foundation and NGO network examined in the companion papers.

A third strand is now physical infrastructure: data centers, electric generation, transmission, water, emissions, and land use. The key finding of this investigation is that strands two and three are no longer separate. The same institutions increasingly treat them as one campaign.

## 3. Something Bigger Than Ordinary Safety Debate Is Happening

In July 2025 a broad coalition of more than 90 organizations launched the People's AI Action Plan specifically to provide a 'counter-weight' to the Trump administration's AI executive orders and AI agenda. The coalition described the administration's strategy as a giveaway to Big Tech, said AI was already harming workers, consumers, and communities, and argued that some AI technologies should never be used at all. [5-6]

The coalition was not limited to technology-policy specialists. It deliberately combined organizations from labor, climate, environmental justice, civil rights, consumer advocacy, media, immigration, education, and anti-monopoly activism. Its signatories included AI Now Institute, Data & Society, Public Citizen, Food & Water Watch, Friends of the Earth, Climate Justice Alliance, Center for Biological Diversity, MediaJustice, Citizens Action Coalition, Color of Change, SEIU, National Nurses United, Open Markets Institute, and an organization literally named Stop Gen AI. [5-6]

This matters because issue bundling is one of the central influence techniques identified in the PsyOp framework. A person may have one legitimate concern—jobs, electricity bills, privacy, climate, children, copyright, misinformation, surveillance—and then encounter a coalition in which all of those concerns are attached to the same target. The psychological effect is larger than any single technical dispute.

The movement's own language makes its strategic purpose clear: it is not merely producing neutral research about AI. It is organizing a political counter-program to the federal AI expansion agenda.

## 4. The Two-Prong Campaign: Attack the Technology and Attack the Infrastructure

The earlier data-center article examined a national coalition of more than 500 NGOs and allied organizations seeking a nationwide halt to new data-center construction. At first glance that may appear to be an environmental campaign distinct from the broader political fight over AI.

The organizational record shows otherwise.

Food & Water Watch, Citizens Action Coalition, and MediaJustice are signatories to the People's AI Action Plan. Food & Water Watch later identified those same organizations—along with several others—as members of the coordinating committee of the Stop Data Centers Coalition. [5, 7]

AI Now provides an even more direct bridge. AI Now helped organize the People's AI Action Plan and states that its policy work is intended to redirect the current trajectory of AI. It also created the North Star Data Center Policy Toolkit, which says it is 'primarily geared toward stopping, slowing, and restricting rampant data center development in the US' and provides organizers and policymakers with bans, moratoria, zoning restrictions, energy rules, water rules, tax policies, and other interventions. [6, 8-9]

That means the two fronts can no longer reasonably be described as unrelated:

- Front One — AI itself: frame rapid AI deployment as a threat to workers, democracy, equality, privacy, safety, children, communities, and public institutions.
- Front Two — AI infrastructure: frame the data centers and energy systems required for AI as threats to water, electricity bills, climate, health, land, and local control.

One front reduces public trust in the technology. The other raises the physical and political cost of building the compute infrastructure needed to develop and deploy it.

No hidden memo is necessary to establish the relationship. The same organizations openly participate in both campaigns, and one of the principal AI-policy organizations has built tools for both the narrative-policy fight and the infrastructure fight.

## 5. How a Legitimate AI Concern Becomes an Influence Campaign

Our earlier PsyOp framework identified a recurring pattern: successful manipulation rarely begins with a completely fabricated claim. It begins with something true, emotionally enlarges it, omits countervailing information, attaches other fears to it, creates an enemy, repeats the message across many apparently independent voices, and converts the resulting emotion into political action.

**Fear amplification:** Worst-case outcomes—mass unemployment, autonomous warfare, children harmed by chatbots, surveillance, democratic collapse, environmental destruction, or human extinction—receive more attention than ordinary beneficial uses.

**Issue bundling:** Jobs, climate, race, surveillance, monopoly, immigration, policing, copyright, children, disinformation, water, energy, and national security are attached to one target: AI.

**Enemy construction:** The debate is repeatedly framed as ordinary people versus “Big Tech,” “tech oligarchs,” “tech bros,” billionaires, or corporations using AI “on us.”

**Omission:** Benefits such as productivity, scientific discovery, medical research, accessibility, cybersecurity, industrial capacity, defense, education, and strategic competition receive less emphasis in activist messaging.

**Coordinated repetition:** Large coalitions distribute overlapping claims across labor, climate, media, civil-rights, consumer, academic, and community organizations.

**Policy conversion:** Concern is translated into moratoria, bans, permitting delays, procurement restrictions, state-by-state regulatory regimes, licensing proposals, reporting mandates, and pressure campaigns.

**Narrative strategy:** AI Now itself describes “discourse-shaping” work, “advanc[ing] narratives,” capitalizing on policy windows, and catalyzing action with a broad coalition of allies. [9]

## 6. A PsyOp-Pattern Analysis of the Major Claims

| Anti-AI Narrative | What Is Actually True | PsyOp / Manipulation Risk |
|---|---|---|
| "AI will take everyone's jobs." | AI will automate some tasks and displace some roles, while creating new tasks, industries, and productivity gains. Effects vary sharply by occupation, adoption pattern, and policy. | High when "some displacement" becomes inevitable mass unemployment. |
| "AI is being used on us, not by us." | Some high-stakes systems are deployed without meaningful consent. But millions of people voluntarily use AI as a tool for writing, research, coding, accessibility, translation, analysis, and work. | High when coercive-use examples are generalized to the technology as a whole. |
| "AI is unsafe and untested." | Some systems are released with unresolved risks. Others are evaluated extensively, and ordinary software also evolves after deployment. Risk depends heavily on use case. | Moderate-High when all systems are treated as equivalent. |
| "AI threatens democracy." | Deepfakes, automated persuasion, and false information can affect politics. AI can also detect fraud, improve access to information, translate material, and assist analysis. | Moderate when only the threat side is presented. |
| "AI is destroying the environment." | AI requires substantial computing, electricity, and infrastructure. Its footprint varies with hardware, model, energy source, cooling, and utilization. AI can also improve grid management, scientific modeling, and efficiency. | High when infrastructure costs are presented without mitigation or benefits. |
| "Big Tech controls AI, so rapid development must be stopped." | Compute and frontier models are concentrated. Competition policy and open models can address concentration. Slowing the entire U.S. ecosystem can also entrench incumbents or shift leadership abroad. | High when anti-monopoly concern becomes an argument against national capacity. |
| "AI will make people unable to think." | Overreliance can weaken some skills. Tools can also extend human capability when used critically, as calculators, search engines, software, and scientific instruments have done. | Moderate-High when a behavioral risk becomes a technological inevitability. |
| "AI data centers harm communities, so expansion should pause." | Poorly structured projects can impose costs. Developer-funded power, grid upgrades, water reuse, ratepayer protections, and better siting can change the outcome. | Very High when regulation and mitigation are displaced by blanket moratoria. |
| "The AI race is just a Big Tech sales pitch." | Companies benefit from expansion, but the U.S. government independently treats AI leadership as an economic and national-security competition with global standard-setting consequences. | High when commercial self-interest is used to erase strategic reality. |

## 7. The Organizational Bridge: The Same NGOs Appear on Both Fronts

| Organization | General AI campaign | Data-center campaign | Documentary significance |
|---|---|---|---|

| | | | |
|---|---|---|---|
| AI Now Institute | Helped launch People's AI Action Plan; advocates narrative and policy strategies to redirect AI trajectory. | Built North Star toolkit to stop, slow, and restrict data centers; hosts organizer trainings. | Direct bridge between AI-policy opposition and infrastructure opposition. |
| Food & Water Watch | Signatory to People's AI Action Plan. | Lead organizer of Stop Data Centers Coalition and national moratorium campaign. | Same national NGO operates on both fronts. |
| Citizens Action Coalition | Signatory to People's AI Action Plan. | Member of Stop Data Centers coordinating committee; organizes state moratorium efforts. | Exact organizational overlap in campaign infrastructure. |
| MediaJustice | Signatory to People's AI Action Plan. | Member of Stop Data Centers coordinating committee. | Media/communications organization sits inside both coalitions. |
| Center for Biological Diversity | People's AI Action Plan signatory. | Organizes against fast-tracking data centers and AI infrastructure. | Links broader AI governance to environmental/permitting front. |
| Friends of the Earth / climate groups | People's AI Action Plan includes Friends of the Earth and multiple climate organizations. | Climate organizations participate in national data-center opposition. | Environmental narrative supplies a second route to slowing AI buildout. |

The important point is not that all organizations have identical motives. They do not need identical motives to create a common political effect. Coalition architecture allows different constituencies to approach the same target through different emotional and policy pathways.

## 8. Follow the Money: Philanthropy Builds the AI-Governance Ecosystem

The anti-corporate AI-policy ecosystem is not operating on volunteer energy alone. Major foundations have openly organized and financed it.

In November 2023, Open Society Foundations joined the David and Lucile Packard Foundation, Democracy Fund, Ford Foundation, Heising-Simons Foundation, MacArthur Foundation, Kapor Foundation, Mozilla Foundation, Omidyar Network, and Wallace Global Fund in a coordinated AI initiative. The participating foundations said they collectively had more than $200 million invested in the field and would coordinate future action, align grantmaking, scale civil-society efforts, build policymaker understanding, support advocacy, and help develop international AI rules and norms. [11-12]

That is not a theory about philanthropic influence. It is the stated purpose of the initiative.

By 2025-2026, an overlapping coalition of foundations launched Humanity AI, a philanthropic initiative intended to drive more than $500 million in new investment over five years. MacArthur described the initiative as designed not only to fund projects but to 'shape public conversations about AI' so that the public and policymakers understand that Silicon Valley's vision is not inevitable. [15-16]

The May 2026 Humanity AI grant announcement illustrates how the money is distributed across the influence system: policy organizations, civil-rights groups, research institutes, journalism organizations, grassroots networks, and public-conversation projects. AI Now received $500,000 through the pooled fund; the Pulitzer Center received $500,000; Data & Society was selected to lead a $3 million AI Civics project designed to bring communities into the governance conversation. [15]

## 9. Soros / Open Society: Direct Funding, Coordination, and International Rules

George Soros's Open Society Foundations is directly inside this AI-policy funding architecture.

Open Society was one of the ten philanthropies in the $200 million coordinated initiative. The group explicitly committed to aligning investments and scaling civil-society efforts around AI governance, worker policy, discrimination and bias, democracy, transparency, accountability, and international rules and norms. [11-12]

Open Society has also funded AI Now directly at the project level. AI Now's 2024 'Lessons from the FDA for AI' project states that it was supported by a grant from Open Society Foundations. The project examines FDA-style regulatory gates for AI and argues that meaningful AI regulation should go far beyond ordinary penalties and create stronger monitoring, documentation, and compliance mechanisms. [13]

Open Society also operates its own AI-and-information work. Its AI in Journalism Futures project convened participants to examine how AI may reshape journalism and civic information over a 5-to-15-year horizon. In 2025 Open Society leadership called the 'unchecked rise of artificial intelligence' an existential-scale human-rights challenge and called for stronger digital rights and transparent AI governance. [24, 29]

This places Open Society in three relevant layers at once: funding AI-governance organizations, coordinating philanthropies and civil society, and shaping the information-policy conversation about AI.

## 10. AI Now: A Narrative-and-Policy Engine at the Center of Both Fronts

AI Now Institute deserves special attention because its own institutional description closely matches the influence-operation mechanics examined in this series.

AI Now states that it develops policy strategies to redirect away from the current AI trajectory. It says its work has set the bar for 'discourse-shaping' analysis and that its approach is to identify and capitalize on actionable policy windows, advance narratives that chart a long-term strategy, and catalyze energy toward action in close partnership with a broad coalition of allies. [9]

Those are not accusations made by a critic. They are AI Now's description of its method.

Its 2025 landscape report makes the strategic frame still more explicit. The report tells advocates to make AI a fight about power rather than progress and offers a roadmap that speaks directly to advocacy campaigns and narrative strategy. Elsewhere the report says today's AI is not merely being used by the public but 'used on us' and argues that AI consistently worsens material conditions and concentrates power. [10, 25-26]

AI Now then operates in the data-center arena as well. Its North Star toolkit is built for organizers and policymakers seeking to stop, slow, and restrict data-center development. Its policy menu includes moratoria, bans, zoning, energy restrictions, water restrictions, tax rules, procurement conditions, and federal interventions. [8]

Financially, AI Now is supported by major foundations. MacArthur awarded AI Now $2 million in general operating support in 2025 and another $2 million aligned with Humanity AI in 2026 for national-security and AI work. Ford identifies AI Now as a grantee, and Open Society funded its FDA-style regulation project. [13-17]

The result is a clean funding-to-research-to-narrative-to-coalition-to-policy pipeline.

## 11. The $500 Million Humanity AI Expansion

Humanity AI demonstrates how rapidly the philanthropic AI-policy field is scaling. MacArthur says the coalition will drive more than $500 million in new investment over five years. Its founding partners include Ford, MacArthur, Mozilla, Omidyar Network, Packard, Kapor, Mellon, Lumina, Doris Duke, and Siegel Family Endowment. [15-16]

The coalition's architecture is broad enough to shape nearly every layer of the public debate:

- Policy: grants to AI Now, Center for Democracy and Technology, and other governance organizations.
- Grassroots and community framing: grants to organizations designed to build public participation and community power around AI.
- Research: support for institutes studying labor, national security, democracy, and societal impacts.
- Media: grants to the Pulitzer Center and other journalism programs.
- Public conversation: Data & Society's AI Civics project and foundation-backed community discussions.
- International framing: support for global forums and rules-oriented work.

Again, none of these activities is inherently improper. The analytical point is concentration. When the same philanthropic network finances the research, the policy advocates, the civic-engagement infrastructure, and the journalism training around one technology, the resulting public conversation is less independent than it may appear from the outside.

## 12. The Media and Storytelling Layer

The AI debate has a funded communications infrastructure in addition to the NGO and policy infrastructure.

The Pulitzer Center's AI Accountability Network says it is funded with support from Open Society Foundations, MacArthur Foundation, Humanity AI, Ford Foundation, Omidyar Network, Park Foundation, and others. Its fellowship program finances journalists, mentorship, training, and—beginning with the 2026-2027 cohort—impact plans designed to help investigations reach strategic audiences through creators, libraries, and multistakeholder discussions. Its AI Spotlight program explicitly trains reporters and editors to cover and shape coverage of AI. [20-21]

Humanity AI separately awarded the Pulitzer Center $500,000 in 2026. MacArthur also awarded $1 million to scale the Center's AI Accountability Network. [15-16]

The Associated Press publicly discloses that Omidyar Network provides financial assistance to support AP coverage of artificial intelligence and its impact on society. AP states that editorial responsibility remains with AP. The funding relationship itself is public and belongs in any map of who finances AI information production. [18-19]

Open Society's AI in Journalism Futures project adds another layer by directly convening people around long-range scenarios for how AI may reshape journalism and civic information. [29]

None of this means a particular news story was ordered by a foundation. It means that foundations involved in AI governance are also financing parts of the journalism ecosystem that decides which AI harms, accountability questions, and policy conflicts receive sustained attention. That is a material influence pathway and should be disclosed as such.

## 13. Clinton / CGI: The Convening Layer, Not the Operational Core

The Clinton network appears in this AI ecosystem primarily as a convening and agenda-setting node rather than as the operational organizer of the People's AI Action Plan.

The Clinton Global Initiative's 2025 Annual Meeting included an 'AI in Action: Closing Critical Gaps' program while featuring leaders from major philanthropies including the Ford Foundation. CGI's 2026 programming likewise announced strategic discussions on democracy, media literacy, and AI's role in public life. [22-23]

That fits the larger network pattern documented in the companion investigation: CGI brings government, philanthropy, business, NGOs, media, and policy leaders into recurring contact. In the AI case, the strongest direct funding and organizing links run through Open Society, Ford, MacArthur, Omidyar, AI Now, Public Citizen, Food & Water Watch, and the coalition organizations described above. CGI is best understood here as part of the recurring cross-sector convening architecture surrounding the issue, not as the principal anti-AI campaign operator.

That distinction strengthens rather than weakens the network analysis: each actor should be placed in the role the documents actually support.

## 14. What Is America Actually Trying to Win?

The counter-narrative becomes easier to evaluate once the strategic objective is stated plainly.

The official U.S. AI strategy says the United States is in a race for global AI dominance and that the country with the largest AI ecosystem will set global standards and receive broad economic and security benefits. The strategy has three pillars: accelerate innovation, build AI infrastructure, and lead international diplomacy and security. [1]

That framing is not merely a slogan about selling chatbots. The federal program now connects AI to scientific discovery, national security, cybersecurity, defense, education, energy infrastructure, export policy, and the ability to establish the technical stack used around the world. [1]

The strategic problem is therefore simple. Frontier AI requires models, chips, talent, power, data centers, capital, research, and rapid deployment. A policy movement that increases friction against the models while a second movement increases friction against the data centers attacks two of the inputs required for national-scale AI leadership.

We do not need to prove that every activist wants China to win in order to evaluate the consequence. If the United States slows model development, fragments regulation across states, pauses data-center construction, delays energy and transmission projects, and raises the cost of deployment while competitors continue building, the relative strategic effect is obvious.

That is why the general AI article and the data-center article belong together. The infrastructure battle is not a separate controversy. It is a second front in the same national question: will the United States build enough technological capacity to remain the leading AI power?

## 15. What Is Left Out of the Fear Narrative?

**Narrative: "AI will eliminate jobs."**
Missing context: AI changes tasks as well as jobs, raises productivity in some occupations, creates new roles, and can make smaller firms and individuals more capable. The relevant question is which tasks change, how quickly, and how workers share in productivity gains.

**Narrative: "AI is a threat to human thinking."**
Missing context: AI can also function as a tutor, translator, coding assistant, research assistant, accessibility tool, and analytical amplifier. Outcomes depend heavily on how it is used.

**Narrative: "AI threatens democracy through misinformation."**
Missing context: AI can also help authenticate media, translate civic information, analyze large public records, identify anomalies, and lower the cost of research. The danger is misuse, not intelligence itself.

**Narrative: "AI is an environmental disaster."**
Missing context: the environmental footprint depends on energy source, chip efficiency, utilization, cooling design, grid investment, and developer-funded infrastructure. AI can also optimize energy systems and accelerate scientific work.

**Narrative: "AI dominance is just a tech-billionaire talking point."**
Missing context: U.S. national-security and technology policy independently treats AI leadership as a geopolitical competition with economic, military, cybersecurity, and standards-setting consequences.

**Narrative: "Environmental groups are independently reacting to data centers."**
Missing context: more than 500 NGOs and allied organizations are participating in a coordinated national coalition, and several of its core organizations also participate in the broader People's AI Action Plan.

**Narrative: "Civil society is merely responding to industry."**
Missing context: major foundations have committed hundreds of millions of dollars to coordinate civil-society AI governance, policy advocacy, public conversations, journalism, and international rules.

The missing information does not make every concern false. It prevents the audience from being forced into a conclusion by selective framing.

## 16. A Better Test for Every AI Claim

| Question | What a Good Analysis Should Show |
|---|---|
| What is the exact AI system? | Model, version, use case, user, and decision context—not "AI" in the abstract. |
| What harm was actually measured? | Observed outcome separated from prediction, anecdote, or worst-case scenario. |
| How common is the harm? | Rate, denominator, baseline, and comparison with non-AI alternatives. |
| What benefits are measured? | Productivity, quality, cost, access, scientific value, security, or other gains. |
| What mitigation is available? | Evaluation, human review, authentication, privacy controls, procurement rules, insurance, liability, technical safeguards. |

| | |
|---|---|
| Who is funding the message? | Foundation grants, corporate funds, government grants, memberships, media funding, or coalition support. |
| Is the source trying to improve AI or stop it? | Regulation, targeted prohibition, licensing, moratorium, infrastructure halt, or industry-wide opposition. |
| Is the claim falsifiable? | What evidence would cause the source to revise its conclusion? |
| What happens if the U.S. slows? | Effects on competition, open-source development, security, standards, investment, talent, and foreign leadership. |
| Are technology and infrastructure being bundled? | Whether concerns about one data center are being converted into an argument against AI generally—or vice versa. |

## 17. The Larger Pattern: Fear Can Regulate a Technology Before People Understand It

The emerging AI influence system follows the same sequence seen in other narrative campaigns examined in this series:

1. Identify a technology with genuine risks and enormous social consequences.
2. Select the most frightening examples and worst-case scenarios.
3. Bundle different constituencies around different fears: jobs, children, democracy, surveillance, monopoly, climate, water, inequality, war, and extinction.
4. Finance research and policy organizations that define the problem and propose governance responses.
5. Fund public engagement, narrative work, journalism, training, coalition building, and strategic communications.
6. Create large cross-sector coalitions so the same conclusion arrives from labor, climate, civil rights, consumer advocacy, academia, media, and community organizations.
7. Translate concern into regulations, moratoria, bans, procurement restrictions, and infrastructure delays.
8. Use the resulting public concern and policy activity as proof that the technology is socially dangerous and requires still more intervention.
9. Extend the campaign from the software to the infrastructure that makes the software possible.
10. Repeat with larger budgets and broader coalitions as the technology becomes more strategically important.

This is not an argument for blind trust in AI companies. Corporate public relations can manipulate, exaggerate, and conceal just as advocacy campaigns can. The appropriate response is not to choose a propaganda team. It is to examine the evidence, funding, incentives, omissions, and requested behavior on both sides.

## Conclusion: From Fear to Discernment

Artificial intelligence has real risks. It also has enormous potential.

That combination makes it unusually vulnerable to political manipulation. A technology capable of affecting work, information, science, defense, medicine, education, finance, and national power supplies almost unlimited material for fear-based messaging. Every constituency can be shown a different reason to be afraid.

The documentary record now shows that these fears are not being amplified only by isolated commentators. A broad civil-society coalition was explicitly created to counter the federal AI expansion agenda. Several organizations in that coalition later became core organizers of the national campaign to halt AI data-center construction. AI Now operates in both arenas: it helped organize the broader AI counter-campaign, describes its own work as discourse-shaping and narrative strategy, and built a policy toolkit specifically to stop, slow, and restrict data centers.

Behind that advocacy layer sits substantial philanthropic capital. Open Society Foundations joined nine other major philanthropies in a coordinated $200 million AI initiative designed to align grants, scale civil-society work, influence governance, and advance international rules and norms. Open Society separately funded AI Now

research. An overlapping Humanity AI coalition is now driving more than $500 million in new investment into policy, civic engagement, research, journalism, and public-conversation infrastructure. The Pulitzer Center's AI Accountability Network is financed by many of the same foundations, while Omidyar Network also funds AI coverage at the Associated Press.

These facts do not require a fantasy of one person issuing orders to every organization. Modern influence systems work through aligned funding, shared frameworks, recurring coalitions, professional networks, convening institutions, communications channels, and mutually reinforcing incentives.

The practical effect is what matters. America is explicitly trying to lead the world in AI. The government says that leadership will determine economic strength, national security, technological standards, scientific capacity, and international influence. At the same time, a well-funded domestic advocacy system is pressuring policymakers to slow or redirect AI development while an overlapping coalition presses to halt the infrastructure needed to run it.

That does not mean every regulation is bad or every data center is good. It means Americans should recognize the full battlefield before deciding what they believe.

Ask who supplied the evidence. Ask who funded the organization. Ask which facts were omitted. Ask whether a specific harm has been generalized to an entire technology. Ask whether the proposed solution fixes the problem or simply stops development. Ask whether the same organizations are appearing across seemingly separate campaigns. And ask what happens strategically if the United States adopts the brakes while its competitors keep building.

The goal should not be uncritical acceleration.

It should be disciplined acceleration: build powerful AI, test it hard, punish actual abuse, protect legitimate rights, make developers pay the real infrastructure costs they create, keep markets competitive, strengthen the grid, educate the public, and preserve America's ability to lead.

Do not let Big Tech sell you a fantasy. But do not let a coordinated fear campaign frighten you into surrendering a foundational technology either.

The data-center fight is the second front. The larger fight is over whether Americans will understand AI well enough to govern it without being manipulated into rejecting it.

## Appendix A — Primary Sources and Documentary Records


**[1] AI.gov — President Trump's AI Strategy and Action Plan.** Official statement that the United States is in a race for AI dominance; three pillars are innovation, infrastructure, and international leadership. Source

**[2] Pew Research Center — How the U.S. Public and AI Experts View Artificial Intelligence.** Documents the large optimism gap between AI experts and the general public. Source

**[3] Pew Research Center — Young U.S. Adults Are Increasingly Wary of AI.** Reports 52% of Americans more concerned than excited in 2026, up from 37% in 2021. Source

**[4] Future of Life Institute — Pause Giant AI Experiments: An Open Letter.** Called for a six-month pause on systems more powerful than GPT-4 and government moratoria if labs did not pause. Source

**[5] People's AI Action Plan.** Coalition statement and signatory list; includes AI Now, Food & Water Watch, Citizens Action Coalition, MediaJustice, Public Citizen, and many others. Source

**[6] AI Now Institute — People's AI Action Plan Launch.** Describes 90+ organization campaign as a counter-weight to Trump AI policy. Source

**[7] Food & Water Watch — The New Coalition Taking on Big Tech's AI Data Centers.** Documents 500+ organization data-center coalition and its coordinating committee. Source

**[8] AI Now Institute — North Star Data Center Policy Toolkit.** States toolkit is geared toward stopping, slowing, and restricting data-center development. Source

**[9] AI Now Institute — About Us.** Describes discourse-shaping work, narrative strategy, policy windows, coalition action, and foundation funding. Source

**[10] AI Now Institute — Artificial Power: 2025 Landscape Report.** Frames current AI trajectory as a power struggle and proposes action strategy. Source

**[11] Open Society Foundations — Open Society and Other Funders Launch New AI Initiative.** Documents $200M coordinated philanthropic AI initiative and goals. Source

**[12] Ford Foundation — Philanthropies Launch New Initiative to Ensure AI Advances the Public Interest.** Lists ten participating foundations and commitments to align investments, scale civil society, and develop international rules and norms. Source

**[13] AI Now Institute — Lessons from the FDA for AI.** AI Now states this regulatory project was supported by a grant from Open Society Foundations. Source

**[14] MacArthur Foundation — AI Now Institute Grant.** $2 million general operating grant in 2025. Source

**[15] Ford Foundation — Humanity AI Announces More Than $18 Million in New Grants.** Lists grants to AI Now, Pulitzer Center, Data & Society-led AI Civics, and other organizations; identifies founding funders. Source

**[16] MacArthur Foundation — $10 Million to Advance AI By and For People.** Says Humanity AI will drive more than $500M and explicitly aims to shape public conversations about AI. Source

**[17] Ford Foundation — AI Now Institute Grantee Profile.** Identifies AI Now as a Ford grantee and calls for greater regulatory friction and campaigning. Source

**[18] Associated Press — AI Training Program and Coverage Expansion With Omidyar Network Support.** AP announcement of Omidyar-supported AI coverage and training. Source

**[19] Associated Press — Supporting AP.** Lists foundation support, including Omidyar Network, and explains philanthropy-funded coverage model. Source

**[20] Pulitzer Center — AI Accountability Fellowships 2026-2027.** Lists Open Society, MacArthur, Humanity AI, Ford, Omidyar, Park and others as funders; includes impact-plan and strategic-audience components. Source

**[21] Pulitzer Center — Reflections on the Inaugural AI Spotlight Series.** Documents MacArthur/Ford-supported training of more than 1,600 journalists and editors from 122 countries. Source

**[22] Clinton Foundation — Clinton Global Initiative 2025 Annual Meeting.** Includes AI in Action programming and major philanthropic participants including Ford Foundation leadership. Source

**[23] Clinton Foundation — CGI Returns to New York City, September 2026.** Announces discussions on democracy, media literacy, and AI's role in public life. Source

**[24] Open Society Foundations — The Future of Human Rights in a Changing World Order.** Calls unchecked AI a major human-rights challenge and advocates stronger digital governance. Source

**[25] AI Now Institute — A Roadmap for Action: Make AI a Fight About Power, Not Progress.** Explicitly discusses advocacy campaigns and narrative strategy. Source

**[26] AI Now Institute — Consulting the Record: AI Consistently Fails the Public.** Illustrates the organization's strong negative framing of present AI deployment. Source

**[27] AI Now Institute — North Star Toolkit: Federal Moratorium.** Recommends federal, state, and local moratorium pathways. Source

**[28] Pew Research Center — What Americans Think About the Global AI Race.** Documents public views on the importance of U.S. AI leadership. Source

**[29] Open Society Foundations — AI in Journalism Futures 2024.** Open Society convening and scenario-planning project on AI's impact on journalism and civic information. Source

## Appendix B — Network Crosswalk: General AI Campaign vs. Data-Center Campaign

| Node | General AI role | Data-center / infrastructure role | Funding / network connection |
|---|---|---|---|
| Open Society Foundations | Coordinated $200M philanthropic AI initiative; funds governance, rights, accountability, international norms; funded AI Now project. | No direct role in Stop Data Centers coordinating committee identified in reviewed records. | Member of ten-foundation AI funding initiative; funder of Pulitzer AI Accountability Network. |
| Ford Foundation | Major funder of AI Now and Humanity AI; funds technology-and-society policy work. | Humanity AI includes climate impacts among AI accountability concerns; AI Now grantee operates data-center toolkit. | Founding Humanity AI funder; CGI participant/convening overlap. |
| MacArthur Foundation | Millions in grants to AI Now, Data & Society, Pulitzer and other AI-governance groups. | Funds AI Now, which runs anti-data-center policy work. | Founding Humanity AI funder and major journalism funder. |
| Omidyar Network | Founding philanthropic AI initiative member; Humanity AI founding partner; tech-regulation advocacy. | Publicly interested in responsible data-center models; funds groups active in AI governance. | Supports AP AI coverage and Pulitzer AI Accountability Network. |
| AI Now Institute | People's AI Action Plan organizer; policy and narrative strategy; foundation funded. | North Star toolkit to stop/slow/restrict data centers; organizer trainings and moratorium proposals. | OSF project funding; MacArthur and Ford support; Humanity AI grantee. |
| Food & Water Watch | People's AI Action Plan signatory. | Lead Stop Data Centers Coalition organizer; national moratorium campaign. | Operates within 90+ AI coalition and 500+ data-center coalition. |
| Citizens Action Coalition | People's AI Action Plan signatory. | Stop Data Centers coordinating committee; state moratorium activity. | Exact organizational overlap. |
| MediaJustice | People's AI Action Plan signatory. | Stop Data Centers coordinating committee. | Exact organizational overlap; information/communications function. |
| Pulitzer Center | AI Accountability Network and AI journalism training. | Covers AI supply chain and data-center impacts through funded journalism. | Funded by OSF, MacArthur, Humanity AI, Ford, Omidyar and others. |

| | | | |
|---|---|---|---|
| Clinton Global Initiative | Cross-sector AI discussions; agenda-setting and convening. | No operational anti-data-center role established in reviewed records. | Recurring convening node connecting philanthropy, policy, business and civil society. |