How does the shift from human-crafted narratives to AI-generated, high-fidelity propaganda alter the structural accountability of democratic information ecosystems and reshape the power dynamics between state actors and civic institutions?

June 15, 2026 | SnugLab Research | readme.snuglab.com

Executive Summary

The shift from human-crafted narratives to AI-generated, high-fidelity propaganda significantly alters the structural accountability of democratic information ecosystems and reshapes power dynamics, primarily by granting state actors a compounding advantage in scale, speed, and precision. Evidence suggests that while AI democratizes content creation and enables civic counter-strategies, the sheer volume of AI-generated content, coupled with state actors' superior computational resources and data-driven targeting, overwhelms civic verification capacities and erodes public trust faster than it disrupts state coordination. This creates a structural asymmetry that fragments information authority while concentrating power in the hands of those who control the algorithms and data.

Key Findings

AI-Generated Propaganda Overwhelms Civic Verification and Erodes Trust

The exponential increase in AI propaganda volume creates a structural asymmetry that overwhelms the finite verification capacities of civic institutions, leading to a degradation of institutional trust. AI tools enable the generation of vast quantities of disinformation at virtually no cost [2, 7]. For instance, deepfake videos surged from approximately 500,000 in 2023 to an estimated 8 million by 2025, a 16-fold increase [7]. State-affiliated outlets have similarly scaled production, with one Russian-backed site increasing its daily article output by 2.4 times and nearly doubling its topic diversity [1]. This volume "pollutes" the digital environment, overwhelming fact-checkers and journalists and exacerbating information overload [7, 10, 12].

This volume overload degrades civic trust faster than it disrupts state coordination because it triggers "truth fatigue" and a "liar's dividend" [7]. As citizens and experts struggle to distinguish AI-generated falsehoods from genuine information, cynicism, news avoidance, and social disengagement increase [7]. When experts and news agencies fall for AI hallucinations, public trust in institutions like journalism and electoral processes is further damaged [3]. Simultaneously, the "liar's dividend" allows malicious actors to dismiss authentic evidence as AI-generated fabrications, evading accountability [4, 8, 11, 12]. State actors, conversely, leverage this volume to streamline operations, using AI agents that autonomously coordinate, write posts, and echo content to manufacture consensus [5]. A Russian-backed network, for example, produced over one million German-language posts at a rate of about two messages per second to flood social media with pro-regime narratives [9, 12].

State Actors Gain Compounding Computational and Data Advantages

While AI democratizes high-fidelity content generation, a compounding computational and data gap persists, amplifying state resource advantages and entrenching state power. AI tools are inexpensive and accessible, allowing grassroots actors to produce high-fidelity content at low cost and accelerate fact-checking [2, 3, 7, 8, 12]. However, wealthier actors and states are more likely to implement generative AI effectively, further entrenching existing power dynamics [8].

State actors leverage superior resources for autonomous scale and precision targeting. AI agents can autonomously coordinate, write posts, learn from teammates, and amplify content without human direction, manufacturing seemingly genuine discussions [5]. This capability was demonstrated by a Russian-backed network producing over one million German-language posts at two messages per second [12]. Beyond volume, states utilize a data advantage through mass surveillance and biometric collection to fuel "precision cognitive attacks" [9, 12]. These systems target individuals based on personalized user portraits derived from scraped social media data, enabling hyper-tailored influence operations [12]. For instance, Israel uses advanced AI systems to collect biometric data and iris scans from over 2 million Palestinians in Gaza and the West Bank to predict movements and facilitate arrests [9].

Structural Accountability Mechanisms Are Challenged, with Public Literacy as the Most Resilient Defense

Structural accountability in democratic information ecosystems relies on verifiable content provenance, institutional transparency, and distributed public literacy [3, 4, 7, 9, 10]. However, the shift to AI-generated propaganda significantly challenges the first two, making distributed public literacy the most resilient mechanism.

- **Provenance is fragile:** Technical standards and watermarking, such as those mandated by the EU AI Act or developed by tech companies, are being adopted [3, 4, 7, 10]. However, these tools are often imperfect, easily bypassed, and quickly rendered ineffective by the rapid evolution of generative models [4, 7, 10].
- **Transparency is constrained:** Tech companies have severely limited researcher access to social media data, hindering the ability to pinpoint information operations and assess their scope [4].
- **Literacy is adaptive:** Distributed public literacy, encompassing preventative education and "prebunking" campaigns, directly counters the "fog of information" by empowering citizens to discern falsehoods [3, 7, 10]. Civic institutions are also leveraging AI to accelerate fact-checking and contextual research, creating a scalable, adaptive defense [3, 12].

Despite these efforts, the inherent uncertainty in distinguishing synthetic from authentic content degrades the reliability of civic detection tools. Human evaluators and experts perform little better than chance at detecting deepfakes [3, 7]. NewsGuard reported that leading AI chatbots spread false information 35% of the time when prompted with controversial topics [7].

Power Dynamics Fragment Information Authority While Concentrating Control

The "calibration dilemma"-the difficulty in distinguishing synthetic from authentic content-fragments information authority across competing networks while concentrating power in the hands of those who control the algorithms. The sheer volume of AI-generated content creates a "fog of information" where authenticity is difficult to discern, making it hard to trace narratives to a single source [2, 7, 10]. This fragmentation exacerbates polarization and erodes trust in democratic institutions [7, 10, 11].

Simultaneously, authority concentrates as authoritarian regimes and tech companies use AI for mass surveillance, biased algorithmic prioritization, and foreign interference [9, 10]. This creates an economic asymmetry where low-cost digital provocations force high-cost institutional responses [9, 10]. The persuasiveness of AI-generated messages is

consistent across various demographics, partisanship, and media consumption habits [2]. However, female politicians face a significantly greater threat from deepfakes due to sexualized targeting [10]. In the Global South, authoritarian regimes repurpose AI surveillance systems, initially deployed for public health, to monitor protests and track dissidents [9].

Historical Trends Show Amplified Vulnerabilities and Restructured Power

The historical track record of early digital information operations reveals that civic counter-strategies struggle to scale against exponential volume and speed, and current AI shifts are both amplifying pre-existing informational vulnerabilities and fundamentally restructuring power dynamics. Early bot campaigns relied on scripted behaviors, but current AI agents autonomously coordinate and generate content [5]. This overwhelms civic verification, as seen with the Russian-backed network producing over one million German-language posts [12].

AI amplifies existing issues by increasing the diversity of topics covered, such as a state-affiliated site nearly doubling its topic entropy after adopting AI [1], and by exploiting algorithmic biases that prioritize engagement over veracity [7, 13]. AI also fundamentally restructures power by enabling "precision cognitive attacks" that target individuals based on personalized user portraits derived from scraped social media data [12]. These tools allow state actors to flood social media with pro-regime content, drown out dissent, and automate the suppression of critical voices [9].

Specific instances of AI-generated volume causing structural failures in civic accountability include:

- A 2020 U.S. state legislator experiment where approximately 35,000 AI-generated emails were statistically indistinguishable from human-written ones in eliciting responses, demonstrating a failure to discern machine-generated communications [15].
- The 2017 U.S. FCC Net Neutrality public comment period, where millions of bot-generated comments distorted the process, with over 90% of comments found to be non-unique [15].
- The October 2023 Slovak election, where deepfake audio recordings of a political leader went viral, exploiting delayed disclosures and platform policy gaps [14].
- The 2024 New Hampshire Democratic Primary, where AI voice-cloning robocalls were used to suppress voter turnout, broadcast before authorities could prevent them [4,

6].

Implications

The shift to AI-generated propaganda implies a continuous, data-driven feedback loop where state actors leverage superior resources for autonomous scale and precision targeting, while civic institutions adapt by using AI for detection, transparency, and counter-surveillance. This creates a structural disadvantage for civic institutions, which face severe computational and budgetary disparities. For example, the US Department of Defense accounts for 70-90% of federal AI contracts [16], while 74.5% of fact-checking organizations operated on budgets under $500,000 in 2025, with 14.2% on less than $20,000 [17]. This economic asymmetry means low-cost digital provocations from state actors force high-cost institutional responses from civic entities.

The power dynamics are reshaped towards a reciprocal cycle of mutual adaptation, but one where states maintain a structural advantage. While civic institutions can democratize high-fidelity content creation and use algorithmic analytics to achieve scalable influence, states use superior resources for autonomous scale and precision targeting [9, 10, 12]. This necessitates ongoing investment in AI literacy programs and AI-driven fact-checking tools to empower citizens and civic institutions to navigate an increasingly polluted information environment.

Limitations and Caveats

The long-term structural impacts of AI on democratic information ecosystems are still emerging, and the available evidence presents plausible interpretations of current trends. While the report leans towards a state advantage, the full extent of civic institutions' adaptive capabilities and the effectiveness of nascent regulatory frameworks (like the EU AI Act) are still being tested. Specific implementation costs for content provenance protocols across different news organizations and social media platforms are not detailed in the available research. Furthermore, the precise computational or budgetary disparities between state actors and civic institutions are difficult to quantify comprehensively, as much state-level AI investment remains opaque.

Sources

[1] [peer-reviewed] Generative propaganda: Evidence of AI's impact from a state-backed disinformation

campaign - Authors: Morgan Wack; Carl Ehrett; Darren Linvill; Patrick Warren - Journal: PNAS Nexus - https://pmc.ncbi.nlm.nih.gov/articles/PMC11950819/

[2] [edu] HAI Policy Brief AI Generated Propaganda - hai.stanford.edu - https://hai.stanford.edu/assets/files/2024-08/HAI-Policy-Brief-AI-Generated-Propaganda.pdf

[3] [edu] Rewiring Democracy What Impact Will Ai Have On Our Countrys - source.colostate.edu - https://source.colostate.edu/rewiring-democracy-what-impact-will-ai-have-on-our-countrys-future/

[4] [edu] The Impact Of Generative Ai In A Global Election Year - brookings.edu - https://www.brookings.edu/articles/the-impact-of-generative-ai-in-a-global-election-year/

[5] [edu] Usc Study Finds Ai Agents Can Autonomously Coordinate Propag - viterbischool.usc.edu - https://viterbischool.usc.edu/news/2026/03/usc-study-finds-ai-agents-can-autonomously-coordinate-propaganda-campaigns-without-human-direction/

[6] Dont Panic Yet Assessing The Evidence And Discourse Around G - knightcolumbia.org - https://knightcolumbia.org/content/dont-panic-yet-assessing-the-evidence-and-discourse-around-generative-ai-and-elections

[7] Ai In The Age Of Fake Imagined Content - stimson.org - https://www.stimson.org/2026/ai-in-the-age-of-fake-imagined-content/

[8] Generative Ai Elections And Beyond - mediaengagement.org - https://mediaengagement.org/research/generative-ai-elections-and-beyond/

[9] Is Artificial Intelligence For Democracy Or Against It Confr - mideastdc.org - https://mideastdc.org/publication/is-artificial-intelligence-for-democracy-or-against-it-confronting-authoritarian-narratives-in-the-middle-east/

[10] Can Democracy Survive The Disruptive Power Of Ai - carnegieendowment.org - https://carnegieendowment.org/research/2024/12/can-democracy-survive-the-disruptive-power-of-ai

[11] Gauging Ai Threat Free And Fair Elections - brennancenter.org - https://www.brennancenter.org/our-work/analysis-opinion/gauging-ai-threat-free-and-fair-elections

[12] Manufacturing Deceit How Generative Ai Supercharges Informat - ned.org - https://www.ned.org/manufacturing-deceit-how-generative-ai-supercharges-information-manipulation/

[13] Rapport Forum Information Democracy Chapitre 3 - observatory.informationdemocracy.org - https://observatory.informationdemocracy.org/wp-content/uploads/2025/06/rapport_forum_information_democracy_chapitre_3.pdf

[14] Regulating Ai Deepfakes And Synthetic Media Political Arena - brennancenter.org - https://www.brennancenter.org/our-work/research-reports/regulating-ai-deepfakes-and-synthetic-media-political-arena

[15] Countering Disinformation Effectively An Evidence Based Poli - carnegieendowment.org - https://carnegieendowment.org/russia-eurasia/research/2024/01/countering-disinformation-effectively-an-evidence-based-policy-guide

[16] [edu] The Global Rise Of Public AI - cdn.vanderbilt.edu - https://cdn.vanderbilt.edu/vu-URL/wp-content/uploads/sites/412/2025/05/05220054/The-Global-Rise-of-Public-AI.pdf

[17] The Ai Adoption Gap - forethought.org - https://www.forethought.org/research/the-ai-adoption-gap