

SLOPAGANDA: HOW AI-GENERATED NOISE IS RECONFIGURING THE DIGITAL INFOSPHERE

SERGIU BORTOȘ

*Department of Sociology, Social Work and Human Resources,
Alexandru Ioan Cuza University of Iași
bortos.sergiu@uaic.ro*

ORCID: 0000-0003-2610-5735

DOI: 10.2478/saec-2026-0002

Title: Slopaganda: How AI-Generated Noise Is Reconfiguring the Digital Infosphere

Abstract: In the wake of generative artificial intelligence (GAI), the traditional heuristic vocabulary of digital sociology, encompassing computational propaganda, disinformation, and fake news, is increasingly inadequate to capture the qualitative shifts occurring within the digital infosphere. This article conceptualizes “slopaganda” as a distinct sociological and media phenomenon characterized by the mass dissemination of low-effort, AI-generated synthetic content (slop) deployed to manipulate collective attention and reshape epistemic environments. Coined to describe the critical intersection of GAI-generated noise and political or ideological influence, the concept of slopaganda shifts the locus of ideological warfare away from literal persuasion and toward sheer semantic saturation, context collapse, and cognitive exhaustion. This article meticulously deconstructs the structural properties of this emerging threat. Also, a taxonomy of the phenomenon is established, structured along three axes: intent and actors, distribution, and format. Furthermore, the systemic sociological impacts of slopaganda are analyzed, focusing on the production of chronic epistemic fatigue, the disintegration of contextual trust within digital publics, and the manifestation of a decentralized algorithmic “firehose of falsehood.” The analysis concludes by outlining empirical agendas to map the long-term cognitive and structural vulnerabilities of algorithmic societies.

Keywords: *slopaganda; generative artificial intelligence; epistemic fatigue; attention economy; algorithmic curation;*

Introduction

In the contemporary media landscape, the lexicon utilized by digital sociology, political science, and media studies to describe informational pathology is undergoing a profound crisis of explanatory power. The standard vocabulary (anchored by terms such as “fake news,” “disinformation,” and “computational propaganda”) was originally formulated to diagnose the distortions of legacy mass media or the structural vulnerabilities of early Web 2.0 architectures. Consequently, these foundational concepts struggle to accurately conceptualize the structural and qualitative shifts brought about by the rapid democratization and

industrialization of generative artificial intelligence (GAI). Legacy propaganda frameworks were constructed on core assumptions of resource scarcity, centralized bureaucratic control, and high production costs. In these classical models, elite actors carefully constructed ideologically coherent narratives to sway public opinion through mass broadcast mechanisms, deliberately aiming for ideological conversion (Howard and Woolley 2016). Similarly, the academic and political panic surrounding “fake news” during the mid-2010s primarily focused on the mimicry of traditional journalistic formats by opportunistic actors seeking to capitalize on programmatic advertising, or state-sponsored cybertroops coordinating micro-targeted political operations designed to alter specific electoral outcomes (Wardle and Derakhshan 2017). Within these established paradigms, the primary threat vector was the intentional, targeted deception of a human audience through a false but superficially plausible truth claim. Furthermore, the emergence of the social media era and participatory web dynamics has radically diluted the public’s perception of conventional journalism. The uncontrolled proliferation of ad hoc content creation and highly fragmented distribution channels undermines the traditional role of journalists as trusted third-party sensemakers, leaving the public sphere highly vulnerable to post-truth relativism (Sălcudean 2020).

However, the advent of sophisticated GAI models (technologies capable of generating practically infinite volumes of high-fidelity synthetic text, hyper-realistic images, and cloned audio at near-zero marginal cost) renders these earlier analytical models fundamentally obsolete. In the era of algorithmic recommendation feeds, the overarching informational threat is no longer characterized merely by the strategic fabrication of discrete lies or the deliberate subversion of specific empirical facts. Instead, digital societies are witnessing the rapid emergence of a novel phenomenon conceptualized as “slopaganda” (Klincewicz, Alfano, and Fard 2025). Coined as a portmanteau of “AI slop” (low-quality, mass-produced digital clutter) and “propaganda,” **slopaganda** refers to unwanted, low-effort, AI-generated synthetic content that is distributed at unprecedented scale and speed to manipulate collective attention, sentiment, and epistemic conditions for ideological, financial, or political ends (Klincewicz, Alfano, and Fard 2025; Institute of Global Politics 2026).

To rigorously understand the mechanisms of slopaganda, digital sociology must transition beyond the paradigm of literal deception. Traditional disinformation studies assume an adversarial but fundamentally rational relationship between truth and falsehood, where the deceptive agent aims to replace a true belief with a false one in the mind of the target (Wardle and Derakhshan 2017). Slopaganda, by contrast, operates on the principle of semantic saturation and cognitive friction. As scholars have observed, the vast majority of this synthetic content is expressive rather than truth-apt; it does not necessarily demand literal belief from its consumers (Klincewicz, Alfano, and Fard 2025). When a political figure circulates a poorly rendered, GAI-generated image of themselves hugging a synthetic animal, or when state-aligned actors flood social networks with low-effort Lego-style animated conflict videos depicting geopolitical adversaries (Alfano and Klincewicz

2026), the objective is rarely to deceive the viewer into believing these scenarios represent physical reality. Nobody genuinely believes that Donald Trump pilots an F-16 or sits in a casino made of Lego blocks (Klincewicz, Alfano, and Fard 2025). Rather, the primary goal is to trigger immediate, pre-reflective affective associations, reinforce tribal alignments, and systematically monopolize the public’s limited attention span through the continuous flooding of the informational zone.

This article posits that slopaganда represents a necessary theoretical category for contemporary digital sociology. It embodies the structural fusion of computational propaganda with the economics of algorithmic attention farming. By substituting volume, velocity, and affective intensity for narrative coherence and factual plausibility, slopaganда exploits the cognitive vulnerabilities of human decision-making alongside the technical vulnerabilities of modern algorithmic recommendation engines. In doing so, it shifts the focus of informational warfare from the targeted manipulation of specific beliefs to the systematic destruction of the shared epistemic environments necessary for democratic deliberation.

Conceptual genealogy: Deconstructing the boundaries of information pathologies

To isolate the specificity of slopaganда as a distinct sociological phenomenon, it is essential to trace its conceptual genealogy. A rigorous deconstruction of the boundaries between classical propaganda, disinformation, fake news, and slopaganда reveals specific evolutionary stages in modern media manipulation across four primary dimensions: intent, production cost, narrative structure, and epistemic mechanism. Table 1 outlines these distinct modalities to clarify the theoretical divergence of slopaganда from its historical predecessors.

Phenomenon	Primary intent	Production cost & barrier to entry	Narrative structure	Epistemic mechanism
Classical propaganda	State/institutional hegemony; ideological alignment	High (requires mass media infrastructure & coordination)	Highly coherent, centralized, systematic worldview	Ideological conversion and persuasion
Disinformation	Deception; covert disruption of adversary systems	High to medium (requires intelligence apparatus, “active measures”)	Targeted truth-value inversion; specific factual fabrications	Cognitive deception and behavioral manipulation
Fake News	Commercial opportunism; ad-revenue farming; partisan baiting	Low to medium (Requires human labor for clickbait writing)	Journalistic mimicry; highly structured clickbait	Epistemic imitation; exploiting institutional trust

Pheno- menon	Primary intent	Production cost & barrier to entry	Narrative structure	Epistemic mechanism
Slopaganda	Attention cap- ture; epistemic disruption; affec- tive saturation	Near zero (auto- mated via generative AI models)	Hyper- fragmented, surreal, visibly synthetic, incoherent	Semantic pollu- tion; cognitive exhaustion; in- ducing cynicism

Table 1. Comparative taxonomy of propaganda, disinformation, fake news, and slopaganda

Classical propaganda is fundamentally characterized by intentional persuasion and centralized orchestration (Lasswell 1927; Ellul 1965). Operating primarily through mass broadcast media, classical propaganda functions as an organized technique utilized by institutional or state elites to influence human behavior through the manipulation of significant symbols and collective representations. It relies on a high degree of narrative coherence, whether the underlying claims are true, half-true, or entirely fabricated, the message is carefully crafted to present a unified, systematic worldview. The overarching epistemic mechanism of classical propaganda is ideological conversion, seeking to instill a specific set of beliefs, values, or commitments in the target audience to authorize state-sanctioned actions (Klincewicz, Alfano, and Fard 2025).

Disinformation, functioning as a distinct subcategory of the broader framework of “information disorder” (Wardle and Derakhshan 2017), refers to the deliberate dissemination of false, inaccurate, or misleading information with the explicit intent to deceive or cause harm. Unlike classical propaganda, which frequently utilizes selective truths to redirect public attention, disinformation inherently relies on truth-value inversion. Historically, its production cost was significant, requiring robust intelligence infrastructure, coordinated media leaks, or specialized “active measures” to infiltrate legitimate news ecosystems and launder narratives (Rid 2020). The epistemic mechanism of disinformation is precise deception: it targets specific informational nodes to alter the recipient’s internal model of reality, coercing them into executing actions based on falsified premises.

Fake news, which evolved into a dominant analytical construct during the Web 2.0 and early social media eras, represents a highly commercialized and decentralized form of journalistic mimicry. Fake news operations specifically mimic the structural affordances, visual layouts, and stylistic conventions of traditional journalism (e.g., formal headlines, bylines, and datelines) to masquerade as credible reporting (Tandoc, Lim, and Ling 2018; Lazer et al. 2018). The underlying motive is frequently opportunistic and pecuniary, driven by the financial dynamics of programmatic advertising networks. While the production cost of fake news was lower than classical propaganda, it still necessitated human agency to author convincing, emotionally evocative articles. Its epistemic mechanism is one of imitation, explicitly relying on the public’s pre-existing cognitive trust in journalistic structures to bypass critical resistance and achieve viral velocity across digital networks.

Slopaganda emerges as the radical culmination and mutation of these earlier pathologies, fundamentally transformed by the structural logic of GAI and algorithmic recommenders. It is defined as unwanted, mass-produced AI-generated content infused with disruptive political or ideological intent (Klincewicz, Alfano, and Fard 2025). Within the mechanics of slopaganda, the traditional relationship between the message and its volume is inverted: volume itself *becomes* the message.

While classical propaganda seeks ideological conversion, and disinformation seeks factual deception, slopaganda aims at systemic epistemic disruption. Its production cost approaches absolute zero, enabling state actors, commercial click-farms, and political opportunists to flood digital platforms instantaneously. Its narrative structure abandons coherence in favor of hyper-fragmentation and affective resonance. Unlike fake news, which mimics the professional standards of journalism to seek legitimacy, slopaganda relies on low-effort, synthetic “slop” that is often visibly artificial, absurd, or surreal (Encyclopaedia Britannica 2025). It rarely attempts to hide its synthetic nature, instead, it exploits the viewer’s willingness to suspend disbelief in exchange for ideological comfort or sheer entertainment, utilizing highly stylized aesthetics to construct immediate affective associations (Alfano and Klincewicz 2026; Institute of Global Politics 2026). Consequently, the epistemic mechanism of slopaganda is neither deception nor conversion, but semantic pollution and cognitive friction, designed to render the search for objective reality so exhausting that the individual ultimately retreats into a state of deep cynicism or unyielding tribal alignment.

Mapping slopaganda onto macro-sociological frameworks

Jean Baudrillard’s seminal theory of simulacra and simulation (1994) provides an indispensable ontological lens for understanding the nature of slopaganda. Baudrillard famously maps the historical regression of the sign through four successive phases: first, where the image reflects a profound reality; second, where it masks and denatures a profound reality; third, where it masks the absence of a profound reality; and fourth, where the sign has no relation to any reality whatsoever, becoming its own pure simulacrum (Baudrillard 1994).

Slopaganda represents the ultimate structural realization of Baudrillard’s fourth phase, operating entirely within the realm of the hyperreal. AI-generated images permeating social media, such as the viral, grotesque “shrimp Jesus” outputs, hyper-stylized synthetic soldiers, or political leaders rendered as Lego figures in combat scenarios, do not seek to represent, nor even to mask, a pre-existing social reality (Alfano and Klincewicz 2026). They are copies without originals. In the hyperreal space of algorithmic feeds, the boundary between the real and the simulated collapses entirely. The audience is presented with a synthetic aesthetic explicitly designed to invoke raw, pre-reflective affect (such as patriotism, religious devotion, dominance, or outrage) while entirely bypassing the cognitive demands of factual verification. As Baudrillard noted, modern simulation is no

longer a question of imitation or duplication, but rather the act of substituting the signs of the real for the real itself (Baudrillard 1994). In slopaganda, the simulated sign acts as a self-referential closed loop, establishing a parallel digital reality that sustains itself on its own reproduction and the continuous feedback of algorithmic engagement.

Where Baudrillard offers an ontological critique of the synthetic image, Luciano Floridi (2013) provides a vital ethical framework for understanding its environmental impact on society. The infosphere is conceptualized as the totality of the global environment composed of all informational entities, properties, relations, and processes (Floridi 2013). Within this framework, humans are viewed as “inforgs” (informational organisms) deeply embedded within and reliant upon this fragile ecosystem. Floridi introduces the critical concept of “semantic pollution”, a form of conceptual and informational “entropy”, to describe the degradation of the infosphere through the introduction of noise, misinformation, and low-quality data.

Slopaganda acts as the industrial-scale engine of semantic pollution in the modern era. By relentlessly flooding the infosphere with infinite volumes of low-effort, unwanted AI-generated content, slopaganda drastically increases the informational entropy of the digital environment. This systematic pollution degrades the quality of the shared cognitive commons. When the ratio of signal (trustworthy, human-curated information) to noise (GAI-generated slop) declines precipitously, the infosphere becomes toxic and uninhabitable for rational agents seeking empirical truth. Floridi’s core ethical imperatives, which dictate that agents ought not to cause entropy, but rather prevent and remove it to promote the flourishing of the infosphere (Floridi 2013), are systematically violated by slopagandists. These actors exploit the zero marginal cost of GAI to prioritize speed and volume, entirely disregarding the semantic health and sustainability of the digital ecosystem. AI slop functions as a form of “plastic pollution” within the digital information ecosystem (Institute of Global Politics 2026). Because synthetic content is highly durable, inexpensive to mass-produce, and difficult to eradicate once distributed, its primary harm stems not from isolated artifacts, but from the systemic, cumulative degradation of the informational environment over time, operating much like ambient environmental contamination.

From the perspective of surveillance capitalism (Zuboff 2019), slopaganda emerges from a digital economy structured around the extraction, prediction, and monetization of attention, which incentivizes the mass production of low-cost, high-engagement synthetic content. Within this framework, contemporary digital economies are fundamentally structured around the extraction of human experience as free raw material, which is then translated into behavioral data for modification, prediction, and sales (Zuboff 2019). This economic model is entirely dependent on the “attention economy” (Williams 2018), where platforms must maximize user engagement metrics (time spent on screen, clicks, infinite scrolls, and emotional reactions) to fuel their predictive algorithms and monetize targeted ad inventory.

Slopaganda functions as the ultimate optimized instrument of surveillance capitalism. GAI radically lowers the cost of producing content that is specifically engineered to capture, retain, and exploit human attention. Because algorithmic feeds prioritize emotionally arousing and hyper-engaging material, often heavily weighted toward outrage, absurdity, or tribal comfort, slopaganda is structurally rewarded and amplified by the underlying architecture of the platforms. The “Big Other” (Zuboff 2019), Zuboff’s term for the pervasive, sensory-computational apparatus of surveillance capitalism, does not care whether content is true, epistemologically coherent, or human-authored, it only measures behavioral responses and engagement yields. Slopaganda leverages this structural indifference, transforming political influence into a highly programmatic, automated process of behavior-modification-by-volume.

Finally, Hannah Arendt’s profound analysis of totalitarian epistemology provides a chilling political-theoretical foundation for understanding the long-term sociological effects of slopaganda. Arendt (1951) observes that totalitarian propaganda does not rely on persuading the mass public of a specific, coherent lie. Instead, its ultimate objective is the systematic destruction of the human capacity to distinguish between truth and falsehood at all. Arendt famously posited that the ideal subject of totalitarian rule is not the convinced ideological fanatic, but rather people for whom the distinction between fact and fiction (the reality of experience) and the distinction between true and false (the standards of thought) no longer exist (Arendt 1951).

Arendtian mass lying is not designed to force belief in a single falsehood, but to induce a deep, paralyzing cognitive cynicism. When citizens are constantly bombarded with conflicting, high-volume, and obviously fabricated narratives, a condition Arendt described as a “lying government” that must constantly rewrite its own history, the consequence is not that they believe the lies, but that they simply stop believing anything (Arendt 1951).

Slopaganda represents the automated, algorithmic realization of this totalitarian epistemology. By flooding the digital public square with an unmanageable volume of low-effort, synthetic noise, it induces a state of chronic skepticism where all truth claims, regardless of their empirical backing, are viewed as equally fabricated or inherently biased. This cultivated epistemic disorientation paves the way for political apathy, the breakdown of shared reality, and ultimately, authoritarian submission.

The taxonomy of slopaganda: A tri-axial structural framework

To extend the conceptual analysis developed above into a more operational framework, this section proposes a tri-axial taxonomy of slopaganda. The purpose of this taxonomy is not to reduce the phenomenon to a single definition, but to map its internal variation across three analytically distinct dimensions: (1) intent and actors, (2) distribution, and (3) format. In this way, the taxonomy links the article’s theoretical discussion to a more systematic apparatus for empirical sociology,

allowing slopaganda to be examined according to who deploys it, how it spreads, and in what form it appears.

The axis of intent/actors: state/institutional vs. opportunistic/commercial

This axis distinguishes between the underlying motivations, resources, and institutional configurations of the actors deploying slopaganda.

First, state/institutional slopaganda is characterized by strategic, long-term geopolitical and ideological objectives. Here, state actors, intelligence agencies, or institutional cybertroops deploy GAI tools to manufacture domestic or international consensus, conduct cognitive warfare, and denigrate geopolitical adversaries (Howard and Woolley 2016). The primary objective is not monetary gain, but the shaping of reality perceptions and the defense or expansion of political power. For instance, the circulation of AI-generated Lego-themed music videos utilized to mock American political and military leadership during the US-Iran conflicts, or the proliferation of GAI-generated images of political candidates shared by official government or party accounts, represent calculated acts of symbolic warfare (Alfano and Klinecicz 2026). In this modality, slopaganda acts as a highly scalable force multiplier for state-sponsored “active measures,” where narrative plausibility is deliberately sacrificed in favor of affective resonance and overwhelming volume.

Conversely, opportunistic/commercial slopaganda is driven purely by economic imperatives, specifically the maximization of programmatic ad revenue and algorithmic engagement-farming. Produced by commercial click-farms, automated spam networks, and individual opportunists, commercial slopaganda directly exploits platform monetization policies and recommendation heuristics. These actors deploy GAI to generate thousands of low-effort stories, bizarre synthetic images, or hyper-sensationalist clickbait videos with no overarching ideological motive other than capturing traffic. A prime example involves instances where media corporations or shell networks churn out thousands of AI-generated “local” news stories each week to capture regional search engine traffic (Klinecicz, Alfano, and Fard 2025). However, sociological analysis reveals that commercial slopaganda is never politically neutral in its *effects*: by flooding the public square with sensationalized, hallucinated content, it systematically pollutes the shared infosphere, preparing the ground for political radicalization and deep epistemic cynicism (Floridi 2013).

The axis of distribution: from the social graph to the algorithmic interest graph

This axis maps the sociotechnical shift in how informational content is disseminated, curated, and discovered within modern platform environments, marking a transition from user-driven networks to opaque machine-learning algorithms (Gillespie 2014).

In early Web 2.0 social architectures, social graph dissemination dominated. Content distribution was primarily bounded by the user's "social graph", their explicitly defined network of follower-following connections and personal relationships (boyd 2010; Marwick and boyd 2011). Information traveled organically through peer-to-peer sharing, retweets, and interpersonal recommendations. While this architecture was certainly vulnerable to coordinated bot networks and early forms of computational propaganda (Howard and Woolley 2016), it still possessed inherent structural checks: users could evaluate the reputational capital of the human sender, and echo chambers were somewhat bounded by pre-existing social boundaries.

Currently, digital ecosystems have shifted toward algorithmic interest graph dissemination. Modern platforms, most notably TikTok, Instagram Reels, and the algorithmic "For you" feeds on Facebook and X, curate and serve content via opaque recommendation algorithms based entirely on real-time behavioral and engagement metrics, operating independently of a user's actual social network or explicit subscriptions. This architecture is highly vulnerable to GAI exploitation. The recommendation system mathematically prioritizes content that maximizes watch time, emotional arousal (such as angry or outraged emoji reactions), and immediate interaction. Slopaganda could be perfectly optimized for this environment: GAI can continuously generate, A/B test, and deploy micro-variations of a message to determine which synthetic iteration triggers the highest algorithmic traction (Gillespie 2014). Because interest-graph feeds completely bypass the user's social network, slopaganda effortlessly circumvents traditional reputational barriers, thrusting unsuspecting users directly into algorithmic rabbit holes of synthetic noise.

The axis of format: textual-hallucinatory vs. visual-synthetic

This final axis categorizes the medium through which the slopagandistic message is expressed and its corresponding cognitive processing pathway within the human brain.

Textual-hallucinatory slopaganda involves the high-volume generation of written text by large language models (LLMs). This format encompasses synthetic blog posts, automated local news stories, and hallucinated search engine summaries. The defining characteristic of this format is the frequent presence of "hallucinations", plausible-sounding but entirely fabricated facts, quotes, or historical timelines. Textual slopaganda explicitly exploits the public's ingrained trust in written text and formal journalistic style, relying on the sheer speed of production to overwhelm search engine indexers and crowd out authentic, human-verified reporting.

Visual-synthetic slopaganda encompasses AI-generated images, low-effort deepfakes, synthetic audio, and algorithmic animations, including phenomena colloquially referred to as "brain rot". Visual slopaganda operates on a much faster, pre-reflective, sensory-affective register. It does not rely on text-based, logical

arguments, instead, it uses hyper-realistic or surreal aesthetics to create instant, gut-level emotional associations (such as portraying a political candidate as a messianic savior or an adversary as literally demonic) (Klincewicz, Alfano, and Fard 2025). This format achieves extreme viral velocity because visual images are processed significantly faster by the human brain than text, directly exploiting the cognitive “illusory truth effect.” Repeated exposure to an aesthetic or image increases its perceived credibility and emotional salience, even when the viewer consciously acknowledges that the image is synthetic. Additionally, the deployment of AI-generated news presenters (including deepfakes mapped onto the likeness of real journalists) offers distinct commercial advantages, such as continuous 24/7 broadcasting and the elimination of language barriers. However, these synthetic formats suffer from severe structural limitations, notably low interactivity, which ultimately renders the news delivery process and degrades the human element of reporting (Cardaș Răduță 2024).

The sociological consequences of slopaganda

The systemic deployment of slopaganda carries profound macro-sociological consequences, altering the epistemic infrastructure of contemporary democratic societies across the axes outlined above. Rather than merely deceiving isolated individuals, it re-engineers the digital social environment to produce chronic epistemic fatigue, collapse contextual trust, and automate the firehose of falsehood at scale. AI-slop reconfigures political communication by blurring the boundaries between entertainment, misinformation, and propaganda; through the rapid diffusion of AI-generated memes and synthetic visuals, political and commercial actors can cultivate false authenticity, amplify polarizing narratives, and bypass traditional media scrutiny (Gross and Colson 2025). Its primary psychological effect is the production of epistemic fatigue, a state of cognitive depletion in which individuals abandon the effort to verify information and distinguish fact from fiction (Siewier 2025). In informational environments where GAI generates an endless stream of plausible but fabricated content, the cognitive cost of verification increases exponentially, while fact-checkers, journalists, and citizens must expend considerable time and labor to debunk each claim. Under these imbalanced conditions, public responses increasingly shift toward rational ignorance, cynicism, apathy, and heuristic reliance on emotional comfort and tribal alignment rather than factual accuracy (Williams 2018; Klincewicz, Alfano, and Fard 2025).

Slopaganda also accelerates the collapse of contextual trust and reputational anchors. In complex sociological systems, trust functions as an epistemic institution sustained by shared contexts, reputational networks, and institutional verification chains (Origgi 2018). By flooding the digital public sphere with low-effort, hyper-realistic, or hallucinated content, slopaganda produces a liar’s dividend, enabling malicious actors to dismiss genuine evidence simply by labeling it as AI-generated slop. At the same time, the omnipresence of GAI noise undermines legitimate information channels: when local newspapers, search engines, and so-

cial media platforms are overrun by synthetic content, the reputational anchors that historically mediated public trust begin to collapse (Origgi 2018). Platforms already suffer from context collapse, where distinct audiences, temporalities, and normative expectations are flattened into a single decontextualized interface (Marwick and boyd 2011; Meyrowitz 1985). Slopaganda weaponizes this condition by merging authentic news, parody, spam, and state-sponsored deception into one rapid-fire algorithmic stream, leaving users unable to recover the original context of any message and pushing trust toward affective and tribal signifiers.

Finally, slopaganda represents the automation and decentralization of the firehose of falsehood model. As theorized by Paul and Matthews (2016), this strategy relies on high-volume, multichannel dissemination, rapid repetition, indifference to factual consistency, and disregard for objective reality. Whereas earlier versions required centralized coordination and substantial resources, GAI now enables virtually any actor, from a lone troll to a state intelligence unit, to deploy a localized firehose targeting specific controversies, groups, or electoral moments. Algorithmic recommendation systems further amplify this dynamic by rewarding engagement over coherence, effectively turning platforms into distribution infrastructures for synthetic cascades. Importantly, the firehose of falsehood does not aim to persuade audiences of a single alternative narrative; rather, it overwhelms cognitive capacity through volume and contradiction, producing epistemic disorientation and exhaustion (Paul and Matthews 2016). The broader sociological result is the erosion of shared verifiable reality and the transformation of democratic societies into fragmented fields of epistemic cynicism, where consensus-building and collective action become increasingly difficult.

Conclusion and future directions

This analysis has conceptualized slopaganda as a distinct, highly disruptive sociotechnical phenomenon representing the radical maturation of generative artificial intelligence in the realm of computational warfare. While pre-AI information pathologies, such as classical propaganda, disinformation, and fake news, depended heavily on the strategic crafting of truth-apt narratives or stylistic journalistic mimicry to persuade or deceive, slopaganda operates almost entirely on a logic of volume, speed, and affective saturation. In the mechanics of slopaganda, the ideological message is superseded by the automated “flooding” of the digital commons, transforming the infosphere into a highly entropic, structurally hostile environment where the search for objective reality is replaced by systemic epistemic fatigue and tribal cynicism.

By anchoring the concept of slopaganda in Baudrillard’s hyperreality, Floridi’s information ethics, Zuboff’s surveillance capitalism, and Arendt’s totalitarian epistemology, it becomes evident that slopaganda is not merely an incremental technological shift. Rather, it represents a structural re-engineering of the public square. It ruthlessly exploits the human cognitive vulnerabilities and the economic imperatives of programmatic advertising, routed through opaque platform interest

graphs that mathematically prioritize affective friction over genuine social connection. The tri-axial taxonomy proposed in this article, categorizing slopaganda across actors (state vs. commercial), distribution channels (social graphs vs. interest graphs), and formats (textual vs. visual), provides a critical, structured analytical toolkit for future empirical research.

To advance the critical sociology of algorithmic regimes and protect the integrity of digital democratic institutions, at least three specific directions for future empirical research are proposed. First, an algorithmic audit of interest graphs should be undertaken by sociologists and computational researchers through rigorous, independent platform-risk assessments that measure how recommendation systems such as those governing TikTok and Instagram Reels prioritize synthetic and GAI-generated content under varying conditions of social polarization. This requires empirically analyzing “filter bubble” dynamics and measuring the velocity at which neutral user feeds are pulled into high-volume, automated “slop” streams (Vision of Humanity 2024). Second, epistemic fatigue ethnographies should employ qualitative and mixed-method approaches to investigate the lived, everyday experience of “epistemic fatigue” and “rational ignorance” among diverse demographic groups. Such studies must map how cognitive overload and constant exposure to hyper-realistic visual simulations degrade citizens’ willingness to engage in political deliberation, fact-checking, and collective civic action (Siewier 2025). Third, the political economy of commercial slop farming requires empirical work tracing the financial pipelines linking programmatic ad exchanges, generative AI content generators, and decentralized click-farms globally. Understanding how commercial actors monetize the systematic pollution of the infosphere is essential for developing targeted regulatory, tax-based, or technical countermeasures, such as cryptographic watermarking and algorithmic demotion, to defend the cognitive sovereignty of democratic societies (Klincewicz, Alfano, and Fard 2025). Only by treating the digital infosphere as a fragile, shared ecological resource can media scholars and sociologists construct the frameworks necessary to preserve the baseline empirical realities required for human flourishing.

Bibliography:

Alfano, Mark, and Michał Klincewicz. 2026. “Slopaganda Wars: How and Why the US and Iran Are Flooding the Zone with Viral AI-Generated Noise.” *The Conversation*.

Arendt, Hannah. 1951. *The Origins of Totalitarianism*. New York: Harcourt.

Baudrillard, Jean. 1994. *Simulacra and Simulation*. Translated by Sheila Glaser. Ann Arbor, MI: University of Michigan Press.

boyd, danah. 2010. “Social Network Sites as Networked Publics: Affordances, Dynamics, and Implications.” In *A Networked Self: Identity, Community, and Culture on Social Network Sites*, edited by Zizi Papacharissi, 39-58. New York: Routledge.

Encyclopaedia Britannica. 2025. “AI Slop.” <https://www.britannica.com/technology/AI-slop>.

Cardaș-Răduță, Dan-Laurențiu. 2024. "The Effectiveness and Limitations of Artificial Intelligence (AI) Use in Current Romanian Journalism." *Saeculum* 57 (1): 111-19.

Ellul, Jacques. 1965. *Propaganda: The Formation of Men's Attitudes*. New York: Alfred A. Knopf.

Floridi, Luciano. 2013. *The Ethics of Information*. Oxford: Oxford University Press.

Gillespie, Tarleton. 2014. "The Relevance of Algorithms." In *Media Technologies: Essays on Communication, Materiality, and Society*, edited by Tarleton Gillespie, Pablo Boczkowski, and Kirsten Foot, 167-94. Cambridge, MA: MIT Press.

Gross, Eduard-Claudiu, and Alicia J. M. Colson. 2025. "AI-Slop and Political Propaganda: The Role of AI-Generated Content in Memes and Influence Campaigns." *EON* 6 (3): 289-98.

Howard, Philip N., and Samuel Woolley. 2016. "Political Communication, Computational Propaganda, and Autonomous Agents – Introduction." *International Journal of Communication* 10: 4882-4890.

Institute of Global Politics. 2026. "AI Slop" and the Information Ecosystem: Insights from a Cross-Sector Convening. New York: Columbia University School of International and Public Affairs. https://igp.sipa.columbia.edu/sites/igp/files/2026-06/AI%20Slop%20and%20the%20Information%20Ecosystem_IGP%20Report.pdf

Klincewicz, Michał, Mark Alfano, and Amir Ebrahimi Fard. 2025. "Slopaganda: The Interaction Between Propaganda and Generative AI." *Filosofiska Notiser* 12 (1): 135-62.

Lasswell, Harold D. 1927. "The Theory of Political Propaganda." *American Political Science Review* 21 (3): 627-31.

Lazer, David M. J., Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, David G. Rand, Filippo Rossi, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, and Jonathan L. Zittrain. 2018. "The Science of Fake News." *Science* 359 (6380): 1094-96.

Marwick, Alice E., and danah boyd. 2011. "I Tweet Honestly, I Tweet Passionately: Twitter Users, Context Collapse, and the Imagined Audience." *New Media & Society* 13 (1): 114-33.

Meyrowitz, Joshua. 1985. *No Sense of Place: The Impact of Electronic Media on Social Behavior*. New York: Oxford University Press.

Origgi, Gloria. 2018. *Reputation: What It Is and Why It Matters*. Translated by Stephen Holmes and Noga Arikha. Princeton, NJ: Princeton University Press.

Paul, Christopher, and Miriam Matthews. 2016. *The Russian "Firehose of Falsehood" Propaganda Model: Why It Might Work and Options to Counter It*. Santa Monica: RAND Corporation.

Sălcudean, Minodora. 2020. "Journalism in the Paradigm of Media Pluralism: Role, Mission, Values." *Saeculum* 50 (2): 29-38.

Siewier, Malwina Anna. 2025. "Resilience as a Strategic Pillar of Cognitive Security." *Humanities and Social Sciences* 32 (4): 169-183.

Tandoc, Edson C., Zheng Wei Lim, and Richard Ling. 2018. "Defining 'Fake News': A Systematic Taxonomy of Scholarly Definitions." *Digital Journalism* 6 (2): 137-153.

Vision of Humanity. 2024. "Why TikTok Isn't Really a Social Media App." <https://www.visionofhumanity.org/why-tiktok-isnt-really-a-social-media-app/>.

Wardle, Claire, and Hossein Derakhshan. 2017. *Information Disorder: Toward an Interdisciplinary Framework for Research and Policy Making*. Strasbourg: Council of Europe.

Williams, James. 2018. *Stand Out of Our Light: Freedom and Resistance in the Attention Economy*. Cambridge: Cambridge University Press.

Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs.