

# AI Slop and the Information Ecosystem

**By Jen Weedon, Camille François, and Jeremie Ponak**

*with Ryan Beiermeister, Ben Brooks, Patricia Cartes, Shane Culloty, Renée DiResta, Alix Dunn, Sourojit Ghosh, Jason Koebler, Rob Reich, Eryk Salvaggio, Artemis Seaford, Sarah Shirazyan, Aidan Walker, and Andrea Wong*



---

## **CONTENTS**

<b>Introduction: Why AI Slop, Why Now?</b>	<b>2</b>
<b>How to Characterize AI Slop</b>	<b>8</b>
<b>Where Slop Matters: Mapping Impacts</b>	<b>17</b>
<b>Why Slop Exists: Supply, Demand, and Incentives</b>	<b>29</b>
<b>What Can Be Done: Intervention Pathways</b>	<b>37</b>
<b>Implications of Slop</b>	<b>42</b>
<b>What We Still Don't Know: A Research Agenda</b>	<b>45</b>
<b>Endnotes</b>	<b>49</b>
<b>Acknowledgments</b>	<b>55</b>
<b>Methodology</b>	<b>56</b>

### **June 2026**

The findings, interpretations, and conclusions expressed in this publication are entirely those of the author(s) and do not necessarily reflect the views of the Institute, its leadership, Columbia University, or the School of International and Public Affairs (SIPA). Any external support for this research did not influence its content.

# AI Slop and the Information Ecosystem

---

**Jen Weedon, Camille François, and Jeremie Ponak**

*with Ryan Beiermeister, Ben Brooks, Patricia Cartes, Shane Culloty, Renée DiResta, Alix Dunn, Sourojit Ghosh, Jason Koebler, Rob Reich, Eryk Salvaggio, Artemis Seaford, Sarah Shirazyan, Aidan Walker, and Andrea Wong*

## **ABSTRACT**

*What is AI slop, and what does its emergence mean for information ecosystems? These proceedings from a March 2026 convening held by Columbia University bring together perspectives from experts and practitioners and provide a systematic overview of the phenomenon of AI slop along with a proposed research agenda. The research agenda is designed to assess the significance of AI slop and, when its effects are negative, to identify measures to protect the health of information ecosystems and democracy.*

---

## INTRODUCTION: WHY AI SLOP, WHY NOW?

Synthetic content—text, images, audio, and video—has long been a part of the online world. The age of generative AI, however, has drastically lowered the barriers to creating such content, and distribution opportunities have expanded. The result is a new era of ambient synthetic content, prevalent across most social media sites and online platforms. Much of this content is optimized to grab attention, rank high in search, or make the most of algorithmic trends, and some research documents these shifts at scale.<sup>1</sup> Generative AI has quietly become part of the default tool kit for quick and easy content generation in media, education, marketing, research, and political communications.

This shift is both technical and economic, as generative AI has collapsed the marginal cost of content production, and when combined with the expansion of distribution platforms, it is reshaping the information environment. Content can now be produced at near-zero cost, and near-infinite scale. Volume itself has become a defining force, overwhelming human attention, institutional filters, and epistemic norms. Long before generative AI, information scientists warned that overload does not merely create inconvenience, but can actively degrade sensemaking, trust, and judgment.<sup>2</sup> Generative systems accelerate these dynamics, producing an informational surplus that is challenging to audit, curate, or even recognize as synthetic.

Enter *AI slop*, a defining feature of an age of cheap and varying-quality synthetic text, images, audio, and video, coupled with easy and near-instantaneous distribution. The term *AI slop* describes content that feels low-effort, repetitive, or strangely hollow, but the *slop* label also raises other questions: History suggests that content initially dismissed as “low quality” or “trash” often gets reevaluated later. Whether tabloids or memes, early internet forums or reality TV, cultural critics have repeatedly written off new forms of mass content, only to later recognize that they mattered, socially or culturally, more than expected. As one recent paper put it, “low” content is often dismissed first and understood later as a legitimate or even important cultural production.<sup>3</sup>

### WHAT IS AI SLOP?

*AI slop* is a term typically used to refer to high-volume synthetic content (text, images, video, audio) generated quickly and optimized for engagement rather than depth. Slop is a subset of AI-generated content (AIGC). Slop content often appears in feeds as repetitive memes, auto-written articles, listicles, or uncanny visuals posted in bulk. As generative AI software has grown from early text-based generators like ChatGPT to today’s video generation tools (e.g., Seedance,<sup>4</sup> Google DeepMind’s Veo<sup>5</sup>), slop’s media has evolved alongside technological developments.

### **WHAT IS AI SLOP? (continued)**

Despite widespread use of the term, there is no consensus definition of *AI slop*, raising important questions about what the concept includes and excludes. Slop is certainly not confined to algorithmically determined social media feeds (that is, feeds determined by automated ranking and recommendation systems that influence what content a user sees based on factors like predicted relevance, likelihood of engagement, or other objectives, rather than chronology alone); concerns about slop have extended to all manner of contexts, including *workslop* (slop in the workplace), academic research slop, AI-generated books and audio, and more.

These developments are why AI slop warrants examination. On the one hand, the growing presence of AI slop appears to embody many of the anxieties associated with platform degradation, information pollution, and what advocate Cory Doctorow has called the “enshittification” of digital ecosystems.<sup>6</sup> On the other hand, slop’s emergence may reflect deeper structural mismatches between cultural demand and human creative supply that are being leveraged for political and economic gain. Treating AI slop solely as a failure risks oversimplifying or blurring a more complex set of social, economic, and cultural dynamics that are only beginning to become clear.

Content moderation, already strained by platform deregulation<sup>7</sup> and the trimming of trust and safety teams,<sup>8</sup> faces new challenges: Slop floods systems with low-cost, ambiguous content that is difficult to classify or justify moderating. Platforms already need to calibrate the trade-offs between encouraging the creation and proliferation of user-generated content (synthetic or otherwise) and taking on the associated cost of moderating that content. These trade-offs are occurring alongside more existential questions involving free speech, agency, the role of private actors in shaping information spaces, and evolving consumer preferences amid rapid technological changes.

### **A SALON ON AI SLOP**

On March 5, 2026, Columbia University School of International and Public Affairs (SIPA)’s Institute of Global Politics (IGP) Technology and Democracy Initiative and the Hewlett Foundation convened an interdisciplinary group of experts to define and discuss the implications of AI-generated content for the health of the information ecosystem and society more broadly. The convening was held under Chatham House Rule, with 20 representatives in attendance from industry (including social media companies and frontier model companies that are developing the most advanced AI systems), academia, investigative journalism, policy advocacy, and internet and culture studies. Five attendees presented short provocations on different aspects and framings of AI slop to surface complexities and disagreements.

Below are the proceedings of this convening—known as the “Slop Salon”—along with key takeaways, informed by a series of pre-convening interviews, a review of the burgeoning literature on this topic, discussions held during the convening itself, and follow-up contributions from participants.

The collective goal of the workshop was not to settle a debate, but to surface the key questions and tensions that arise in the process of defining and understanding slop. If *slop* is to be a useful term, it must be clear what slop is and what it isn't; where slop causes real harm and where it produces joy or satire. A clear understanding of slop might permit a broader examination of where and when slop serves unexpected functions, and clarity in thinking will aid the efforts of different communities—researchers, journalists, platforms, and policymakers—grappling with both the upsides and the downsides of slop in practice. The following questions organized the substantive work of the Slop Salon and these proceedings:


- 1. What is AI slop, and what isn't?** The proceedings articulate some of the ambiguities any working definition must navigate.
- 2. Under what conditions is AI slop useful as a term?** Despite semantic imprecision, the term has become shorthand for consternation about content scale, authenticity, and interplay with platform incentives.
- 3. Where should we draw the lines between slop and useful automation, creative assistance, and digital pollution?** This question reflects legitimate disagreements about slop's value and purpose. The proceedings also explored cases where slop serves unexpected functions, such as novel expression and political satire, alongside cases where it operates as infrastructure for scams. None of these categories is easily dismissible, and the categories are not mutually exclusive.
- 4. How are platforms and third-party actors currently encouraging or responding to slop, or both? What gaps remain in governance approaches?** Platforms (defined broadly as the technologies, systems, and infrastructure that enable users to interact, communicate, or exchange goods and services) have begun to experiment with intervention strategies involving AI-generated content. The proceedings map some of these responses and identify where governance breaks down, particularly in ambiguous areas where content is of low quality but not obviously harmful, and where moderation frameworks have the least traction.

To capture the many different efforts made by commentators, scholars, and practitioners to sketch the outlines of an understanding of AI slop, these proceedings are filled with anecdotes, comparisons, and frameworks. The authors of this report chose to lay out multiple frameworks and metaphors, capturing the wide range of plausible interpretations, their tensions, and opportunities for them to be reconciled.

## THE EMERGENCE OF THE TERM *AI SLOP*

The term *AI slop* crystallized in the wake of rapid advances in generative AI, particularly after the public release of OpenAI’s ChatGPT in late 2022 and open image models like Stable Diffusion. The phrase gained wider cultural visibility as mainstream outlets chronicled the flood of AI-generated books on Amazon,<sup>9</sup> AI-automated YouTube channels,<sup>10</sup> and synthetic images circulating during real-world crises and election cycles, raising concerns not only about misinformation but about saturation, discernment fatigue, and aesthetic degradation. In 2025, *AI slop* was dubbed a “word of the year” by Merriam-Webster,<sup>11</sup> and inspired a slop-focused episode of comedian John Oliver’s HBO show *Last Week Tonight*,<sup>12</sup> signaling its arrival as a shorthand for anxieties about scale, authenticity, and platform incentives rather than merely a meme or insult.

**Figure 1: Google Trends in terms from January 2024 to March 2026, just prior to the Slop Salon. Screenshot captured on May 21, 2026.<sup>13</sup>**



The implications of vague definitions and conceptual ambiguity can be long lasting and potentially counterproductive. As with earlier terms like *fake news*, which began as a descriptive label for fabricated stories but collapsed into a partisan weapon, loose usage can erode analytical clarity and sometimes provoke backlash.<sup>14</sup> When “misinformation” and “disinformation” became catch-all categories, particularly after events like the 2016 US presidential election, they were applied to issues as varied as coordinated inauthentic behavior and ordinary disagreement on topics such as politics, public health, and integrity in elections. The definitional elasticity enabled important research and policy interventions of varying efficacy, but also fueled claims of censorship, platform overreach, and ideological bias.

### **THE EMERGENCE OF THE TERM *AI SLOP* (continued)**

A parallel dynamic could unfold around the term *AI slop*. If the term's use collapses heterogeneous phenomena (spam, aesthetic distaste, automation at scale, industrialized political messaging, experimental creativity, etc.) into a single pejorative bucket, it risks obscuring important distinctions and unduly influencing interventions. Being very broad in *slop*'s definitional qualities may encourage interventions that target volume, style, or quality over other *slop* characteristics and impacts, and it could invite claims of bias or arbitrariness. Similarly, too much risk aversion in mitigating *AI slop* may exacerbate downstream impacts. Thus, the definitional contours of *slop* have meaningful implications for institutional responses across sectors.

**Treating *AI slop* solely as a failure risks oversimplifying or blurring a more complex set of social, economic, and cultural dynamics that are only beginning to become clear.**

**Figure 2: Domains of slop discussed at the workshop**

Domain	Research Questions	Elements of Discussion
<b>Definitions, Contours and Frames</b>	What is slop? How is the concept defined and where are its limits?	Definitions, scope, and boundaries Historical and conceptual lineage and comparisons Areas of ambiguity and contestation
<b>Incentives and Optimization</b>	Why does slop emerge? What are the behavioral and economic drivers of slop creation and production at varying scales? What incentives and optimization processes drive its creation?	Supply, demand, and monetization dynamics Platform incentives, including ranking systems, optimization features for engagement, and monetization
<b>Systems and Infrastructure</b>	What systems and infrastructure enable slop’s scale, persistence, and evolution?	Content generation pipelines (human-curated, automated, and hybrid) Distribution, amplification and targeting networks Recursive feedback loops, including models being trained on synthetic data
<b>Aesthetics, Meaning, and Cultural Signals and Implications</b>	Why and when does slop resonate? What makes slop persuasive, legible, or emotionally resonant, and in which contexts?	Visual and narrative patterns and aesthetic styles Cultural resonance and memes Perception, plausibility and “vibes” as signals of meaning
<b>Impacts and Intervention Points</b>	What are the consequences and impacts of slop for different audiences, domains and timelines? How can these consequences and impacts be addressed? What are the levers for response in different parts of the ecosystem?	Epistemic, normative, and social effects Degradation of information quality and increased costs on users and systems to determine accuracy and relevance Policy, product and market solutions and interventions

The conceptual framework shown in Figure 2 organizes the study of AI slop into layered analytical domains from definitional framing and incentive structures to production systems, cultural meaning, and downstream impacts. It aims to highlight how AI slop emerges from interconnected technical, economic, and social dynamics rather than any single source.

---

## HOW TO CHARACTERIZE AI SLOP

In the following section, we explore some of the theoretical framings used by researchers, as well as those offered up by participants in our workshop, to clarify the characteristics of AI slop.

One example is the “7Vs” framework, which characterizes AI slop along seven dimensions: volume (the scale of production), velocity (the speed of generation and circulation), variety (the range of forms and genres), value (erosion or transformation of cultural and epistemic worth), verification (problems of truth, trust, and provenance), visibility (algorithmic amplification and discoverability), and virality (meme-like diffusion and replication).<sup>15</sup>

Another approach focuses on qualities that slop content tends to share: superficial competence (slop’s veneer of quality is belied by a deeper lack of substance), asymmetrical effort (it takes vastly less effort to generate such content than would be the case without AI), and mass producibility (it is part of a digital ecosystem of widespread generation and consumption).<sup>16</sup>

*The Indicator*, a practitioner-focused investigative publication that provides original reporting, research, and tutorials on online harms and digital threats, coauthored by Alexios Mantzarlis and Craig Silverman, published a tentative slop taxonomy based on its observations. *The Indicator’s* taxonomy is practitioner-focused and framed for regulators, researchers, and trust and safety workers navigating AI slop.<sup>17</sup> Mantzarlis’s taxonomy addresses some of AI slop’s form and purpose, where slop’s form can be described as expressive or deceptive. Expressive slop “serves to visualize the creator’s ideas in a captivating but overtly artificial form ... the creator is typically self-labeling their content prominently and audiences are for the most part aware that the content is a manifestation of the creator’s wishes, a fantastical recreation of something they cannot as captivatingly represent through other means.”<sup>18</sup> Deceptive slop, in contrast, may not intentionally signpost its synthetic nature, and can be “bombastic and relatively easy to debunk, but potentially believable when you quickly scroll past it on a digital feed.”<sup>19</sup>

Mantzarlis proposes two primary motivations for purveyors of slop emerging from his investigations: sociopolitical or economic. The purpose of sociopolitical slop is to advocate, attempt to change minds, or otherwise pursue some kind of social or political advantage. Economic slop seeks to make money directly by scams or defrauding, directly or indirectly by building large followings on platforms that can be monetized through advertisements and subscriptions. It is worth noting here that the potential creative, countercultural purposes of slop explored below in the proceedings are not captured in this early taxonomy.

**Figure 3: A tentative taxonomy of AI slop from *The Indicator*<sup>20</sup>**

		FORM	
		Expressive	Deceptive
		Content that serves to visualize the creator’s ideas in a captivating but overtly artificial form	AI slop that doesn’t intentionally signpost its synthetic nature to ensnare the digitally naive
GOAL	Social / Political	Content whose goal is to advocate for a position or change people’s minds	AI slop that doesn’t intentionally signpost its synthetic nature to ensnare the digitally naive
	Economic	Content whose goal is to make money directly or indirectly	AI slop that doesn’t intentionally signpost its synthetic nature to ensnare the digitally naive

This taxonomy offers an early glimpse of attempts to differentiate slop by intent—an exercise that is repeated by other scholars, and valuable for platforms in some use cases, but difficult to discern at scale.

**Slop does not always operate as the harmful end point itself, but rather as the connective tissue that sustains exploitative infrastructures.**

### **SHRIMP JESUS: WHEN SLOP GOES VIRAL**



“Shrimp Jesus” was one of the bizarre AI-generated images blending Jesus with shrimp and other surreal elements that spread widely on Facebook in 2024 and catalyzed the slop conversation.<sup>21</sup> Investigations showed many of these posts were created by spam networks optimizing for engagement and later redirected users to ad-heavy or scam websites.

To complement scholars’ efforts to classify slop, Chantal Shaib et al. of Northeastern University proposed a framework for measuring AI slop; in their research, they defined slop as generic, low-quality, or misleading AI-generated content. The researchers organized measurement along three dimensions, roughly categorized as (1) how useful the information is, (2) how accurate and appropriately framed it is, and (3) how well written it is. Drawing on annotations from professional editors across news articles and question-and-answer (Q&A) responses, the researchers found that although “is this slop?” judgments are inherently subjective, the editors’ assessments consistently reflected aspects of these three underlying dimensions. The authors found that which dimension matters most varies by context (for example, style and relevance dominate in news contexts, while accuracy and structure matter more in Q&A content). Their research also found that existing automated tools failed to reliably replicate human judgment.<sup>22</sup>

### **THE CHALLENGE OF DEFINING SLOP: INSIGHTS FROM THE CONVENING**

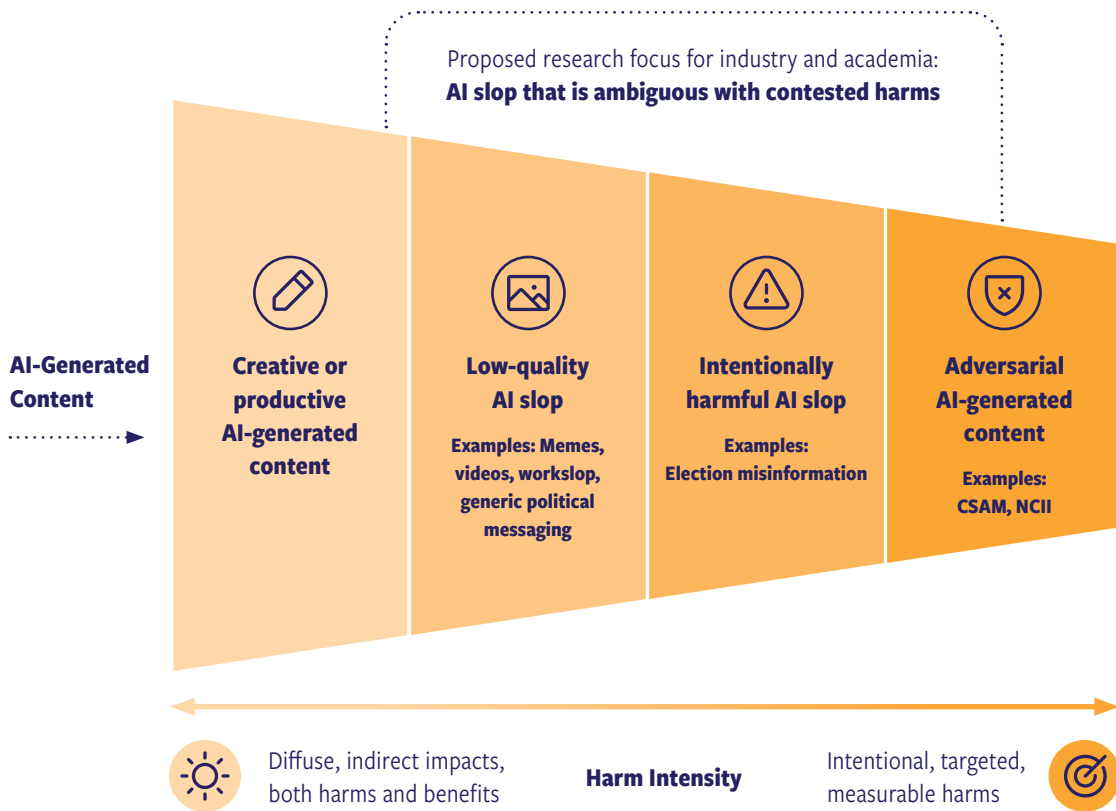
The most persistent disagreements from the convening were definitional and conceptual rather than empirical, reflecting genuine differences in framing across research, practitioner, and platform

communities. A recurring theme in the discussions challenged the utility of the concept and vocabulary of slop and questioned whether *AI slop* as a term does, or should, carry inherent normative judgment. The group noted that some content dismissed as slop has genuine value for both producers and consumers. Using *slop* in a pejorative way in harm-focused work usefully flags problems that platforms and policymakers face, but risks flattening content that serves expressive or other positive functions. A value-neutral conceptualization might be analytically preferable to some, but this remained an area of disagreement. Consider, for example, an analogous debate about the capacity of AI models to induce shifts in user preferences or behaviors. Some distinguish between persuasion and manipulation, where the former is salutary and the latter problematic. Others study the phenomenon under the banner simply of “persuasion studies,” a field of study where persuasion is considered value-neutral, referring only to the capacity to change preferences or behaviors, not to whether such changes are desirable or undesirable.<sup>23</sup>

Salon participants developed the Slop Funnel (Figure 4) as a shorthand framework during a breakout session that distinguishes AIGC types (including slop) by harm intensity and intent; this framework puts slop squarely between benign creative uses and clearly harmful content that can be AI-generated, such as non-consensual intimate imagery (NCII) and child sexual abuse material (CSAM). Governance is particularly difficult in this designated middle territory, where content moderation frameworks have varying levels of traction, and where team responsibilities may not have clean oversight. This was, across salon participants’ efforts, the most concrete attempt at conceptualizing the universe of AI-generated content and creating clear definitional boundaries to differentiate AI slop from other forms of AIGC. The funnel’s purpose is to visualize aspects of slop and its potential harms, and to propose some prioritized areas of focus for platforms and independent researchers to consider.

The Slop Funnel begins with the broad universe of AIGC, which includes creative, productive, and benign uses such as art, writing, work outputs, and productivity tools. Moving through the funnel to the second category, content becomes increasingly narrow, low-effort, and mass-produced, optimized for engagement and algorithmic reach. This is the first category meaningfully described as AI slop: often not explicitly harmful, but contributing to attention capture, information clutter, and degradation of the broader information environment. Further along lies intentionally harmful AIGC produced or deployed with strategic objectives such as election misinformation and coordinated influence operations, where harm is no longer incidental but instead purposeful and directional, and arguably not slop at all. At the narrowest end sits adversarial AIGC, defined by explicit intent to deceive, exploit, or cause harm; this includes scams, harassment, NCII, and CSAM.

**Figure 4: The Slop Funnel**



The two ends of the spectrum present platforms with clearer positions. AIGC that is not algorithmically amplified, or that is “high-quality,” may create issues for platforms, creators, or economic incentives, but these impacts could be diffuse, longer-term, and difficult for platforms to moderate. Adversarial AI content including CSAM, NCII, fraud, and scams is obviously nefarious, and platforms are broadly expected to address it, even when they do so ineffectively.

Several participants noted that *slop* has already come to encompass a much wider range of content—including material generated without algorithmic distribution—and that the term has reached sufficient adoption to resist the slop funnel’s narrowed conceptions of slop areas for prioritized consideration. For example, non-algorithmically distributed AIGC can still cause real harm, such as poor work products that consume employee time, excess energy consumption, and negative labor market effects. The challenge is that many of the assumed costs associated with AI slop may be spread across society rather than concentrated in one obvious event or victim. For example, people may spend more time verifying information, moderators and journalists may face growing workloads, search and recommendation systems may become less reliable, and trust in online content may gradually erode. These harms may be cumulative and systemic rather than singular and measurable, and excluding them from prioritized slop categories risks defining the problem according to what platforms can most easily measure or moderate—such as clearly illegal or overtly malicious content—rather

than according to where the burdens and harms actually fall. In practice, this could shift attention away from broader degradations to the information environment simply because they are harder to quantify, attribute, or govern.

In sum, these diverse theoretical frames and accompanying convening discussions underscored the value of examining AI slop through multiple lenses and units of analysis, ranging from individual content artifacts (such as a single image, video, meme, or post), to account- or creator-level dynamics, to the broader sociotechnical systems that produce, distribute, amplify, and reward slop at scale. The discussions also highlighted that different units of analysis may yield different conclusions about the nature, significance, and potential harms of AI slop.

### **METAPHORS FOR SLOP: POLLUTION, PLASTIC, OR PLATFORM LOGIC?**

Scholars have proposed metaphors to situate AI slop, framing it as a form of systemic by-product akin to junk food or plastic, rather than an isolated failure of individual quality or taste. From this perspective, generative AI systems are like other technologies that simultaneously lower barriers to production and introduce widespread negative externalities. Like the prominence of junk food in food deserts, AI-generated content can be inexpensive, versatile, abundant, and easy to produce and distribute. Also like plastic in the physical environment, it risks degrading digital environments through cognitive overload, misinformation, bias, and algorithmic homogenization (e.g., when AI systems and recommendation algorithms converge on the same content, gradually narrowing cultural and intellectual diversity and representation at scale).<sup>24</sup>

Importantly, metaphors emphasize that AI slop is as much a product of the systems, incentives, and platforms that prioritize volume and velocity as it is of the AI models that generate the content itself. Like junk food, AI slop may serve as a rational recourse for individuals facing structural constraints—including not only economic disadvantage, but also time scarcity, excess cognitive load, and unequal access to high-quality alternatives.<sup>25</sup> This suggests that slop could offer a way to participate in cultural and informational production that would otherwise be inaccessible due to limited digital resources or institutional gatekeepers. Slop is thus not merely exploitative or degrading, but symptomatic of unmet demand within contemporary information economies, fitting into a larger dynamic whereby social costs inflicted by industries (whether the tech industry's impact on the information ecosystem or the fast-food industry's impact on the environment and workers) indirectly create the demand for the patches they provide.

Although describing slop as a system rather than an individual creation of content offers many epistemological benefits, it also imposes some limitations. For one, human consumers and platforms rarely have access to the full supply chain context of a given piece of content, and may rely on snap judgment or incomplete heuristics, respectively, to decide whether an artifact is slop or not.

## EXTENDING THE METAPHORS

Several metaphors have emerged for making sense of AI slop and communicating its stakes to policymakers and the public, each illuminating different dimensions of its economics, incentives, environmental effects, and cultural dynamics—and each with its own limitations. The most prominent metaphors are explored in Figure 5.

**Figure 5: Metaphors for AI slop**

Metaphor	What It Illuminates
<b>Plastic Pollution</b>	Like plastic, AI slop is durable, inexpensive to produce, and presumed to be difficult to remove once it spreads across the ecosystem. Individual pieces of plastic (slop) may seem trivial, but the cumulative effects degrade the informational environment over time. The harm that emerges is less from a single artifact and more from the systemic accumulation and challenges of “cleanup.” The proliferation of microplastics in brain tissue <sup>26</sup> mirrors the flooding of attention feeds with AI slop; both are forms of invisible, ambient contamination that degrade the systems meant to process what matters.
<b>Spam</b>	AI slop echoes discussions from the spam domain: Slop is cheap to produce but costly for platforms and users to filter and manage at scale, and the impacts can vary. The imbalance between production costs and moderation costs creates an economic imbalance.
<b>Tobacco</b>	The tobacco analogy highlights how institutions normalize products that are profitable, are considered addictive, and are culturally embedded even as their supposed harms remain initially disputed or poorly understood. AI slop similarly benefits from economic incentives and platform engagement dynamics that obscure its long-term societal effects. As with tobacco, debates may evolve from questions of individual consumption toward broader concerns about systemic public health and environmental impact. In March 2026, a California jury found Meta and YouTube negligent in designing platforms to “addict” young users, a verdict some legal experts have compared to the fall of Big Tobacco, considered one of the most consequential corporate liability shifts in modern American history. <sup>27</sup>
<b>Junk Food / Fast Food</b>	This analogy reflects both the supply side and the demand side. Fast food is always available, is relatively cheap, is authentically liked by many, can serve social functions (warming shelter, community center, nutrition access, jobs provider), and operates as infrastructure allowing other parts of communities and economies to function in the absence of a better alternative. AI slop has many similar qualities.
<b>Multilevel Marketing Schemes (MLMs)</b>	AI slop resembles a multilevel marketing scheme for attention. Just as MLMs recruit participants to sell (usually low-quality) products to extract value from downstream recruits, the slop ecosystem recruits participants (human or AI) to produce and circulate low-quality content to extract value from algorithmic distribution and advertising systems. Slop’s value thus does not come from the content itself, but rather from how it is situated within a distribution system, and it is the platforms that extract value from the attention economy.

## SLOP'S PUNK ROCK ETHOS

Despite having shared characteristics, slop can also vary substantially: Enter Italian Brainrot. *Brainrot* is a popular online term used to describe the effects of consuming large volumes of low-value digital content, frequently associated with perpetual scrolling and meme culture. While it is often used humorously or self-deprecatingly, the term captures broader cultural anxieties about attention, overstimulation, and the design of online media ecosystems. It is also used in slop-inspired games popularized on Roblox.<sup>28</sup> *Brainrot* is also used to label an aesthetic category of avant-garde and often purposefully discordant memes, made with or without the use of AI tools.

Italian Brainrot is a specific manifestation of this phenomenon. Spreading quickly on TikTok through remixing and reposting, Italian Brainrot is made up of three key components: an anthropomorphized object or animal such as a ballerina with a cappuccino cup head, a nonsensical Italian-sounding name (“Ballerina Cappuccina”), and an audio narration performed by a stereotypical-sounding Italian AI voice. Its post-ironic humor and canny, crowdsourced parody of AI slop aesthetics is what made it so popular in early 2025.

Absurd and rapidly created, Italian Brainrot demonstrates some of slop’s better qualities: It is participatory, democratized, and ironic, and its contrast with other kinds of slop serves as a critique of algorithms and their platforms.<sup>29</sup> In an interview held with creator and artist Fabian Mosele to help shape the Slop Salon and prepare these proceedings, Mosele conceptualized slop as a genre that functioned as a means of reclaiming what felt inevitable as a result of the ubiquity of AI-generated content, tools, and optimization. This counterculture and almost punk rock ethos, for Mosele, represents a rejection of top-down imposition of aesthetic norms and technocratic solutions. Similarly, Mosele shared their perspective on creating characters and franchises outside dominant corporate intellectual property as a way to circumvent corporate control over creative agency.



Image: Examples of popular Italian Brainrot characters

This view complicates negative characterizations of AI slop as digital trash. As Cody Kommers et al. caution,<sup>30</sup> dismissing slop outright risks obscuring its social and cultural functions, and foreclosing avenues for meaningful study. AI slop can act as a supply-side response to persistent demand for content—explanations, entertainment, commentary, and expression—at a scale that human labor alone cannot meet.

These definitional debates also shape how we think about consequences. Understanding slop's impacts requires recognizing that the same artifact can appear as low-quality content to one observer and as valued expression to another; graffiti is an applicable metaphor for this on a content level.

---

## WHERE SLOP MATTERS: MAPPING IMPACTS

If defining slop is complicated, assessing its impact is even more complex. Discussions of AI slop often move quickly toward claims of harm, yet what constitutes harm is itself shaped by underlying assumptions about quality, expertise, and cultural value. A comprehensive understanding of impact requires considering both its positive and negative dimensions, examining them through different analytical frames and at various scales.

### INDIVIDUAL-LEVEL IMPACTS

At the social and individual level, concerns about AI slop frequently focus on its impact on cognition and attention, and what form this takes across different demographics. Critics argue that generative systems increasingly flood digital environments with engaging but ultimately low-substance content, capturing attention while offering limited informational value or depth.<sup>31</sup> Emerging research on generative AI use similarly suggests that heavy reliance on automated writing and summarization tools may weaken critical reasoning or independent problem-solving skills, particularly when users substitute AI output for reflection rather than using it as a scaffold.<sup>32</sup> These concerns intersect with broader attention economy dynamics: In environments already saturated with stimuli, the addition of high-volume synthetic material may increase cognitive load and reduce the perceived value of information signals. These risks are especially salient for younger users, whose developmental environments are shaped by algorithmically recommended content streams across platforms and gaming ecosystems. Studies have already demonstrated that extensive social media use can be linked to diminished cognitive load,<sup>33</sup> and that this phenomenon is pernicious among younger users.<sup>34</sup>

#### AI SLOP AND KIDS' CONTENT

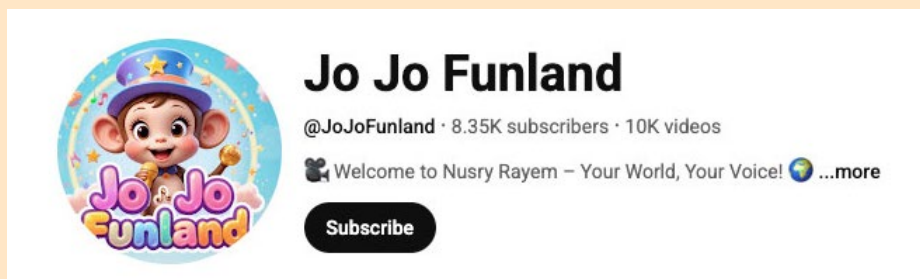


Image: The YouTube profile of JoJo Funland, a channel producing AI slop aimed at children<sup>35</sup>

### **AI SLOP AND KIDS' CONTENT (continued)**

Children's content has become one of the most visible frontiers of the AI slop economy. Searches for popular and trusted terms like the ever-present "Cocomelon" or "Peppa Pig" increasingly surface algorithmically generated knockoffs filled with uncanny animations, distorted character voices, and nonsensical storylines, produced at high volume by anonymous channels chasing autoplay revenue. Jo Jo Funland, one such channel, has generated up to 50 new videos a day and accumulated millions of cumulative views.<sup>36</sup>

The promotion of this content appears to be structural. A February 2026 *New York Times* investigation analyzed over 1,000 YouTube Shorts recommended to young children and found that the algorithm was heavily pushing AI-generated content explicitly targeted at toddlers and preschoolers.<sup>37</sup> Because young children rarely type queries themselves, the content is engineered to exploit recommendation systems and reach kids through autoplay rather than active search. Subjecting children to repetitive, low-quality content of this kind may carry real developmental costs, including hyperactivity and immature cognitive processing, and pediatric researchers have begun raising alarms about what the proliferation of slop on YouTube means for early childhood development.<sup>38</sup>

Although YouTube has recently taken steps to reduce spammy, inauthentic content on the platform, moderation at this scale poses difficulties. Creators are asked only to voluntarily label "altered and synthetic content," and an investigation by the advocacy group Fairplay found that this labeling system does not extend to YouTube Kids at all.<sup>39</sup> Provider-side watermarks, of the kind created automatically by the Sora application when used by the platform, can easily be covered up.

Civil society and the courts are now applying pressure on multiple fronts. In April 2026, more than 200 child advocacy groups and experts—coordinated by Fairplay and including the American Federation of Teachers, the American Counseling Association, and researchers such as Jonathan Haidt—sent an open letter to YouTube CEO Neal Mohan and Google CEO Sundar Pichai demanding that AI slop be barred from YouTube Kids and that Google halt investments in AI children's content studios.<sup>40</sup>

Because March 2026 court decisions left Meta and Google liable for design choices deemed harmful to their users,<sup>41</sup> YouTube may continue to face questions on its handling of this content and its impacts on children.

## **INSTITUTIONAL-LEVEL IMPACTS: SLOP’S EFFECTS ACROSS SCIENCE, ACADEMIA, AND BUSINESS KNOWLEDGE INFRASTRUCTURES**

Concerns about slop extend to collective informational environments. Scholars have warned that generative systems could contribute to the gradual degradation of knowledge infrastructures, from online discourse to scientific communication, by introducing plausible but weakly grounded material into circulation.<sup>42</sup> Similar anxieties appear in discussions of institutional writing and research production, where AI-generated summaries, reports, or administrative text may accumulate faster than they can be meaningfully evaluated.<sup>43</sup>

Businesses are not immune to these dynamics. Emerging discussions of “workslop” suggest that AI-generated workplace content can create new forms of organizational noise, making it more difficult to distinguish actionable insight from filler content and contributing to wasted time, decision fatigue, and lower-quality communication between coworkers.<sup>44</sup> It is worth pointing out that AI-generated content created by employees at work is not algorithmically amplified and thus does not meet the definitional scope of *slop* the proceedings mostly focused upon; workslop, however, is worth mentioning as a phenomenon in this overall debate. Slop-like data has reportedly emerged in unexpected places, such as municipal data sources, as well.<sup>45</sup> These developments point toward a broader pattern in which generative abundance complicates information filtering across domains rather than producing isolated failures; this is an area warranting additional focus and research.

## **POLITICAL-LEVEL IMPACTS: SLOPAGANDA AND AI SLOP’S INTERSECTION WITH POLITICAL DOMAINS**

The political domain introduces additional stakes. Scholars argue that generative content has contributed to shifts in political communication by amplifying performative, emotionally resonant, or meme-driven messaging styles, potentially reinforcing populist narratives and reshaping expectations about authority and expertise.<sup>46</sup> AI slop may thus not simply distort information but also alter the communication norms through which legitimacy and persuasion operate, representing a new political frontier for propaganda. Journalists have also framed the increasing phenomenon of “political memefare” as the latest terrain for strategic communications and public diplomacy, citing AI-generated content such as the Donald Trump administration’s use of AI slop videos and imagery to depict detained immigrants as cartoon characters,<sup>47</sup> its creation of satirical videos of the administration dropping feces on protesters,<sup>48</sup> depictions of real estate development and a FIFA soccer field in Gaza,<sup>49</sup> and “nostalgia slop” that harks back to a more “traditional,” Norman Rockwell–style past.<sup>50</sup> This phenomenon is not unique to the Trump administration; an AI-generated video circulated on Chinese social media sites depicted obese American workers struggling on manufacturing assembly lines.<sup>51</sup>

The discourse around slop’s role in politics has intensified since the February 2026 war between Israel, the United States, and Iran. Official Iranian diplomatic social media accounts<sup>52</sup> have posted AI-generated videos mocking President Donald Trump and Prime Minister Benjamin Netanyahu in

registers borrowed directly from meme culture and commercial AI content, including Trump being punched by Jesus Christ and falling into a fiery pit; Trump depicted as the comic-book character Venom; and AI-generated images of US astronauts burying the Jeffrey Epstein files on the moon.

What these cases share, and what distinguishes them from earlier aesthetics of political misinformation, is a register that recent commentary has begun to describe as *slopaganda*. Slopaganda is political communications that inherit the visual language, production economics, and uncanny aesthetics of commercial AI slop, working through remix, absurdity, and cultural references rather than straightforward claims of fact. The March and April 2026 Iranian-created slopaganda videos were intentionally created, targeted at an American audience, and meant to provoke without being straightforwardly deceiving.<sup>53</sup>

Two features serve to make slopaganda more difficult to govern when using conventional content moderation frameworks. First, cartoon content and stylized content may be less likely to be flagged or removed by platform moderation systems than realistic footage, and may function as a policy or moderation bypass whether or not it was intended as one.<sup>54</sup> Platforms rarely regulate “propaganda” as a distinct content category, focusing instead on such behaviors as coordinated inauthentic activity, undisclosed state sponsorship, and incitement. Second, because this content type mirrors the model used by commercial slop infrastructure, it is increasingly indistinguishable from other forms of slop. State actors may weaponize this ambiguity in the future to disguise the provenance of slopaganda.

### **SYNTHETIC DISSENT: A NEW FORM OF POLITICAL EXPRESSION**

Slopaganda is not the only form of AI-generated political content with consequences; in other contexts, synthetic dissent shows how the same tools serve expression. Content that is quickly written off as slop can serve an important expressive function, especially when it enables political satire and dissent.

Recent examples show that generative AI can give people more room to express political criticism where speaking up openly is risky, while also helping to obscure authorship under repressive regimes. In Belarus, for example, the opposition created “Yas Gaspadar,”<sup>55</sup> an AI-generated virtual candidate chatbot, ahead of tightly controlled elections. Although the AI-generated candidate could interact with voters, it also functioned as a satirical political statement. An opposition figure joked that the AI candidate could speak freely without risking arrest.<sup>56</sup>

Türkiye and Venezuela offer other examples of how synthetic media can help people express dissent in repressive contexts. In March 2025, Turkish authorities jailed opposition leader Ekrem İmamoğlu. This led to the country’s largest protest wave in a decade.<sup>57</sup> While he was in prison,

## **SYNTHETIC DISSENT: A NEW FORM OF POLITICAL EXPRESSION (continued)**

his party used AI to generate a speech in his voice and broadcast it to thousands of demonstrators. Younger activists also turned to AI image tools to produce satire targeting the Recep Tayyip Erdoğan government. Among the viral images that circulated were AI-generated depictions of riot police arresting fictional characters like Pikachu.<sup>58</sup> Some images suggested that in repressive environments, even a cartoon character could be detained. In 2024, following Nicolás Maduro's disputed reelection, media organizations launched *Venezuela Retweets*, a news show with synthetic anchors La Chama ("The Girl") and El Pana ("The Dude").<sup>59</sup> The AI anchors delivered reporting produced by human journalists and helped shield them from the government's crackdown.

A recent commentary calls these AI-mediated forms of expression "synthetic dissidents."<sup>60</sup> The term refers to AI-generated personas and chatbots that bring dissenting speech into the public eye while making the people behind them harder to identify and punish. But long before generative AI, dissidents relied on low-cost, attribution-obscuring forms of expression, such as graffiti, samizdat, and coded political memes, to criticize power when direct speech was dangerous. Satire and artistic expression have been central to that tradition, especially in closed societies.

International human rights law provides broad protections for artistic expression, including satire and other uses of art to challenge people and institutions. Under Article 19 of the Universal Declaration of Human Rights, "Everyone has the right to freedom of opinion and expression; this right includes freedom to hold opinions without interference and to seek, receive and impart information and ideas through any media and regardless of frontiers."<sup>61</sup> The same protection appears in Article 19 of the International Covenant on Civil and Political Rights, which explicitly covers expression "in the form of art" and through "any other media."<sup>62</sup>

Generative AI allows people to exercise their freedom of expression rights in novel ways, including by empowering satire and artistic expression. Humor is one of the most powerful weapons in the fight against authoritarianism. Throughout history, satire and parody have served as a powerful form of defiance against authoritarian regimes around the world.<sup>63</sup> The internet redefined satire by providing a platform for creators to produce and disseminate satirical content, amplifying their voices and reach, and generative AI is further expanding the conversation about inequality, injustice, and other power dynamics.

As we think collectively about the risks associated with AI slop, we should not lose sight of its expressive value. In certain cases, it can serve as a vehicle for satire and political criticism that would otherwise be harder, or more dangerous, to express. Efforts to govern AI-generated content should take care not to suppress those forms of dissent in the process.

*Written by Sarah Shirazyan*

## SLOP'S ROLE IN THE SCAM ECOSYSTEM

Slop often crosses over into fraud and scams, particularly when it is used as a vector to take advantage of online populations with limited digital literacy. Researchers have documented how AI-generated content can function as an upstream layer in scam ecosystems, serving to attract attention, build audiences, and create a veneer of legitimacy before directing users toward deceptive schemes. Analysis provided by Renée DiResta and Jason Koebler as part of their workshop provocations and their related published work shows that networks of social media pages already use AI-generated imagery and text to generate engagement at scale, often routing users toward monetized spam, data-harvesting sites, or low-quality commercial funnels.<sup>64</sup> In this context, slop does not always operate as the harmful end point itself, but rather as the connective tissue that sustains exploitative infrastructures.

DiResta expanded on this concept during the convening through her provocation, arguing that an important function of some slop is to provide cover for adversarial actors. The sheer volume of AI-generated content serves to normalize synthetic material in the information ecosystem, acting as camouflage inside which targeted manipulation can hide. DiResta referenced a January 2026 incident to illustrate this: The White House posted a digitally altered photograph of civil rights attorney Nekima Levy Armstrong being arrested at an anti-ICE protest, manipulating her face to appear tearful and darkening her skin, with no watermark, label, or indication of editing.<sup>65</sup> The framing worked in part because of a year of cartoonish AI-generated political content that allowed White House officials to dismiss the post as only a meme instead of anything more consequential.

DiResta's perspective is that slop is not just a symptom of a degraded information environment but is itself also a resource or tool for actors who may have more nefarious intentions. Slop aesthetics can act as camouflage, increase the costs of verification, and provide support for and reinforce cheap, reusable content distribution systems designed to maximize clicks, views, and advertising revenue. These nuances get missed when content is described simply as "low quality."

DiResta's provocation further distinguished among four categories of AI slop, defined by two characteristics (Figure 6): *specificity* (is the content generic or aimed at a real person, institution, place, or event?) and *linkage to action* (does the content end at shaping attention, or does it route the viewer toward a tangible next step, such as to click on a link, pay something, install software, submit information, harass someone, or show up somewhere?). This builds on Mantzarlis's "expressive vs. deceptive / political vs. economic" taxonomy, which similarly looks past surface characteristics to focus on intent. DiResta's framework adds a further dimension, however, moving beyond topic and intent toward *purpose* and *operational function*: what the slop is built to do.

**Figure 6: Renée DiResta’s proposed framework for distinguishing among different types of AIGC, particularly relating to scams and fraud, by specificity and action linkage**

Specificity	Action	
	No Action Linkage	Action Linkage
<b>Generic</b>	<p><b>Kitsch and Filler Content</b></p> <p>Ambient feed pollution: synthetic filler optimized for engagement; no referent, and no next steps. Example: Shrimp Jesus.</p>	<p><b>Conversion Spam</b></p> <p>Generic content routed to funnels: affiliate links, phishing, installs, and lead generation. Automated posting, A/B testing, and link rotation drive scale.</p>
<b>Referent-Specific</b>	<p><b>Targeted Narrative Shaping</b></p> <p>Insinuation about a real person or event with “just a meme” deniability. Often bot-amplified to simulate consensus or salience.</p>	<p><b>Weaponized Persuasion and Predation</b></p> <p>Operational abuse: impersonation, extortion, tailored harassment, false procedural instructions. Example: Deepfakes. Bots and coordinated networks can provide scale and speed.</p>

DiResta also provided insights into how to apply known trust and safety frameworks to slop, particularly the Actors, Behaviors, Content framework.<sup>66</sup> Researchers in the adversarial influence operations field have used these three vectors as analytical lenses, and DiResta’s work demonstrated how to apply this framework to AI slop:

- **Actors: Diagnostic by presence or absence.** For content generated by content farms, actor identity is largely irrelevant because the logic is pure economic arbitrage, and identifying who runs the operation may matter little to the overall threat assessment. For deceptive slop like fake reviews, actor identity begins to matter: Whether reviews are generated by the brand itself, a hired farm, or a competitor changes both enforcement targets and intervention rationale. For state-linked influence operations (such as those using coordinated inauthentic behavior), attribution becomes critical for interventions, and *slop* becomes an unhelpful label. This suggests a diagnostic heuristic: If actor identity changes the threat assessment, one may no longer be dealing with slop. Actor relevance itself signals operational context.
- **Behaviors: Worthy intersections of analysis.** Behavioral signals such as volume, automation patterns in distribution, evasion techniques, platform manipulation, and coordination can help distinguish what is considered harmful from what is considered merely annoying. A single fake review can be irrelevant; the appearance of 10,000 constitutes a different problem. For DiResta, most genuinely novel adversarial developments enabled by AI are behavioral, not content-level, with examples being cheap consensus simulation via reply swarms or the ability to solve CAPTCHAs while flooding public comment systems. The content may be mediocre, but the deployment behavior determines its impact.

- **Content: Where everyone looks first.** Content-level analysis identifies deception mechanisms, for example, whether a video impersonates a journalist or a fictional image claims to depict something real. But focusing on creation modality only risks misplacing the emphasis on authenticity, rather than the context of the content. A synthetic image and a decontextualized real photograph could perform identical functions in campaigns; likewise, a deepfake and a cheapfake (a deepfake made from inexpensive AI tools) could produce equivalent audience effects. As synthetic material saturates digital environments, “made with AI” may become an increasingly uninformative risk signal.

## SYSTEMIC-LEVEL IMPACTS

Cutting across these domain-specific impacts, salon participants worked toward a structural distinction between first-order harms and second-order harms. First-order harms are direct, observable, and typically traceable to a specific piece of content: AIGC that facilitates harassment, scams, NCII, and CSAM, alongside low-effort, high-volume content optimized for algorithmic reach rather than informational value. Second-order harms surface at the level of the ecosystem, with diffuse impacts that emerge from accumulation rather than from any single piece of content. These types of impacts include epistemic degradation, labor displacement, and effects on youth development. These effects may be further compounded by the Ouroboros effect, wherein generative models trained on existing low-quality outputs subsequently produce more of it, contaminating future training datasets through synthetic data feedback loops.<sup>67</sup> This represents a form of “epistemic collapse” that could reshape the trajectory of knowledge production.

### HOW WILL WE KNOW WHICH HISTORY IS REAL?

Among the more insidious frontiers of the AI slop economy is the takeover of historical content on YouTube. A September 2025 investigation by 404 Media documented a sprawling ecosystem of AI-generated channels.<sup>68</sup> These channels have been pumping out three-hour-plus narrated videos on topics such as medieval peasant life and the American frontier. One Sleepless Historian video<sup>69</sup> on medieval peasants surviving cold nights had accumulated 2.3 million views before journalist Jason Koebler caught its AI narrator glitching in mid-sentence: *“In the end, Anne Boleyn won a kind of immortality... FEEEEEEEEEE.”*

Although AI hiccups like this give critics humorous examples to point to, the proliferation of such content can carry real consequences. In the most obvious sense, videos like Sleepless Historian’s are replacing work that amateur historians and content creators have traditionally spent weeks

### **HOW WILL WE KNOW WHICH HISTORY IS REAL? (continued)**

scripting, editing, and—crucially—fact-checking. Pete Kelly, who runs the popular *History Time* channel, told 404 Media: “It used to be enough to spend your entire life researching, writing, narrating, editing ... but now someone can come along and they can do the same thing in a day instead of it taking six months, and the videos are not accurate.”<sup>70</sup> The French Whisperer, another long-running history creator, reported a roughly 60 percent drop in views over the past year, which he attributes to the AI slop surge.<sup>71</sup>

Setting aside outright displacement, the slop tax on viewers is the loss of recognizability. AI history videos are engineered to mimic the aesthetic conventions of legitimate documentary work, and less discerning viewers than Koebler will likely lose themselves in their tight narrative and polished prose. As this kind of content floods YouTube, a student searching, for example, “history of Pompeii” will have an increasingly hard time distinguishing carefully crafted scholarship from AI slop, either mistaking one for the other, or spending considerably more time than before finding a reliable source to draw on.

But the harms go beyond directly impacting creators and consumers. The deeper concern is that AI-generated historical content does not merely simplify—it actively encodes outdated and incorrect understandings. A peer-reviewed study published in *Advances in Archaeological Practice*<sup>72</sup> used Neanderthals as a test case to measure the gap between AI-generated content and current scholarship. Their findings captured the biases that ChatGPT repeated, for example, framing outdated research as still relevant. A world filled with AI history risks genericizing real historical events, removing key details, carefully chosen examples, and revelatory biases. History, as the oft-cited phrase says, is written by the victors—and AI writing and narration will only further obscure uncommon perspectives.

Two cross-cutting issues framed this discussion and remain open in the literature. The first involves the crowding-out effect, which refers to the phenomenon of what changes or gets lost when slop displaces other content in algorithmic ranking, or distorts the relative signal of high-quality content by sheer volume. The second involves measurement. Second-order harms are diffuse by definition, and platform transparency limitations constrain the independent research needed to establish baselines or detect inflection points.

These perspectives suggest that when one is considering impact and harm, AI slop is best understood not as a singular content phenomenon that causes harm, but as a set of distributed pressures on information ecosystems. Its impacts can range from individual cognitive effects to institutional filtering challenges, economic redistribution, and shifts in communication norms. Rather than producing

dramatic failures, slop may instead operate through accumulation: gradually reshaping how knowledge, authority, and value are perceived across cultural, scientific, and political landscapes.

At the same time, evidence for systemic degeneration remains uneven and difficult to measure. Some nascent empirical work suggests that AIGC is not yet as dominant in certain sectors as public discourse implies. Similar research on its impact in the 2024 global elections cycle suggested that generative AI was largely not leveraged for large scale misinformation occurring on platforms.<sup>73</sup>

The cultural dimension muddies a purely negative framing. Research on audience reception suggests that some forms of generative content are not only tolerated but actively embraced, with users interpreting artifacts like Spider Jesus (a variant of the aforementioned Shrimp Jesus imagery) as humorous, ironic, or culturally meaningful.<sup>74</sup> Similarly, artists and creators have begun experimenting with generative excess in explicitly subversive ways, using rapid production, absurdity, and algorithmic aesthetics to critique the very systems that enable them.<sup>75</sup> Such practices reflect a shift identified in recent scholarship: Generative tools lower barriers to cultural participation, enabling more people to produce and circulate creative work.<sup>76</sup>

Although this strand of experimentation remains underexplored in the literature, it suggests that slop may not function solely as a site of cultural negotiation or, at the other extreme, a symptom of cultural decline. The reality lies likely somewhere in between, and is complicated due to independent researchers often lacking the tools or data required to adequately capture discourse, polarization, or content degradation patterns or prevalence on platforms. As has been well-documented, most notably with the loss of CrowdTangle (a public insights tool acquired by Meta that enabled research on public content on Facebook that was deprecated in August 2024), platforms today maintain tight control over data and algorithm access.<sup>77</sup> This asymmetry makes it difficult to determine the extent of any decline in information quality, and whether these changes are attributable to content and medium, or other factors such as algorithmic design. The absence of shared measurement tools also means that platforms themselves remain the primary arbiters of what counts as degradation, amplification, or harm, shaping both the research agenda and the public understanding of the phenomenon.

Asking what slop is and what harm it causes leads naturally to a second set of questions on why slop is produced at such scale, and what conditions make that production economically rational.

## SLOP MUSIC: WHEN THE ALGORITHM ATE THE CHARTS



Image: The cover of “No More You” by Eddie Dalton, an AI-created musician whose songs topped the iTunes charts in April 2026<sup>78</sup>

Until recently, platforms could police waves of low-effort audio content with straightforward spam policies; spammy music was duplicative, was easy to spot, and rarely found an audience. Generative AI changed this, as unlike earlier waves of automated content, current generative audio tools like Suno and Lyria generate songs that increasingly resemble human-made recordings, leaving platforms without reliable tools to differentiate authentic content from low-effort synthetic output. Platforms struggle to maintain quality control, and the competitive stakes for human artists are clear, as every songwriter patiently honing their craft must now compete with users who prompt-engineer 10 new albums a week.

Is this simply democratization, encouraging competition and opportunity? Artists have long seized on new technology in order to create new sounds. Popular music tools, such as Apple’s Logic Pro, have for years used machine learning to create drum patterns that sound real and spare the cost of a session percussionist.

For those who have observed the rise of AI-generated songs, the key question is not whether prompt-based music is truly *music*, but rather, what the impact will be on the delicate system of incentives that constitutes today’s creator economy. The volume of new music being created is astonishing, averaging over 100,000 uploads per day, according to one source.<sup>79</sup> Will human creators see their slice of the pie shrink? Not because their music is inferior, but rather because they have focused on their craft instead of studying algorithms and search engine optimization

### **SLOP MUSIC: WHEN THE ALGORITHM ATE THE CHARTS (continued)**

tactics? If most of these tracks represent hasty low-effort attempts to strike it lucky in the algorithm, have we democratized music, or simply democratized spam?

Today's streaming platforms are ultimately technology companies, and may feel the pressure to encourage AI adoption even as it undercuts the human authenticity that makes their services so popular. If a platform can shift consumption toward "platform-owned" AI audio, for which they won't need to pay royalties, their margins will improve. How can we ensure that platforms are not forced to compete with one another by downranking pricier non-synthetic music?

As Universal Music Group CEO Lucian Grainge put it: "Validating business models that fail to respect artists' work and creativity, and promote the exponential growth of AI slop on streaming platforms, is a grave disservice to artists, songwriters, and all of us who work in music."<sup>80</sup> If the volume of AI slop music proves difficult to control, human creators will find themselves struggling to compete on an increasingly uneven playing field, working harder for shrinking revenues, while the music industry becomes less human by the day.

*Written by Shane Culloty*

---

## WHY SLOP EXISTS: SUPPLY, DEMAND, AND INCENTIVES

### AI SLOP AS A SUPPLY-SIDE RESPONSE

If slop is not just a content category but a systemic outcome, the next question is why it is produced at such scale. The proceedings examined the proliferation of AI slop as a supply-side response to structural conditions in digital information economies. Some research suggests that interventions that are focused solely on content removal or verification often miss the deeper drivers of demand, such as emotional salience, identity reinforcement, identity distillation, and platform engagement incentives.<sup>81</sup> As societal mechanisms push individuals to seek out non-mainstream theories or find like-minded voices, misinformation can offer a supply-side solution to this instinct.<sup>82</sup>

This perspective has increasingly been applied to generative media, including political deepfakes and synthetic campaign content. In this context, scholars argue that demand for emotionally resonant or attention-grabbing material often precedes and shapes the production of misleading or low-quality content.<sup>83</sup> Generative tools do not create demand for slop so much as respond to it, supplying inexpensive, scalable content that fits existing appetites within digital ecosystems. Slop can also provide content that is more heavily personalized and abundant than was previously possible.<sup>84</sup>

Furthermore, generative tools may expand opportunities for expression, publication, or communication in certain contexts. Researchers have noted that generative AI tools can expand access to editorial support, translation, and native-language assistance that has historically been costly or inaccessible to many writers and creators.<sup>85</sup> From this perspective, AIGC can be understood as a supply-side response to the demand from academics and content creators to tap into the English-speaking information ecosystem.

One of the authors of this paper, Jason Koebler, contributed a provocation that highlighted his investigations into the economic drivers underpinning AI slop. Sharing examples from his reporting, Koebler suggested that much AI slop content is produced not for deception, but for monetization, fueled by low-cost courses teaching viral content strategies for industries including pornography creators and social media influencers. From Koebler's viewpoint, the proliferation of freely available educational resources on how to monetize content contributes to the industrialization of content production at scale. Others in the media have also highlighted this dynamic, with implications for political discourse. A recent story covered the phenomenon of a medical school student creating Make America Great Again (MAGA)-focused AI influencer content of attractive young women to make thousands of dollars.<sup>86</sup> In another example, a Burkina Faso-based teen created an AI-generated synthetic video of a coup in France that did not actually occur. This synthetic video reached 12 million Facebook views and gained attention from the office of the president of France due to the deceptive nature of the content, particularly after Meta refused to take the video down.<sup>87</sup>

### **THE €7 (\$8.13)<sup>88</sup> PARIS COUP, A LOUD DEBUT IN SLOPREPRENEURSHIP**

The Burkinabè teen behind the December 2025 AI-generated video of a French coup that did not occur in real life, which racked up over 12 million views on Facebook and prompted an unsuccessful intervention by French President Emmanuel Macron after Meta refused to remove it, is illuminating not for the viral artifact itself, but for what bookends it.<sup>89</sup>

Before the coup video, the teenager was reportedly not a political content creator, nor did he seem to care much about Macron, France, or politics in general. The 17-year-old high school student in Ouagadougou had spent late 2025 throwing AI-generated content at the algorithm to see what would stick: a man transforming into a snake, rain of whales, lions attacking pedestrians in Paris, a dragon piercing the clouds. He had tried motivational videos first, which were unsuccessful. His prior biggest payout (around €10, or \$11.60 as of this writing<sup>90</sup>) came from a clip of a man putting his head into a futuristic haircutting machine. He taught himself video generation through YouTube tutorials and produced everything on his smartphone using Sora 2. Only in early December 2025 did he pivot to fake breaking news formats, beginning with a fabricated coup attempt in Benin, before extending the template to France. The fabricated French coup AI-generated video earned him €7, or approximately \$8.13 as of this writing.<sup>91</sup>

After the video went viral and his Facebook page surged past 279,000 followers (boosted in part by Macron's own denunciation, which amplified it), the teenager quickly converted the notoriety into a coaching business, advertising AI video creation training at 7,000 CFA francs per hour (roughly €10) via WhatsApp. He proudly told the Agence France-Presse, a global news agency, perhaps strategically using the interview itself as a marketing opportunity, that at least five clients had contacted him in the week after the video blew up. One of his earlier (fake) Radio France Internationale-branded inauthentic coup videos had even featured an AI anchor breaking the fourth wall to advertise the WhatsApp number directly. The fake coup was already doubling as an ad for the sloprepreneurship coaching business.

Although some platforms extended monetization opportunities to selected countries on the African continent in 2024, namely Nigeria and Ghana,<sup>92</sup> dozens of countries have long been excluded from the creator economy. As a result, creators focused on creating this content in excluded countries, like the Burkina Faso-based teen behind the viral AI-generated video depicting a coup that did not actually occur, explicitly recommend using a virtual private network (a tool that routes internet traffic through a private encrypted connection and can mask a user's location) to create accounts from Europe or North America.<sup>93</sup>

The path to slop creation is simple: Content creators looking for side gigs or ways to monetize social media algorithms find a culturally important or sensitive topic (which often can be political or socio-cultural) and then use cheaply generated AI content to plug the gap and make money. The primary motivator for most slop creation may not be the desire to spread political messaging, to exacerbate wedge issues, or to otherwise shape the information environment, but it can serve those purposes if it is sufficiently lucrative. Platforms (e.g., YouTube, Meta, or TikTok) generate revenue in part through advertising, which means that content generating a lot of views, clicks, watch time, and engagement can produce revenue even if the content itself is low quality. As discussed, generative AI dramatically lowers the cost and effort required to create large volumes of content, creating a kind of arbitrage. (Economic arbitrage is the practice of exploiting a mismatch between cost and value in a system to generate profit or advantage.<sup>94</sup>) In a slop use case, producers can cheaply mass produce attention-grabbing AI content and still earn advertising revenue if enough users view it. This dynamic can incentivize the large-scale production of slop because the cost of generating the content can be much lower than the potential revenue or traffic it produces. Over time, creators, marketers, and content farms may then optimize for whatever types of AI-generated content perform well within platform recommendation systems. The same production and distribution systems used for entertainment-oriented or spam-like slop can also generate political content, which does not always require deliberate political manipulation. As creators respond to audience demand, engagement incentives, trending topics, and algorithmic rewards, the ecosystem may naturally produce politically themed, polarizing, misleading, or emotionally provocative AI-generated content because that content attracts attention and performs well within engagement-driven systems.

YouTube cost per mille<sup>95</sup> (CPM) rates reportedly range from about \$1 USD per 1,000 views in Pakistan, to \$5–\$7 per 1,000 views in the US, and €4–5 (or \$2.30–\$4.60 USD) in France; thus, creators producing content calibrated to American or Western European audiences can capture disproportionate returns. Journalist Jean-Marc Manach uncovered 8,300-plus AI-generated fake news websites in French alone (plus over 300 in English, and over 150 in German), run by search engine optimization (SEO) specialists gaming Google Discover (a personalized content feed shown on Google’s mobile app, home page and some Android devices) for advertising revenue. Google Discover is driven by algorithmic ranking systems, and in one instance promoted fabricated stories to millions of mobile users claiming that French banknotes would cease to exist, or that a 25,000-year-old pyramid had been found under a mountain, to millions of mobile users. Content creators can earn thousands of euros or equivalent US dollars per day through Google AdSense.<sup>96</sup>

Koebler’s research also described how “hustle bros” actively recruit and train new entrants to capture value from Western, and especially US, attention markets (the underlying economic systems and associated business models that online platforms use to capture and monetize users’ time and focus<sup>97</sup>) to benefit from their relatively lucrative advertisement ecosystems. This dynamic raises broader concerns about how globalized, incentive-driven content production pipelines are reshaping information ecosystems. The optimization of these monetization pathways could be fueling a decentralized, self-replicating, global content production infrastructure that warrants critical inquiry.

## HOW DEMAND FOR AI SLOP GETS BUILT

Generative tools expand supply; the demand for AI slop is shaped by the infrastructures through which content circulates. Observers across journalism, platform research, and policy analysis have noted that slop often reflects incentive structures rather than purely authorial intent. Monetization programs, engagement metrics, ranking algorithms, and automation pipelines collectively reward scale, speed, and emotional resonance, creating the conditions under which generative abundance becomes rational, rather than pathological. At the same time, large platforms must expend resources dealing with adversarial actors and scaled abuse, with organized groups of spammers and scammers seeking to overwhelm audiences with engagement bait. This cohort tends to reliably create volumes of edgy (but non-violative) content in the hope that some of it will be rewarded by the algorithm, at which point they will focus their resources into that particular content niche.

AI slop can be considered, then, a predictable outcome of optimization within algorithmically mediated attention markets. As one analysis notes, this content often succeeds because it is well suited to the metrics through which platforms measure relevance and engagement.<sup>98</sup> Platforms themselves increasingly experiment with tools that automate or encourage synthetic content production, such as generative ad systems and AI-assisted posting features, suggesting that slop is not merely tolerated but structurally integrated into platform strategies. Meta recently released an AI-generated content feed called Vibes,<sup>99</sup> and in the October 2025 earnings call, Mark Zuckerberg framed<sup>100</sup> AI-generated content as a definitional attribute of the next generation of social media and online engagement (the first two generations, as he explained, being defined by content created by friends and family, and then by human creator and influencer content). xAI, an artificial intelligence company founded in 2023 by Elon Musk focused on developing advanced AI systems and large language models, has been flooded with viral sexualized deepfakes, as Grok, its AI assistant and chatbot developed by xAI and integrated into X (formerly Twitter) allowed the content to be quickly pumped out and its algorithm spread the images across its platform.<sup>101</sup> Sora2 is OpenAI's text-to-video model,<sup>102</sup> which can generate short, highly realistic video clips from written prompts.

Unlike hybrid environments where AI tools are embedded within existing social ecosystems, OpenAI's short-lived Sora stand-alone application positioned generative media as the central mode of creation and consumption, effectively removing the traditional constraints on audiovisual production by allowing content supply to scale independently of human labor, and be micro-personalized based on the consumer's preferences. The Sora application reflected what some commentators have begun to describe as a form of "slop capitalism," in which platforms stumble into automating cultural production itself, integrating generation, distribution, and monetization into a single optimized pipeline.<sup>103</sup>

## THE DEATH OF SORA: WHAT IT TELLS US ABOUT THE SLOP ECONOMY

In March 2026, OpenAI announced it would shut down Sora, its AI video generation platform.<sup>104</sup> After Sora's launch, its worldwide user count peaked at around a million and then collapsed to fewer than 500,000, while the app burned through roughly \$1 million every day.<sup>105</sup> Downloads plunged by nearly 75 percent from their November 2025 peak,<sup>106</sup> and OpenAI pulled the plug six months after release, dissolving a three-year, \$1 billion partnership with Disney in the process.

Sora's collapse complicates the assumption that consumer demand for slop-as-destination—content that is openly and exclusively AI-generated, consumed for its own sake—tracks the platform-side enthusiasm for producing it. Embedded and mixed into human-created feeds, synthetic content appears to thrive; as a stand-alone product, however, it struggles to sustain interest, at least in Sora's case.

Compounding the economic problem, Sora launched into a US market where public sentiment toward AI had already soured sharply. Half of U.S. adults told Pew in June 2025 that the increased use of AI in daily life makes them feel more concerned than excited, with just 10% more excited than concerned.<sup>107</sup> Journalist Brian Merchant has argued that the ubiquity of AI slop is directly fueling the broader AI backlash, including its more violent edges.<sup>108</sup> For most people, the primary encounter with generative AI has been a feed full of uncanny videos, synthetic engagement bait, spam comments, and increasingly AI-mediated search results.

In that environment, a flagship consumer app devoted to AI-generated video could become a visible liability, concentrating anxieties about slop, job displacement, copyright, and environmental harms into a single product experience. Commentators have accordingly characterized the shut-down as “a reality check moment for the makers of AI video tools,”<sup>109</sup> as OpenAI and competitors like ByteDance reweigh the IP and reputational liabilities of consumer-facing generative video. Read this way, killing Sora may have been a reputational retreat as much as a financial one.

The relationship between AI providers and platforms remains under-explored, especially as these previously separate technologies converge. Figure 7 attempts to provide theoretical categories of different interaction types, but further research is needed into the velocity and kind of content produced by platform type as it relates to AI slop.

**Figure 7: Generative systems classified by distribution and attention markets**

Category	Examples	Relationship Between AI and Platform	How Content Flows
<b>Separate Systems</b>	ChatGPT → Instagram, Midjourney → X, Claude → Reddit	AI tools exist outside the platform; users generate content elsewhere and upload manually. It’s worth noting that at the time of writing, many sloprepreneurship tutorials are focused on how to “chain” together multiple AI tools (Claude for background research, ElevenLabs for a voiceover track on a “Faceless YouTube” channel, Kling for video generation, etc.) to automate the slop creation pipeline end-to-end	Human decides what to post; platform only ranks and distributes
<b>Tool-Integrated Platforms</b>	TikTok AI filters, Instagram AI stickers, Canva AI tools	Platform offers AI tools for creation but does not auto-generate feeds	Users create within the platform, but outputs remain user-directed
<b>Hybrid AI-Social Platforms</b>	Grok inside X, Meta AI posting tools, Google AI search summaries	AI generation and distribution share infrastructure; outputs can be rapidly amplified	Creation and ranking are tightly linked; algorithms scale outputs quickly
<b>AI-Native Content Ecosystems</b>	Meta Vibes feed, Sora-style video ecosystems	AI generation is built into the platform’s core content supply or feed logic	Synthetic content becomes a primary or automated component of the ecosystem

For many platforms, the economic incentives are clear: Viral content often translates into advertising revenue, audience growth, and data generation, creating opportunities for both platforms and creators to profit from high-volume production.<sup>110</sup> Research documenting how scammers and content farms leverage AI-generated imagery to build audiences further illustrates how generative tools can be embedded directly into monetization pipelines.<sup>111</sup> AI slop is thus being generated in response to platform incentives. Investigations into Facebook’s recommendation systems, for example, have shown that networks of Facebook Pages posting AI-generated imagery at high frequency are systematically rewarded with visibility and audience growth, which can then be converted into advertising revenue, affiliate marketing income, or traffic to ad-saturated content farms.<sup>112</sup> These pages often originate in regions where the cost of producing content at scale is a fraction of the ad revenue it can earn from Western audiences. Algorithmic demand thus does not merely distribute slop but actively shapes its supply by creating feedback loops in which engagement metrics incentivize its outputs.

A broader interpretation situates AI slop within more long-standing critiques of platform capitalism and extraction economies. Some observers argue that the proliferation of synthetic content reflects a familiar trajectory in which platforms progressively internalize value chains, optimizing for full end-to-end control over production, distribution, and monetization. This logic echoes analyses of “enshittification”<sup>113</sup> or the extraction economy,<sup>114</sup> in which digital platforms gradually shift from serving users and creators to prioritizing extractive monetization strategies once markets are locked in. AI slop can be understood not only as a response to existing incentives, but as a potential precursor to deeper vertical integration, in which platforms increasingly rely on internally generated content to reduce dependence on human creators altogether. Amid this backdrop, different sorts of creators, including artists,<sup>115</sup> authors,<sup>116</sup> editors,<sup>117</sup> animators,<sup>118</sup> musicians,<sup>119</sup> influencers,<sup>120</sup> and others, have varyingly organized and spoken out against the impacts of AI-generated content on their livelihoods and culture more broadly. A UNESCO report on cultural rights from December 2025 echoed concerns about flattening and monoculture, positing that generative AI as driven by commercial interests may restrict cultural diversity, promote homogenization, and amplify mainstream aesthetics while sidelining regional, Indigenous, and non-Western modes of expression.<sup>121</sup>

### **THE CIRCULAR ECONOMY OF TECHNO-DEPENDENCE**

Think of the social media scroll as load-bearing infrastructure. The platforms inundated with AI slop today are essential for most businesses, researchers, and media. For many people, they are also the primary conduits for support, connection, information, and care in society.

Slop’s impact on these systems has been made through its entanglement in a circular economy of harm, where the social cost that technology inflicts becomes the demand for its services. AI has stumbled into a position where its role is to *cause or exacerbate social problems, then offer partial remedies that entrench the underlying dynamic*. As social media platforms become alienating to users and harder for creators to use to make a living, AI slop becomes the tool they use to claim their slice of a shrinking pie. As the education system deteriorates (in part because people use AI to cheat), as people get lonelier (in part because of chatbots substituting for social contact), as jobs are eliminated from the economy (in part by automation), and as income inequality accelerates (turbocharged by AI-generated wealth creation), many people will be displaced. Free or low-cost AI has proven, in all these instances, to be the infrastructure they fall back on.

Stopping this race to the bottom requires reflection on whether technology actually proposes solutions or merely postpones them. Platforms, models, and labs can’t do everything, but each day they are able to do more and more. Slop is a highly visible early problem of AI’s

### **THE CIRCULAR ECONOMY OF TECHNO-DEPENDENCE (continued)**

social-institutional alignment, and it will not be the last. The task of aligning these systems is political in the Aristotelian sense rather than the partisan sense—it is a holistic, rigorous search for a shared “good life.” The only decent and effective way to conduct that search is in the open, with the public authentically involved.

*Written by Aidan Walker*

The proliferation of slop also raises questions of an environmental iceberg: Although visibly posted AI slop signals a potential problem, publicly visible AI-generated content may represent only a small fraction of total synthetic production, since generative systems encourage high-volume iterative creation in which many outputs are never shared, published, or archived. Thus, even without distribution, slop creation carries computational and resource implications. AI content that never appears on any feed or never happens to go “viral” can still contribute to the industry’s carbon and water footprint.<sup>122</sup>

**Each platform type therefore encounters slop not as a uniform content problem but as a modality- and incentive-specific manifestation of optimization logic within its own product architecture.**

---

## WHAT CAN BE DONE: INTERVENTION PATHWAYS

If slop emerges from structural incentives, interventions must operate across multiple layers, including platforms, tools, users, and institutions. The following section does not aim to recommend specific actions or normative solutions. Rather, it sketches a preliminary map of potential leverage points across platforms, creators, and users that may influence how AI slop circulates within digital ecosystems.

### PLATFORM POLICIES, INCENTIVES, AND CONTROLS

One of the most immediate leverage points is platform algorithmic decisions and monetization strategy: how platforms rank, monetize, and surface content. This is different from traditional content policy and trust and safety work, which interacts with content based on the content, actor, or behavior. AI slop is not “just” a content challenge; it represents a complex interplay between platform affordances, user and creator incentives, and consumer preferences, whether these preferences are genuinely organic or more algorithmically curated. The same “low quality” AI-generated video or article may appear identical at the content level, but the implications may differ depending on whether it is created by a teenager experimenting with a tool, a sophisticated monetized clickbait farm optimizing for ad revenue, a coordinated fraud network, or a head of state. The “quality” of the slop is not the anchoring principle here; it’s what the slop enables that counts.

Likewise, the behavior surrounding slop (including engagement farming, impersonation, spam distribution, coordinated amplification, or using slop as a launching point for fraud) often drives impact and perceived harm more than the mere fact that a given piece of content happens to be AIGC. A content-focused policy for slop risks focusing too much on aesthetics or origin, while an actor–behavior lens can help distinguish annoyance from manipulation, clutter from fraud, and experimentation from exploitation. Further, a platform’s content moderation team may attempt to fit different varieties of slop into its existing content verticals, but this risks splitting up a common problem across different functions and teams, delaying coordinated action and duplicating effort. Applied to AI slop, a framework that focuses on the function of slop along with what it enables could encourage more aligned interventions across production, distribution, and enforcement layers.

AI slop can also be manifested differently across platform types as it is shaped by distinct incentive structures and modes of distribution:

- On large-scale social media platforms like Facebook or X, slop often appears as engagement-optimized content such as memes, hyper-emotional imagery, AI-generated political cartoons, or spammy influencer posts designed to trigger amplification within algorithmic feeds.
- Search and discovery platforms (e.g., search engines, Pinterest-style inspiration sites) may face different challenges: Slop degrades utility by flooding results with generic, low-value synthetic content, reducing trust and usefulness over time.

- Messaging platforms and closed groups can function as rapid-fire circulation channels, where slop is remixed and repurposed within interpretive communities.
- Creator platforms and monetized content ecosystems may incentivize volume over quality, encouraging automation at scale.
- Productivity and knowledge platforms may face subtler forms of slop, as AI-generated summaries and white papers risk introducing synthetic redundancy into professional discourse.

Each platform type therefore encounters slop not as a uniform content problem but as a modality- and incentive-specific manifestation of optimization logic within its own product architecture.

### **EXPLORING PLATFORMS' DIVERGENT POLICIES ON SLOP**

Recent policy shifts on YouTube illustrate how relatively targeted adjustments can influence the economics of generative output. By mid 2025, technology journalists described YouTube as being flooded with low-effort AI-generated videos and started reporting on YouTube's evolving approaches to managing this content.<sup>123</sup> By mid-2025, the platform blocked ad revenue sharing for channels clearly engaged in automated slop production, and in early 2026 it directly addressed the issue of slop, announcing plans to build on its anti-spam tools in order to mitigate slop's spread.<sup>124</sup> These measures suggest that modest adjustments to ranking signals, monetization thresholds, or eligibility rules can alter creator incentives without requiring a categorical ban on generative content itself.

Other platforms have taken varying stances on slop. Meta has made clear that it sees AI-generated content as the future of its platform, even as its subsidiary Instagram has taken a more cautious approach.<sup>125</sup> Spotify has launched a new spam filter, referring to this move as an "AI protection," though its external messaging focuses mainly on behavioral signals over sloppy content.<sup>126</sup> TikTok, for its part, is considering rolling out a new feature allowing users to select the amount of AI-generated content they encounter on their feed, although it is stopping short of allowing users to block AI content entirely.<sup>127</sup> Microsoft CEO Satya Nadella has implored the population writ large to think about AI as post-slop—as a tool rather than content, while at the same time his newly appointed head of Xbox has promised its customers it won't be relying on "soulless AI slop" in games.<sup>128</sup> For Reddit, which has garnered attention for being a training ground for many generative AI providers, AI slop poses a real challenge to the moderation community,<sup>129,130</sup> which ultimately decides subreddit by subreddit whether to allow AI-generated content on submissions. Many subreddits have banned AI slop, even as the company explores its own integrations of generative AI content.<sup>131</sup>

### **EXPLORING PLATFORMS’ DIVERGENT POLICIES ON SLOP (continued)**

The lack of clarity regarding AI slop and its permissiveness, boundaries, and governance raises broader questions about whether platforms will ultimately adopt explicit policies or continue to manage the issue through incremental adjustments. For now, most companies appear to be balancing experimentation with reputational risk, creating an opening for journalists, researchers, and civil society actors to continue to advocate for greater transparency in platform intentions and their efficacy.

### **TRANSPARENCY AND PROVENANCE SIGNALS**

A second set of potential interventions concerns verification. Efforts like watermarking, provenance tracking, or identity verification systems aim to provide audiences with clearer cues about how content was produced and by whom, and platform leaders have pointed to labeling as one pragmatic path forward. Instagram head Adam Mosseri has suggested a practical approach may be to label what is “real media” vs. labeling what is synthetic.<sup>132</sup> For Instagram, the goal is not to eliminate AIGC, but to clarify its status and origin so that users can interpret it appropriately. Over time, stronger provenance signals could reshape expectations about authenticity, attribution, and credibility in ways that alter both creator behavior and audience response.

Although provenance and labeling might help in some instances, they are not a cure-all. Labels can be inconsistent, easy to overlook in busy interfaces, and sometimes applied incorrectly even to real content, all of which erodes user trust and can lead users to ignore labels. Technical signals like watermarks or provenance metadata can be stripped, faked, or not adopted uniformly. Even when labeling “works,” research shows that telling people something is AI-generated often changes their opinions about who they think created it, but not how persuasive or influential they find the content, so the core issues for certain functions of slop remain intact.<sup>133</sup> A sampling of the AI slop reviewed as part of the salon proceedings indicated that visible watermarks were easily obfuscated by creator branding.

Efforts to create “nutrition labels” for information, such as source credibility ratings from NewsGuard,<sup>134</sup> transparency indicators from the Trust Project,<sup>135</sup> or crowdsourced context systems like Community Notes,<sup>136</sup> aim to reduce discernment costs by giving users structured signals about provenance, authorship, and trustworthiness. These models assume relatively stable publishers and bounded pieces of content, however, whereas in an AI slop environment, content is inexpensive to create, often ephemeral, and regularly detached from durable identities, while distribution and monetization dynamics play a central role in shaping impact. An “AI-generated” badge may clarify origin, but it expresses little about the broader context in which a piece of content was created and distributed. The nutrition label analogy is both appealing and limited: Labels may help to orient users, particularly

in some contexts such as in a crisis or where content is attempting to persuade, but meaningful intervention likely requires integrating signals and incentive structures, not just content attributes.

## **POLITICAL ACTOR SELF-REGULATION**

A third intervention path sits upstream of both platforms and creator tools: the political actors who commission or distribute AI-generated campaign content in the first place. A proposal in Quebec offers an early case study. In October 2025, IVADO and CEIMIA, two Montreal-based AI research and policy organizations, proposed a code of conduct for the responsible use of AI by Quebec political parties,<sup>137</sup> covering labeling of AI-generated content, commitments against deepfakes and targeted misinformation, data protection standards, and staff training on AI-related elections risks. In March 2026, the National Assembly of Quebec unanimously adopted a motion calling on all parties represented in the legislature to adhere to the proposal.<sup>138</sup>

As a voluntary instrument enforceable through reputation rather than sanction, the Quebec code is modest in scope. But it is still structurally noteworthy, operating on the demand side, binding political actors to commitments about what they will and will not commission or circulate. Voluntary codes of this kind face the familiar limitations of soft law instruments: Compliance is uneven, holdouts face no penalty beyond public criticism, and enforcement depends on a media and civil society ecosystem capable of identifying breaches. Whether the code meaningfully shapes the 2026 Quebec election will be an early empirical test of whether party-level self-regulation can meaningfully constrain slopaganda.

## **CREATOR-SIDE FRICTION AND DESIGN CHOICES**

Interventions may also arise upstream, at the level of production tools and creator workflows. Generative platforms already exercise some control over what their systems produce through content policies and safety filters, particularly in domains such as political persuasion, religion, and sensitive public-interest topics. Tool designers could also introduce forms of friction that discourage purely automated content pipelines, such as nudges encouraging disclosure, limits on bulk publishing, and constraints on fully automated posting integrations. These design choices would not prevent legitimate experimentation or creative use of generative tools but could increase the effort required to operate large-scale slop farms.

Economic dynamics may also shape this landscape over time. Although many providers have thus far offered access to generative models at low or zero marginal cost in order to drive adoption, shifts toward monetization or usage-based pricing could alter the viability of slop-driven business models. If the cost of large-scale generation rises, actors motivated primarily by advertising revenue may be pushed toward different strategies or platforms. The future prevalence of AI slop may depend as much on market structure and pricing decisions as on moderation policies.

## DIGITAL LITERACY INITIATIVES

Digital literacy is another oft-cited strategy for navigating today’s high-velocity, AI-infused information environment, but its historical focus on source evaluation and fact-checking will likely need to evolve in a slop-saturated space. Scholars argue that digital literacy now includes not only the ability to evaluate news accuracy but also skills like lateral reading (the practice of leaving a webpage or source to check what other trusted sources say about it before deciding whether it is credible),<sup>139</sup> critical awareness of platform incentives, and an understanding of how AI works in content production and distribution.<sup>140</sup> The concept of “multiliteracy” further suggests that literacy must span multiple modes of communication and meaning-making in digital spaces, not just conventional text evaluation. In the age of generative AI, educators and researchers emphasize that learners need frameworks that help them triage signals from noise, manage increasing cognitive overload, and evaluate not only “truth” but relevance, intent, and credibility.<sup>141</sup>

**In the age of generative AI, educators and researchers emphasize that learners need frameworks that help them triage signals from noise, manage increasing cognitive overload, and evaluate not only “truth” but relevance, intent, and credibility.**

---

## IMPLICATIONS OF SLOP

Beyond immediate moderation and governance challenges, the proliferation of AI slop raises deeper questions about the future structure of the internet itself. Some commentators have linked the rise of synthetic content to long-standing concerns captured by the so-called Dead Internet theory, the possibility that large portions of online discourse may gradually become dominated by automated or commercially optimized material, rather than genuine human interaction.<sup>142</sup>

Equity considerations further complicate this picture. Research on generative AI adoption suggests that although these technologies may offer powerful learning and productivity benefits, their advantages may accrue disproportionately to users in the Global North, where infrastructure, education systems, and language compatibility are strongest.<sup>143</sup> If generative systems reshape knowledge production while reinforcing existing linguistic, cultural, or educational hierarchies, they may deepen, rather than shrink, digital divides.

Finally, the recursive dynamics of generative systems introduce long-term epistemic concerns. Studies on model training warn that as AI-generated content increasingly enters the datasets used to train future models, feedback loops may lead to homogenization, degradation, or loss of information diversity over time.<sup>144</sup> This possibility, combined with growing calls for “zero trust” information environments in which users assume content may be synthetic unless verified, suggests that AI slop may ultimately reshape not only media ecosystems but also the baseline assumptions through which knowledge and authenticity are evaluated online.

As the research agenda in the next section of this report makes clear, addressing the community’s measurement gaps with respect to slop is foundational to many of the proceeding’s identified open questions.

## THE AI SLOP OUROBOROS

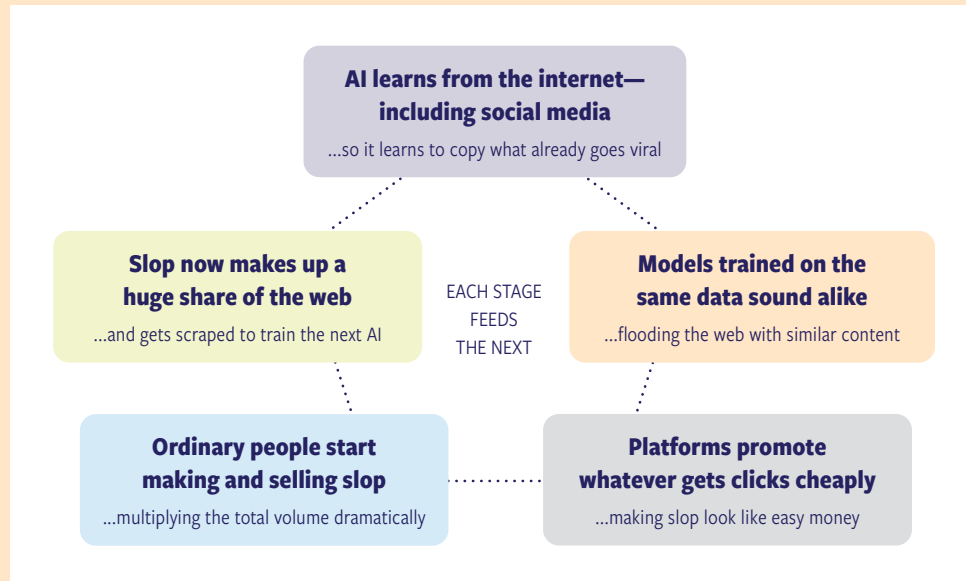


Image: the AI slop Ouroboros

The concept of an AI slop Ouroboros—generative models contaminating their own training data by ingesting their previous outputs—was formalized by Shumailov et al. in their 2024 *Nature* paper, which showed that the “indiscriminate use of model-generated content in training causes irreversible defects in the resulting models, in which tails of the original content distribution disappear.”<sup>145</sup> The salon’s collective discussions surfaced a larger, socially mediated version of the same loop, one in which platform amplification logic plays an active role, functioning as follows:

- 1. How AI models are built: the role of training data.** Foundation models ingest text, images, and video scraped from the open web to learn language, images and styles, much of it from social media platforms. The aesthetic and stylistic conventions of the content on those platforms thus become part of what the models will reliably produce in the future. (One can at least hope that open-source models fine-tuned across many cultural contexts will hold open the kind of unruly understorey that monocultures, by their nature, cannot see.)
- 2. AI generated content at scale (including slop):** Because mainstream models are trained on overlapping sources and reward signals, their outputs can tend to sound and look alike. When the models generate variations of the same wording conventions, image styles, and sentence rhythms, this accelerates a monoculture of content.

## THE AI SLOP OUROBOROS (continued)

- 3. Platform amplification.** Slop is structurally well-suited to engagement metrics: cheap to produce, easy to A/B test, and infinitely iterable. One of slop's distinctive features is the gap between the attention it captures and the attention it sustains—enough to trigger algorithmic amplification, even if it quietly frustrates users who may not want to drown in it.
- 4. Easy money - recruitment into slopreneurship.** As amplification rewards volume, growing numbers of ordinary people (not propagandists or troll farms, but people looking for income) get recruited into producing slop and even selling slop tutorials. This may not be to manipulate anyone, but to buy themselves watches and cars. This decentralized production model feels genuinely new, and it scales.
- 5. The feedback loop repeats.** As the volume of slop rises, that content becomes a more significant share of what's on the web, which makes it a larger share of future training data. The loop cycles back into the beginning, the training data stage, and the loop tightens with each iteration.

The downstream consequence is that social platforms may drift away from their original purpose. Trends that predate generative AI—algorithmic feeds displacing chronological ones, creators chasing metrics over communities, ad-tech logic colonizing every surface, etc.—get sharply compressed in time.

Of course, the empirical magnitude of all this remains contested, and evidence for systemic degeneration is uneven and difficult to measure (notably under current platform data access conditions). But the structural logic is clear, and it points somewhere specific.

Social media platforms caught in this AI slop Ouroboros are no longer functioning as the global public squares they were once described as. They are becoming something else: closed circuits in which humans optimize content to exploit algorithms, and algorithms optimize distribution to exploit humans, while real audiences tune out. Just as researchers ask whether recursive training data leads to model collapse, it is worth asking whether the AI slop Ouroboros is leading us toward the collapse of social media as we have known it.

And whether that, in the end, might not be such a bad thing.

*Written by Camille François*

---

## WHAT WE STILL DON'T KNOW: A RESEARCH AGENDA

The Slop Salon uncovered a broad set of open questions. Some were contested conceptual questions, where disagreement reflects competing normative framings rather than gaps in data. Even with unlimited evidence, the field may remain divided on what “slop” is and whether the term should exist at all. Other questions were empirical, answerable in principle but constrained by current methodological limits and platform data access barriers. Still others were governance and intervention questions that sit at the boundary of research and policy, where relevant evidence will emerge only through platform experimentation, regulatory intervention, or both. Treating these three categories of questions distinctly matters for the research agenda that follows, as they require different methods, different audiences, and different relationships between researchers and the platforms whose content and phenomena they study.

### CONTESTED CONCEPTUAL QUESTIONS

#### NORMATIVE QUESTIONS

The first open question concerns whether the word *slop* should carry a built-in normative judgment. Using the term in harm-focused work assumes low quality or negative impact, which is useful for describing the challenges that platforms and policymakers are trying to address, but risks dismissing AIGC that audiences actively value or that functions expressively, as in the Italian Brainrot and synthetic dissent cases described above. A neutral definitional vocabulary may be analytically cleaner, but it would require shared terms the field does not currently have. Conversely, preserving the normative valence of *slop* may be the only way to capture what is at stake in harm-focused work, at the cost of ruling out the content some consumers embrace.

#### LENSES AND UNITS OF ANALYSIS

Another, related open question concerned the unit of analysis: which properties of slop content (e.g., format, intent, aesthetics, distribution patterns, or production methods) should be prioritized when observing and analyzing AI slop and its potential impacts. The Slop Funnel and the 7Vs frameworks treat slop primarily as a property of content, while the supply-and-demand and systems discussions frame it as a property of infrastructure: the incentive structures, ranking systems, and monetization pipelines that would produce slop-like outputs regardless of who generates them. Depending on which frame is adopted, the research object differs: a single artifact, a creator’s production pattern, a platform’s ranking behavior, or the sociotechnical ecosystem as a whole.

There is an unresolved question about where AI slop sits in relation to longer traditions of low-cost, low-status media. Memes, tabloids, and pulp fiction have all been dismissed in the past as degraded forms before being reappraised as culturally or politically significant. Whether AI slop is best understood as part of a longer lineage of low-cost, high-volume media production, or as a genuinely new category whose scale and dynamics make such comparisons misleading, remains an open ques-

tion across media studies, platform research, and political theory. It has direct implications for how seriously to take the expressive and participatory framings that emerged from the convening, and it cannot be settled by studying slop in isolation from history.

Progress on these questions will depend on sustained conceptual work across disciplines, necessitating close engagement between researchers who approach slop as content, those who approach it as infrastructure, and those situating it within longer cultural histories. Progress also depends on a willingness to let competing framings coexist in the literature rather than forcing premature consensus or convergence.

## EMPIRICAL QUESTIONS

Empirical questions are in principle answerable but are currently constrained by insufficient data access and measurement infrastructure. Most of the empirical questions uncovered during the convening and in the broader literature remain open because of structural constraints on independent research rather than because they are conceptually difficult.

**1. How can second-order harms be measured?** First-order harms from AI slop like scams, harassment, NCII, and content optimized for engagement farming are typically observable at the level of individual artifacts, even if moderation at scale remains challenging and enforcement is imperfect. Second-order harms, including epistemic degradation, labor market displacement, effects on cognition and youth development, and the environmental costs of scaled generation, are diffuse by definition and emerge from accumulation rather than from any single piece of content. Researchers studying these phenomena currently lack shared baselines against which to detect these effects, lack methodologies capable of isolating slop's contribution from other features of the information environment, and lack longitudinal platform data that would allow inflection points to be identified. The Ouroboros effect is the clearest example of a second-order dynamic that is theoretically well-specified, but whose real-world trajectory in deployed systems remains largely unknown.

**2. When slop fills a feed, search result, or recommendation surface, what gets displaced?**

The answer matters for assessing whether slop is additive to the information environment or substitutive for higher-quality alternatives. This is especially important for such domains as scientific communication, local journalism, and educational content, where displacement could be consequential. Answering this question requires visibility into platform ranking behavior and counterfactual comparisons.

**3. What relationship is evolving between AI providers and platforms as these previously distinct elements converge?** The typology of separate systems, tool-integrated platforms, hybrid AI-social platforms, and AI-native content ecosystems introduced earlier (Figure 7) is a starting point rather than a settled matter. How content generation and distribution interact across these categories, including the velocity of production, the share of feed composition, and the feedback loops between generative tools and the platforms that distribute their outputs is not currently

observable at the scale needed. The Sora shutdown offers one data point about consumer demand for slop-as-destination versus slop-as-embedded, but it is a single case, and whether its lesson will generalize across platform types and markets is unclear.

Further empirical questions follow from the supply-and-demand framing. Koebler’s reporting on the sloprepreneur and “hustle bro” ecosystem opens a research agenda on the political economy of slop. How large are these content-creation networks, how durable are they, and how do they respond to platform policy shifts or changes in generative model pricing? What would a typology of slop creators look like? The labor-market effects of scaled generation on creators, translators, and editorial workers are similarly under-measured, as are the downstream effects on audiences whose primary encounter with generative AI is through slop rather than through productivity tools. Shaib et al. have begun to formalize slop measurement, as mentioned, but the broader question of whether reliable automated detection is achievable, and what the costs of false positives and negatives are in practice, remains open.

Progress on this cluster depends less on conceptual clarification than on research infrastructure. The loss of CrowdTangle and the broader contraction of platform data access described earlier remains a constraint on this work.

## **GOVERNANCE AND INTERVENTION QUESTIONS**

A third topic area concerns the efficacy of interventions, which depends on analysis of policy decisions that are made outside the research community. Evidence and analysis should be feasible after deployment and observation.

The labeling and provenance question is the most developed. Early evidence, including the Isabel Gallegos et al. finding that disclosure changes inferences about authorship but not about persuasiveness,<sup>146</sup> suggests that labels alone are unlikely to address the core dynamics of slop. It is unclear whether this finding holds under different slop-saturation conditions, across content types, or when labels are integrated with other signals such as source credibility ratings or crowdsourced context. The question of whether stronger provenance infrastructure changes these dynamics will be answered only through sustained deployment and independent evaluation.

A parallel question concerns creator-side friction. The logic of interventions at the generation layer, including through usage-based pricing, bulk-publishing limits, disclosure nudges, and constraints on fully automated posting, is that raising production costs should reduce the viability of arbitrage-driven slop production without constraining legitimate use. Whether this logic holds in practice is unknown, and it depends on such elasticities as how price-sensitive slop production is, how quickly actors make substitutions across generative tools or platforms, and whether friction displaces activity rather than reducing it, elasticities that have not been systematically measured. The question becomes particularly important if generative model access shifts from free or low cost toward metered pricing, which would alter slop economics independent of any platform intervention.

The convening raised, but did not resolve, the question of how agentic AI (the autonomous systems that can plan and act independently to complete complex tasks with limited human oversight) will change the landscape. Slop today is still regularly produced through human-directed generation (vs. fully outsourced to automated systems) even when volume is high; the creator still has to prompt, select, and post. Agentic systems that are capable of running full production pipelines identifying topics, generating content, posting across platforms, and optimizing in response to engagement signals would change the relationship between generation and distribution in ways that existing frameworks are not built to capture. Whether agentic slop production remains a niche experiment or becomes the dominant modality is a consequential open question for near-term research.

Finally, the digital literacy question remains under-specified. Traditional media literacy frameworks emphasize source evaluation, fact-checking, and lateral reading, all of which assume relatively stable publishers and verifiable claims. In a slop-saturated environment, the relevant literacy may concern pattern recognition across the ecosystem rather than evaluation of individual artifacts, and it may need to account for the declining information value of provenance signals themselves. What this form of literacy will look like in practice, how it can be taught at scale, and how it will interact with platform-level interventions rather than substituting for them make up a research agenda that has not yet been fully articulated.

Progress here depends substantially on decisions made outside the research community. Which interventions get deployed, at what scale, and with what transparency will determine what can be studied independently.

---

## ENDNOTES

- 1 Hans W. A. Hanley and Zakir Durumeric, “Machine-Made Media: Monitoring the Mobilization of Machine-Generated Articles on Misinformation and Mainstream News Websites,” preprint, arXiv, May 2023, <https://arxiv.org/abs/2305.09820v5>.
- 2 Madeleine I. G. Daepf, Alejandro Cuevas, Robert Osazuwa Ness, Vickie Yu-Ping Wang, Bharat Kumar Nayak, Dibyendu Mishra, Ti-Chung Cheng, Shaily Desai, and Joyojeet Pal, “Generative Propaganda,” preprint, arXiv, September 23, 2025, <https://doi.org/10.48550/arXiv.2509.19147>.
- 3 Cody Kommers, Eamon Duede, Julia Gordon, Ari Holtzman, Tess McNulty, Spencer Stewart, Lindsay Thomas, Richard Jean So, and Hoyt Long, “Why Slop Matters,” preprint, arXiv, December 23, 2025, <https://doi.org/10.48550/arXiv.2601.06060>.
- 4 “Seedance 2 AI Video.” Seedance 2, 2026, <https://www.seedance2ai.io/>.
- 5 Google. “Ve0.” Google DeepMind, 2025, <https://deepmind.google/models/ve0/>.
- 6 Cory Doctorow, *Enshittification: Why Everything Suddenly Got Worse and What to Do About It* (New York: Farrar, Straus and Giroux, 2025).
- 7 Raluca Csernaton, “The EU’s AI Power Play: Between Deregulation and Innovation,” Carnegie Endowment for International Peace, May 20, 2025, <https://carnegieendowment.org/research/2025/05/the-eus-ai-power-play-between-deregulation-and-innovation>.
- 8 Clare Duffy, “Meta Is Getting Rid of Fact Checkers. Zuckerberg Acknowledged More Harmful Content Will Appear on the Platforms Now,” *CNN*, January 7, 2025, <https://www.cnn.com/2025/01/07/tech/meta-censorship-moderation>.
- 9 Kate Knibbs, “Scammy AI-Generated Book Rewrites Are Flooding Amazon,” *Wired*, January 10, 2024, <https://www.wired.com/story/scammy-ai-generated-books-flooding-amazon/>.
- 10 Aisha Down, “More than 20% of Videos Shown to New YouTube Users Are ‘AI Slop’, Study Finds.” *The Guardian*. December 27, 2025, <https://www.theguardian.com/technology/2025/dec/27/more-than-20-of-videos-shown-to-new-youtube-users-are-ai-slop-study-finds>.
- 11 Anna Furman, “Merriam-Webster’s Word of the Year for 2025 Is AI ‘Slop.’” *PBS News*, December 15, 2025, <https://www.pbs.org/newshour/nation/merriam-websters-word-of-the-year-for-2025-is-ais-slop>.
- 12 Last Week Tonight. “AI Slop: Last Week Tonight with John Oliver (HBO).” YouTube, June 22, 2025, <https://www.youtube.com/watch?v=TWpg1RmzAbc>.
- 13 “AI slop.” Google Trends, accessed May 21, 2026, <https://trends.google.com/explore?q=AI%20slop&date=2024-01-01%202026-03-02&geo=US&gprop=web>
- 14 Wardle, Claire. “Fake News. It’s Complicated.” *First Draft News*, February 16, 2017, <https://firstdraftnews.org/articles/fake-news-complicated/>.
- 15 Dag Øivind Madsen and Richard W. Puyt, “The 7Vs of AI Slop: A Typology of Generative Waste,” SSRN Scholarly Paper no. 5558018, *Social Science Research Network*, October 2, 2025, <https://doi.org/10.2139/ssrn.5558018>.
- 16 Kommers, Cody, Eamon Duede, Julia Gordon, Ari Holtzman, Tess McNulty, Spencer Stewart, Lindsay Thomas, Richard Jean So, and Hoyt Long. “Why Slop Matters.” Preprint, arXiv, December 23, 2025. <https://doi.org/10.48550/arXiv.2601.06060>.
- 17 Alexios Mantzarlis, “A Tentative Taxonomy of AI Slop,” *Indicator*, July 16, 2025, <https://indicator.media/p/a-tentative-taxonomy-of-ai-slop>.
- 18 Ibid.
- 19 Ibid.
- 20 Ibid.
- 21 Dani Di Placido, “Facebook’s AI-Generated ‘Shrimp Jesus,’ Explained,” *Forbes*, April 28, 2024, <https://www.forbes.com/sites/danidiplacido/2024/04/28/facebooks-surreal-shrimp-jesus-trend-explained/>.
- 22 Chantal Shaib, Tuhin Chakrabarty, Diego Garcia-Olano, and Byron C. Wallace, “Measuring AI ‘Slop’ in Text,” preprint, arXiv, 2025, <https://arxiv.org/abs/2509.19163>.
- 23 Stefano Faraoni, “Persuasive Technology and Computational Manipulation: Hypernudging out of Mental Self-Determination,” *Frontiers in Artificial Intelligence* 6 (July 2023), <https://doi.org/10.3389/frai.2023.1216340>; Michael Klenk, “Ethics of Generative AI and Manipulation: A Design-Oriented Research Agenda,” *Ethics and Information Technology* 26, no. 1 (2024), <https://doi.org/10.1007/s10676-024-09745-x>; Marcello Ienca, “On Artificial Intelligence and Manipulation,” *Topoi* 42 (June 2023), <https://doi.org/10.1007/s11245-023-09940-3>; S. C. Matz et al., “The Potential of Generative AI for Personalized Persuasion at Scale,” *Scientific Reports* 14, no. 1 (2024): 4692, <https://doi.org/10.1038/s41598-024-53755-0>.

- 24 Jasper Roe, Leon Furze, and Mike Perkins, “GenAI as Digital Plastic: Understanding Synthetic Media through Critical AI Literacy,” preprint, arXiv, February 12, 2025, <https://doi.org/10.48550/arXiv.2502.08249>.
- 25 Jason Miklian and Kristian Hoelscher, “A New Digital Divide? Coder Worldviews, the ‘Slop Economy,’ and Democracy in the Age of AI,” *Information, Communication & Society* 0, no. 0 (2026): 1–17, <https://doi.org/10.1080/1369118X.2025.2566814>.
- 26 Su-jun Fang, Zhao-di Yin, Li-fan Li, Qi Cai, Peng-fei Zheng, and Li-zhen Chen, “Overall Effects of Microplastics on Brain,” *Frontiers in Toxicology* (2025), <https://doi.org/10.3389/ftox.2025.1619096>.
- 27 Kurt Wagner and Alexandra S. Levine, “Meta, Google Risk Big Tobacco-like Fallout after Addiction Trial,” *Bloomberg*, March 26, 2026, <https://www.bloomberg.com/news/articles/2026-03-26/meta-google-risk-big-tobacco-like-fallout-after-addiction-trial>.
- 28 Kara Phillips, “One of Roblox’s Biggest Experiences Right Now Is a Bizarre Italian Brainrot Character-Stealing Simulator And, Lord Help Me, I Now Understand Why It’s so Popular,” *PC Gamer*, September 2025, <https://www.pcgamer.com/games/sim/one-of-robloxs-biggest-experiences-right-now-is-a-bizarre-italian-brainrot-character-stealing-simulator-and-lord-help-me-i-now-understand-why-its-so-popular/>.
- 29 Fabian Mosele, “Italian Brainrot: When AI Slop Becomes Culture,” *Medium*, May 11, 2025, <https://medium.com/@fabianmosele/italian-brainrot-when-ai-slop-becomes-culture-4cff8c344d66>.
- 30 Kommers et al., “Why Slop Matters.”
- 31 Nesrine Malik, “With ‘AI Slop’ Distorting Our Reality, the World Is Sleepwalking into Disaster,” *The Guardian*, April 21, 2025, <https://www.theguardian.com/commentisfree/2025/apr/21/ai-slop-artificial-intelligence-social-media>.
- 32 Hao-Ping (Hank) Lee et al., “The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers,” *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, April 26, 2025, 1–22, <https://doi.org/10.1145/3706598.3713778>.
- 33 Xin Wang, Xuanzhu Zhao, and Changsheng Yu, “The Influence of Information and Social Overload on Academic Performance: The Role of Social Media Fatigue, Cognitive Depletion, and Self-Control,” *Revista de Psicodidáctica (English Ed.)* 30, no. 2 (2025): 500164, <https://doi.org/10.1016/j.psicoe.2025.500164>.
- 34 Vaishnavi Subhash Naik, Edlin Glane Mathias, Priyanka Krishnan, and Vanitha Jagannath, “Impact of Social Media on Cognitive Development of Children and Young Adults: A Systematic Review,” *BMC Pediatrics* 25 (October 2025): 826, <https://doi.org/10.1186/s12887-025-06041-5>.
- 35 “JoJo Funland,” YouTube, accessed May 9, 2026, [https://www.youtube.com/channel/UCFzQRN0YkQJc\\_nBoQlpykTQ](https://www.youtube.com/channel/UCFzQRN0YkQJc_nBoQlpykTQ).
- 36 Emily Tate Sullivan, “AI ‘Slop’ Is Flooding Children’s Media. Parents Should Be Very Alarmed,” *The 74*, March 18, 2026, <https://www.the74million.org/zero2eight/ai-slop-is-flooding-childrens-media-parents-should-be-very-alarmed/>.
- 37 Arijeta Lajka, “How A.I.-Generated Videos Are Distorting Your Child’s YouTube Feed,” U.S., *The New York Times*, February 26, 2026, <https://www.nytimes.com/2026/02/26/us/ai-videos-children-youtube.html>.
- 38 Tiffany Munzer, Joanna Parga-Belinkie, Libby Matile Milkovich, Suzy Tomopoulos, Taiwo Ajumobi, Corinn Cross, Roslyn Gerwin, and Sheri Madigan, on behalf of the American Academy of Pediatrics Council on Communications and Media, “Digital Ecosystems, Children, and Adolescents: Policy Statement,” *Pediatrics* 157, no. 2 (2026): e2025075320, <https://doi.org/10.1542/peds.2025-075320>.
- 39 Ashwin Verghese, “YouTube: Stop ‘AI Slop’ for Kids, Says Letter from Fairplay, over 200 Experts, Including Jonathan Haidt,” *Fairplay*, April 1, 2026, <https://fairplayforkids.org/youtube-stop-ai-slop-for-kids-says-letter-from-fairplay-over-200-experts-including-jonathan-haidt/>.
- 40 Fairplay et al., letter to Sundar Pichai and Neal Mohan, 2026, <https://fairplayforkids.org/wp-content/uploads/2026/03/YouTube-Letter-AI-Slop.pdf>.
- 41 Bobby Allyn, “Jury Orders Meta and Google to Pay Woman \$3 Million in Social Media Addiction Trial,” NPR, March 25, 2026, <https://www.npr.org/2026/03/25/nx-s1-5746125/meta-youtube-social-media-trial-verdict>.
- 42 Iris van Rooij, *AI Slop and the Destruction of Knowledge*, August 12, 2025, <https://doi.org/10.5281/zenodo.16905560>.
- 43 Jane Tullis, “Sifting Through the Slop: How Generative AI Created a Market for Lemons for Text-Based Works,” SSRN Scholarly Paper no. 5266660, Social Science Research Network, April 1, 2025, <https://doi.org/10.2139/ssrn.5266660>.
- 44 Kate Niederhoffer et al., “AI-Generated ‘Workslop’ Is Destroying Productivity,” *Harvard Business Review*, September 22, 2025, <https://hbr.org/2025/09/ai-generated-workslop-is-destroying-productivity>; Aidan Walker, “Slop Capitalism,” *How To Do Things With Memes* (Substack), April 15, 2025, <https://howtodothingswithmemes.substack.com/p/slop-capitalism>.
- 45 Scott Markley and Diana Drogaris, “Zoning and Municipal Data Brokers: Toward a Political Economy of Data Slop,” SSRN Scholarly Paper no. 5188578, Social Science Research Network, March 21, 2025, <https://doi.org/10.2139/ssrn.5188578>.
- 46 Rhys Crilley and Robert A. Saunders, “The Resonant Slop Machine: Public Diplomacy and Strategic Narratives in the Age of Artificial Intelligence,” *Place Branding and Public Diplomacy*, ahead of print, December 15, 2025, <https://doi.org/10.1057/s41254-025-00419-z>.

- 47 Nesrine Malik, “With ‘AI Slop’ Distorting Our Reality, the World Is Sleepwalking into Disaster.”
- 48 Allan Smith, “Trump Posts AI Video Showing Him Dumping on No Kings Protesters,” *NBC News*, October 20, 2025, <https://www.nbcnews.com/politics/donald-trump/trump-posts-ai-video-dumping-no-kings-protesters-rcna238521>.
- 49 Sarah McCammon, Hadeel Al-Shalchi, and Anas Baba, “Trump’s Social Media Video Garners Pushback from Arabs and Muslims in U.S. and Gaza,” *NPR*, February 26, 2025, <https://www.npr.org/2025/02/26/nx-s1-5309695/trump-gaza-video>.
- 50 Charles Pulliam-Moore, “Who Is AI Nostalgia Slop Even For?,” *The Verge*, November 21, 2025, <https://www.theverge.com/ai-artificial-intelligence/825498/openai-sora-video-slop-nostalgia-mister-rogers>.
- 51 Ethan Karp, “The Truth About A Chinese AI Video Mocking Efforts To Revitalize U.S. Manufacturing,” *Forbes*, April 28, 2025, <https://www.forbes.com/sites/ethankarp/2025/04/28/the-truth-about-the-chinese-ai-video-mocking-efforts-to-revitalize-american-manufacturing/>.
- 52 Maria Flannery, “Slopaganda: How Iran Is Using Memes to Troll Trump,” *Eurovision News Spotlight*, April 17, 2026, <https://spotlight.ebu.ch/p/iran-ai-slopaganda-trump-memes>.
- 53 AFP, “With Legos, Trolling and Twain, Iran Pushes War Narrative on Social Media,” *France24*, April 7, 2026, <https://www.france24.com/en/live-news/20260407-with-legos-trolling-and-twain-iran-pushes-war-narrative-on-social-media>.
- 54 “Disclosing Use of Altered or Synthetic Content,” YouTube Help, Google, 2026, <https://support.google.com/youtube/answer/14328491?hl=en&co=GENIE.Platform%3DAndroid>.
- 55 Sviatlana Tsikhanouskaya (@Tsihanouskaya), post on X, February 23, 2024, <https://x.com/Tsihanouskaya/status/1761000189664375144>.
- 56 Allie Funk, Kian Vesteinsson, and Grant Baker, “The Struggle for Trust: Online Freedom on the Net 2024,” Freedom House, 2024, <https://freedomhouse.org/report/freedom-net/2024/struggle-trust-online#key-findings>.
- 57 Al Jazeera Staff, “Why Are There Protests in Türkiye? What to Know,” *Al Jazeera*, March 23, 2025, <https://www.aljazeera.com/features/2025/3/23/why-are-there-protests-in-turkiye-what-to-know>.
- 58 Selin Bengi Gümrükçü, “Artificial Intelligence–Based Aesthetics of Dissent in,” *Public Seminar*. June 18, 2025, <https://publicseminar.org/2025/06/turkey-protesters-using-ai/>.
- 59 Stefano Pozzebon, “In Venezuela, AI news anchors aren’t replacing journalists. They’re protecting them,” *CNN World*, September 18, 2024, <https://www.cnn.com/2024/09/18/americas/venezuela-retweets-ai-news-maduro-intl-latam>.
- 60 Asker Bryld Staunæs and Michael Maly, “Synthetic Dissidents: How AI Protects Dissent under Repression,” *The Loop*, February 13, 2026, <https://theloop.eopr.eu/synthetic-dissidents-how-ai-protects-dissent-under-repression/>.
- 61 “Universal Declaration of Human Rights,” OHCHR, <https://www.ohchr.org/en/human-rights/universal-declaration/translations/english?LangID=eng>.
- 62 United Nations, “International Covenant on Civil and Political Rights,” OHCHR, December 16, 1966, <https://www.ohchr.org/en/instruments-mechanisms/instruments/international-covenant-civil-and-political-rights>.
- 63 Srdja Popovic and Mladen Joksic, “Why Dictators Don’t Like Jokes,” *Foreign Policy*, April 5, 2013, <https://foreignpolicy.com/2013/04/05/why-dictators-dont-like-jokes/>.
- 64 Renée DiResta and Josh A. Goldstein, “How Spammers and Scammers Leverage AI-Generated Images on Facebook for Audience Growth,” *Harvard Kennedy School Misinformation Review*, ahead of print, August 15, 2024, <https://doi.org/10.37016/mr-2020-151>.
- 65 Monica Alba, Jennifer Jett, and Rebecca Shabad, “White House Shares an Altered Photo of Arrested Minnesota Protester Nekima Levy Armstrong,” *NBC News*, January 23, 2026, <https://www.nbcnews.com/politics/white-house/white-house-shares-altered-photo-arrested-minnesota-protester-nekima-l-rcna255595>.
- 66 Camille François, “Actors, Behaviors, Content: A Disinformation ABC Highlighting Three Vectors of Viral Deception to Guide Industry & Regulatory Responses,” *Annenberg Public Policy Center of the University of Pennsylvania*, September 20, 2019, [https://www.annenbergpublicpolicycenter.org/wp-content/uploads/ABC\\_Framework\\_TWG\\_Francois\\_Sept\\_2019.pdf](https://www.annenbergpublicpolicycenter.org/wp-content/uploads/ABC_Framework_TWG_Francois_Sept_2019.pdf).
- 67 Iliia Shumailov, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson, “The Curse of Recursion: Training on Generated Data Makes Models Forget,” preprint, arXiv, April 14, 2024, <https://doi.org/10.48550/arXiv.2305.17493>.
- 68 Jason Koebler, “AI Generated ‘Boring History’ Videos Are Flooding YouTube and Drowning out Real History,” *404 Media*, September 3, 2025, <https://www.404media.co/ai-generated-boring-history-videos-are-flooding-youtube-and-drowning-out-real-history/>.
- 69 “Sleepless Historian,” YouTube, Accessed May 9, 2026, <https://www.youtube.com/channel/UCy3Hgwj12x3fDQUcutNm2DA>.
- 70 “History Time,” YouTube, <https://www.youtube.com/@HistoryTime>.
- 71 Jason Koebler, “AI Generated ‘Boring History’ Videos Are Flooding YouTube and Drowning out Real History,” *404 Media*, September 3, 2025, <https://www.404media.co/ai-generated-boring-history-videos-are-flooding-youtube-and-drowning-out-real-history/>.

- 72 Matthew Magnani and Jon Clindaniel, “Artificial Intelligence and the Interpretation of the Past,” *Advances in Archaeological Practice*, December 2025, 1–16, <https://doi.org/10.1017/aap.2025.10110>.
- 73 Sayash Kapoor and Nadine Arvind, “We Looked at 78 Election Deepfakes. Political Misinformation Is Not an AI Problem,” Knight First Amendment Institute, December 13, 2024, <https://knightcolumbia.org/blog/we-looked-at-78-election-deepfakes-political-misinformation-is-not-an-ai-problem>.
- 74 Nina Lutz, Joseph Schafer, Julie Vera, Sourojit Ghosh, and Kate Starbird, “*Spider Jesus: Christian Visual AI Slop and Scams on Facebook*”, working paper, 2026.
- 75 Fabian Mosele, “Italian Brainrot: When AI Slop Becomes Culture,” *Medium*, May 11, 2025, <https://medium.com/@fabianmosele/italian-brainrot-when-ai-slop-becomes-culture-4cff8c344d66>.
- 76 Kommers et al., “Why Slop Matters.”
- 77 Raffael Heiss and Isabelle Freiling, “Addressing Social Media Platforms’ Influence on Academic Research,” *Humanities and Social Sciences Communications* 13, no. 1 (2026): 192, <https://doi.org/10.1057/s41599-026-06690-6>.
- 78 Alia Shoaib, “New Singer Dominating iTunes Chart Is AI-Generated,” *Newsweek*, April 7, 2026, <https://www.newsweek.com/new-singer-dominating-itunes-chart-is-ai-generated-11791886>.
- 79 Murray Stassen, “Music Streaming Platforms Now Host Quarter of a BILLION Tracks. Where Does It End?” *Music Business Worldwide*, January 14, 2026, <https://www.musicbusinessworldwide.com/quarter-of-a-billion-tracks-now-sit-on-music-streaming-services-where-does-it-end/>.
- 80 Murray Stassen, “Sir Lucian Grainge Talks AI Music, Superfans in 2026 Memo, Taking Swipe at Firms ‘Validating Business Models That Fail to Respect Artists’ Work,’” *Music Business Worldwide*, January 8, 2026, <https://www.musicbusinessworldwide.com/sir-lucian-grainge-talks-ai-music-superfans-in-2026-memo-taking-swipe-at-firms-validating-business-models-that-fail-to-respect-artists-work/>.
- 81 Ceren Budak et al., “Misunderstanding the Harms of Online Misinformation,” *Nature* 630, no. 8015 (2024): 45–53, <https://doi.org/10.1038/s41586-024-07417-w>.
- 82 Jon Roozenbeek, “Social Identity and Misinformation,” *Lawfare*, August 29, 2025, <https://www.lawfaremedia.org/article/social-identity-and-misinformation>.
- 83 Felix M. Simon, Sacha Altay, and Hugo Mercier, “Misinformation Reloaded? Fears about the Impact of Generative AI on Misinformation Are Overblown,” *Harvard Kennedy School Misinformation Review*, ahead of print, October 18, 2023, <https://doi.org/10.37016/mr-2020-127>.
- 84 Kommers et al., “Why Slop Matters.”
- 85 Jasper Roe, Leon Furze, and Mike Perkins, “GenAI as Digital Plastic: Understanding Synthetic Media Through Critical AI Literacy,” arXiv:2502.08249, preprint, arXiv, February 12, 2025, <https://doi.org/10.48550/arXiv.2502.08249>.
- 86 EJ Dickson, “This Scammer Used an AI-Generated MAGA Girl to Grift ‘Super Dumb’ Men,” *Wired*, April 21, 2026, <https://www.wired.com/story/ai-generated-maga-girls/>.
- 87 AFP, “Burkinabe Teen behind Viral French ‘coup’ Video Has No Regrets,” *France 24*, September 19, 2025, <https://www.france24.com/en/live-news/20251219-burkinabe-teen-behind-viral-french-coup-video-has-no-regrets>.
- 88 As of this writing, the exchange rate: <https://www.oanda.com/currency-converter/en/?from=EUR&to=USD&amount=7>
- 89 AFP, “Burkinabe Teen behind Viral French ‘coup’ Video Has No Regrets,” *France 24*, September 19, 2025, <https://www.france24.com/en/live-news/20251219-burkinabe-teen-behind-viral-french-coup-video-has-no-regrets>.
- 90 As of this writing, the exchange rate: <https://www.oanda.com/currency-converter/en/?from=EUR&to=USD&amount=10>
- 91 As of this writing, the exchange rate: <https://www.oanda.com/currency-converter/en/?from=EUR&to=USD&amount=7>
- 92 Ifeoma Joy Okorie, “Meta Approves Monetisation for Content Creators in Nigeria, Ghana,” *Techpoint Africa*, June 28, 2024, <https://techpoint.africa/news/meta-monetisation-nigeria-ghana/>.
- 93 Shelton Trending, “How to Monetize Tiktok from South Africa and Anywhere Else,” YouTube, December 16, 2024, <https://www.youtube.com/watch?v=d4woFC8Ks4s>.
- 94 Cedric Thompson, “What Is Arbitrage? Definition, Example, and Costs” *Investopedia*, updated April 22, 2025, <https://www.investopedia.com/ask/answers/what-is-arbitrage/>.
- 95 Veronica Kovalenko, “YouTube CPM Rates: Changes You Need to Know (2026),” *Mediacube – Blog*, February 23, 2026, <https://mediacube.io/en-US/blog/youtube-cpm>.
- 96 Rob Waugh, “Journalist Says 4,000 Fake AI News Websites Created to Game Google Algorithms,” *Press Gazette*, July 11, 2025, <https://pressgazette.co.uk/news/french-journalist-who-uncovered-4000-fake-ai-news-websites-warns-uk-could-be-next/>; Jean-Marc Manach, “Nous Avons Découvert Des Milliers de Sites d’info Générés Par IA,” *Next.Ink*, February 17, 2026, <https://next.ink/173214/recap-nous-avons-decouvert-des-milliers-de-sites-dinfo-generes-par-ia-tous-nos-articles/>.

- 97 David S. Evans, *The Economics of Attention Markets* (April 15, 2020), Available at SSRN: <https://ssrn.com/abstract=3044858> or <http://dx.doi.org/10.2139/ssrn.3044858>.
- 98 Caiwei Chen, “How I Learned to Stop Worrying and Love AI Slop,” *MIT Technology Review*, December 23, 2025, <https://www.technologyreview.com/2025/12/23/1130396/how-i-learned-to-stop-worrying-and-love-ai-slop/>.
- 99 Aisha Malik, “Meta Launches ‘Vibes,’ a Short-Form Video Feed of AI Slop | TechCrunch,” *TechCrunch*, September 25, 2025, <https://techcrunch.com/2025/09/25/meta-launches-vibes-a-short-form-video-feed-of-ai-slop/>.
- 100 “Meta Reports Third Quarter 2025 Results,” *Meta Investor Relations*, 2025, <https://investor.atmeta.com/investor-news/press-release-details/2025/Meta-Reports-Third-Quarter-2025-Results/default.aspx>. Transcript: [https://s21.q4cdn.com/399680738/files/doc\\_financials/2025/q3/META-Q3-2025-Earnings-Call-Transcript.pdf](https://s21.q4cdn.com/399680738/files/doc_financials/2025/q3/META-Q3-2025-Earnings-Call-Transcript.pdf).
- 101 Kate Conger, Dylan Freedman, and Stuart A. Thompson, “Musk’s Chatbot Flooded X With Millions of Sexualized Images in Days, New Estimates Show,” *The New York Times*, January 22, 2026, <https://www.nytimes.com/2026/01/22/technology/grok-x-ai-elon-musk-deepfakes.html>; Daep et al., “Generative Propaganda.”
- 102 OpenAI, “Sora: Creating Video from Text,” *OpenAI*, 2015, <https://openai.com/index/sora/>.
- 103 Walker, “Slop Capitalism.”
- 104 Osmond Chia, “Sora: OpenAI Closes AI Video App and Cancels \$1bn Disney Deal,” March 25, 2026, <https://www.bbc.com/news/articles/c3w3e467ewqo>.
- 105 Connie Loizos, “Why OpenAI Really Shut down Sora,” *TechCrunch*, March 30, 2026, <https://techcrunch.com/2026/03/29/why-openai-really-shut-down-sora/>.
- 106 Nitish Pahwa, “Sam Altman Is Finally Admitting Something No One Else Wants To,” *Slate Magazine*, March 25, 2026, <https://slate.com/technology/2026/03/ai-openai-sam-altman-disney-sora-shutdown.html>.
- 107 Michelle Faverio and Emma Kikuchi, “Key Findings About How Americans View Artificial Intelligence,” *Pew Research Center*, March 12, 2026, <https://www.pewresearch.org/short-reads/2026/03/12/key-findings-about-how-americans-view-artificial-intelligence/>.
- 108 Brian Merchant, “Why the AI Backlash Has Turned Violent,” *Blood in the Machine*, April 14, 2026, <https://www.bloodinthemachine.com/p/why-the-ai-backlash-has-turned-violent>.
- 109 Anthony Ha, “Sora’s Shutdown Could Be a Reality Check Moment for AI Video,” *TechCrunch*, March 29, 2026, <https://techcrunch.com/2026/03/29/soras-shutdown-could-be-a-reality-check-moment-for-ai-video/>.
- 110 Kate Crawford, “Intensification - Eating the Future: The Metabolic Logic of AI Slop,” *E-Flux Architecture*, September 2025, <https://www.e-flux.com/architecture/intensification/6782975/eating-the-future-the-metabolic-logic-of-ai-slop/>; Max Read, “The Internet’s AI Slop Problem Is Only Going to Get Worse,” *NYMag*, September 30, 2024, <https://nymag.com/intelligencer/article/ai-generated-content-internet-online-slop-spam.html>.
- 111 Renee DiResta and Josh A. Goldstein, “How spammers and scammers leverage AI-generated images on Facebook for audience growth,” arXiv preprint arXiv:2403.12838 (2024).
- 112 Jason Koebler, “Facebook’s Algorithm Is Boosting AI Spam That Links to AI-Generated, Ad-Laden Click Farms,” *404 Media*, March 19, 2024, <https://www.404media.co/facebook-s-algorithm-is-boosting-ai-spam-that-links-to-ai-generated-ad-laden-click-farms/>.
- 113 Doctorow, *Enshittification*.
- 114 Tim Wu, *The Age of Extraction: How Tech Platforms Conquered the Economy and Threaten Our Future Prosperity* (New York: Alfred A. Knopf, 2025).
- 115 Aaron Short, “Can Artists Stop the AI Slop Machine?” *Hyperallergic*, September 17, 2025, <https://hyperallergic.com/can-artists-stop-the-ai-slop-machine/>.
- 116 Dan Milmo, “Thousands of Authors Publish ‘Empty’ Book in Protest over AI Using Their Work,” *The Guardian*, March 10, 2026, <https://www.theguardian.com/technology/2026/mar/10/thousands-authors-publish-empty-book-protest-ai-work-copyright>.
- 117 H. Holden Thorp, “Resisting AI Slop,” *Science* 391, no. 6780 (2026): 5, <https://doi.org/10.1126/science.aee8267>.
- 118 Matt Ferguson, “Animators Against Slop,” *Life in the Machine* (Substack), August 5, 2025, <https://lifeinthemachine.substack.com/p/animators-against-slop>.
- 119 Aria Groove, “Musicians Organize ‘NO AI SELLOUT’ Protests Against AI Licensing Deals.” *Electronic Groove*, May 19, 2026, <https://electronicgroove.com/musicians-organize-no-ai-sellout-protests-against-ai-licensing-deals/>.
- 120 Sofia Aira, “Head of Instagram Faces Backlash from Influencers Claiming It ‘Failed to Protect’ Them against ‘AI Slop,’” *Hello Partner*, January 2, 2026, <https://helloworldpartner.com/2026/01/02/head-of-instagram-faces-huge-backlash-from-influencers-claiming-the-platform-failed-to-protect-them-against-ai-slop-in-2025/>.
- 121 Irianna Lianaki Dedouli and Clare Stark, *Societal Implications of AI: Key Insights from Emerging Signals* (Paris: UNESCO, 2025).

- 122 Sasha Luccioni, Bruna Trevelin, and Margaret Mitchell, *The Environmental Impacts of AI - Policy Primer*, September 2024, [https://www.sashaluccioni.com/AI%20+%20Environment%20Primer%20\(Hugging%20Face\).pdf](https://www.sashaluccioni.com/AI%20+%20Environment%20Primer%20(Hugging%20Face).pdf).
- 123 Amanda Silberling, "YouTube Prepares Crackdown on Mass-Produced and Repetitive Videos as Concern over AI Slop Grows," *TechCrunch*, July 9, 2025, <https://techcrunch.com/2025/07/09/youtube-prepares-crackdown-on-mass-produced-and-repetitive-videos-as-concern-over-ai-slop-grows/>.
- 124 Neal Mohan, "From the CEO: What's Coming to YouTube in 2026," *YouTube Official Blog*, January 21, 2026, <https://blog.youtube/inside-youtube/the-future-of-youtube-2026/>.
- 125 The Tech Buzz, "Zuckerberg Ditches Metaverse, Goes All-In on AI Feeds," January 29, 2026, <https://www.techbuzz.ai/articles/zuckerberg-ditches-metaverse-goes-all-in-on-ai-feeds>.
- 126 Lauren Peterson, "Spotify Strengthens AI Protections for Artists, Songwriters, and Producers," Spotify Newsroom, September 25, 2025, <https://newsroom.spotify.com/2025-09-25/spotify-strengthens-ai-protections/>.
- 127 Dani Di Placido, "TikTok Tackles AI-Generated Slop with New User Controls," *Forbes*, November 26, 2025, <https://www.forbes.com/sites/danidiplacido/2025/11/26/tiktok-is-pushing-back-against-ai-slop/>.
- 128 Zack Zwiezen, "Xbox's New Boss Promises Not To Flood It With 'Soulless AI Slop,'" *Kotaku*, February 20, 2026, <https://kotaku.com/xboxes-new-boss-promises-not-to-flood-it-with-soulless-ai-slop-2000672087>.
- 129 Rithula Nisha, "How Reddit Is Cutting through the Noise of 'AI Slop,'" *AI Magazine*, February 11, 2026, <https://aimagazine.com/news/reddit-cutting-through-ai-slop>.
- 130 Kat Tenbarge, "AI Slop Is Ruining Reddit for Everyone," *WIRED*, December 5, 2025, <https://www.wired.com/story/ai-slop-is-ruining-reddit-for-everyone/>.
- 131 Scharon Harding, "Reddit Mods Are Fighting to Keep AI Slop off Subreddits. They Could Use Help," *Ars Technica*, February 17, 2025, <https://arstechnica.com/gadgets/2025/02/reddit-mods-are-fighting-to-keep-ai-slop-off-subreddits-they-could-use-help/>.
- 132 Karissa Bell, "Instagram Chief: AI Is So Ubiquitous 'It Will Be More Practical to Fingerprint Real Media than Fake Media,'" *Engadget*, December 31, 2025, <https://www.engadget.com/social-media/instagram-chief-ai-is-so-ubiquitous-it-will-be-more-practical-to-fingerprint-real-media-than-fake-media-202620080.html>.
- 133 Isabel O. Gallegos, Chen Shani, Weiyang Shi, Federico Bianchi, Izzy Benjamin Gainsburg, Dan Jurafsky, and Robb Willer, *Labeling AI-Generated Content May Not Change Its Persuasiveness*, July 30, 2025, <https://hai.stanford.edu/policy/labeling-ai-generated-content-may-not-change-its-persuasiveness>.
- 134 "Frontpage," *NewsGuard*, <https://www.newsguardtech.com/>.
- 135 "Frontpage," *The Trust Project*, 2017, <https://thetrustproject.org/>.
- 136 X, "Community Notes: A Collaborative Way to Add Helpful Context to Posts and Keep People Better Informed," Community Notes, <https://communitynotes.x.com/guide/en/about/introduction>.
- 137 "Proposed Code of Conduct for the Use of AI by Québec Political Parties," *Ivado*, 2025, <https://ivado.ca/en/repertoire-des-ressources/proposed-code-of-conduct-for-the-use-of-ai-by-quebec-political-parties/>.
- 138 "AI and Elections: The National Assembly Supports IVADO and CEIMIA's Code of Conduct," *Ivado*, 2026, <https://ivado.ca/en/2026/03/26/an-ai-code-of-conduct-for-political-parties-ivado-and-ceimia-launch-an-introductory-video/>.
- 139 "Lateral Reading," *Princeton University Library*, accessed May 22, 2026, <https://libguides.princeton.edu/medialiteracy/lateralreading>.
- 140 Brian Merchant, "Why the AI Backlash Has Turned Violent," *Blood in the Machine*, April 14, 2026, <https://www.bloodinthemachine.com/p/why-the-ai-backlash-has-turned-violent>.
- 141 Jasper Roe, Leon Furze, and Mike Perkins, "GenAI as Digital Plastic: Understanding Synthetic Media Through Critical AI Literacy," arXiv:2502.08249, Preprint, arXiv, February 12, 2025, <https://doi.org/10.48550/arXiv.2502.08249>.
- 142 S.E. Smith, "How to Disappear Completely," *The Verge*, December 18, 2024, <https://www.theverge.com/24321569/internet-decay-link-rot-web-archive-deleted-culture>.
- 143 Jasper Roe, Leon Furze, and Mike Perkins, "GenAI as Digital Plastic: Understanding Synthetic Media Through Critical AI Literacy," arXiv:2502.08249, Preprint, arXiv, February 12, 2025, <https://doi.org/10.48550/arXiv.2502.08249>.
- 144 Iliia Shumailov, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson, "The Curse of Recursion: Training on Generated Data Makes Models Forget," arXiv:2305.17493, preprint, arXiv, April 14, 2024, <https://doi.org/10.48550/arXiv.2305.17493>.
- 145 Iliia Shumailov, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson, "The Curse of Recursion: Training on Generated Data Makes Models Forget," arXiv:2305.17493, preprint, arXiv, April 14, 2024, <https://doi.org/10.48550/arXiv.2305.17493>.
- 146 Isabel O. Gallegos, Chen Shani, Weiyang Shi, Federico Bianchi, Izzy Benjamin Gainsburg, Dan Jurafsky, and Robb Willer, *Labeling AI-Generated Content May Not Change Its Persuasiveness*, July 30, 2025, <https://hai.stanford.edu/policy/labeling-ai-generated-content-may-not-change-its-persuasiveness>.

---

## ACKNOWLEDGMENTS

The Slop Salon was made possible by the Hewlett Foundation. The authors would like to thank Eli Sugarman and Amber Miller for their support of this work and participation in the Slop Salon.

The authors acknowledge with gratitude the provocateurs and interviewees whose presentations contributed to the workshop immeasurably: Renée DiResta, Sourojit Ghosh, Jason Koebler, Eryk Salvaggio, and Aidan Walker.

The following researchers and practitioners also graciously shared their time and insights to inform the paper assembled ahead of the workshop and distributed to all participants to inform the Slop Salon: Nina Lutz of the University of Washington, whose research was central to Sourojit Ghosh's presentation; Alexios Mantzarlis; Fabian Mosele; and Kate Starbird from the University of Washington.

We are deeply grateful to the participants and provocateurs of the Slop Salon for bringing intellectual generosity, rigor, candor, and humor to the discussion:

Ryan Beiermeister

Ben Brooks — Berkman Klein Center, Harvard

Patricia Cartes — CantinaAI

Elena Cryst — Stanford University

Shane Culloty — Stanford University, McCoy Family Center for Ethics in Society

Renée DiResta — Georgetown University

Alix Dunn — The Maybe

Sourojit Ghosh — University of Washington

Zach Graves — Foundation for American Innovation

Jason Koebler — 404 Media

Josh Lawson — OpenAI

Rob Reich — Stanford University

Eryk Salvaggio — University of Cambridge

Artemis Seaford

Sarah Shirazyan — Stanford University

Aidan Walker — Researcher

Andrea Wong — ByteDance

Camille François — Columbia University School of International and Public Affairs (SIPA) (Host)

Jeremie Ponak — Columbia University School of International and Public Affairs (SIPA) (Host)

Eli Sugarman — Hewlett Foundation (Host)

Jen Weedon — Columbia University School of International and Public Affairs (SIPA) (Host)

---

## **METHODOLOGY**

This report synthesizes insights from the Slop Salon, a closed-door convening of researchers, practitioners, and policymakers on AI slop held on March 5, 2026, at the Hewlett Foundation in Menlo Park, California. The workshop combined short provocations, facilitated discussions, and structured synthesis exercises, informed by pre-reads and some participant interviews. The analysis draws on organizers' notes, prepared remarks in the form of provocations, and referenced prior research. No verbatim transcript was produced; quotations are paraphrased or lightly edited for clarity.

The convening operated under the Chatham House Rule. Specific participant contributions are therefore anonymized, unless explicit permission for attribution was granted.

This report distinguishes between cited research, participant perspectives (with acknowledgment of the contributions of those who presented provocations, noted above), and authors' synthesis. These proceedings do not aim to establish consensus, but are intended to map key themes, areas of alignment, and points of divergence within a rapidly evolving field.



**IGP** Institute of  
Global Politics

**International Affairs Building**  
420 West 118th Street  
New York, NY 10027

**212-853-4720**  
[igp.sipa.columbia.edu](http://igp.sipa.columbia.edu)  
[igp@sipa.columbia.edu](mailto:igp@sipa.columbia.edu)