

Why AI Slop Matters, but Not Like That

SACHITA NISHAL, Northwestern University, USA

MARIJN SAX, University of Amsterdam, Netherlands

KIMON KIESLICH, University of Amsterdam, Netherlands and University of Hohenheim, Germany

This is a response to the paper “Why Slop Matters”. By offering both immanent and external critique, we argue that the authors’ reasoning neglects the socio-technical context of AI slop. Our paper presents an ethical and social science informed response that centers the debate on the social function and aesthetic value of AI slop. We conclude that AI slop is an important research subject but call for a contextual and culturally-grounded debate on the issue. To that end, we discuss some key elements of an agenda for future research on the phenomenon of AI slop.

CCS Concepts: • **Applied computing** → *Law, social and behavioral sciences; Arts and humanities*; • **Theory of computation**;

Additional Key Words and Phrases: AI Slop; Big Tech; Cultural Value; Communicative Meaning; Societal Impact

1 Introduction

In their provocative article ‘Why Slop Matters,’ Kommers et al. [16] argue that contrary to the common idea that AI-generated slop is a form of pollution, it serves a social function and has its own aesthetic value. We share their conviction that AI slop warrants rigorous study, but we argue that their framework, which is based on individual preference satisfaction, forecloses the social and political aspects of AI slop that matter most. To that end, in this piece, we deconstruct their argument through immanent and external critique, and propose a more critical research agenda for studying AI slop.

The satisfaction of individual desires through AI-generated, hyper-personalized content is a pervasive frame in both industry¹ and academic discourses on generative AI. Kommers et al.’s analysis exemplifies this approach too, since their argument for slop’s social function asks whether it satisfies users’ preferences (the kid-father example), and their aesthetic argument asks whether consumers might come to value it (the kitsch analogy). This frame carries implicit assumptions about the nature of people’s preferences, AI’s (social) costs, and aesthetic value, and thus it shapes what questions can and cannot be asked about AI slop.

For one, people are not always preference-satisfying consumers with static, self-derived choices who encounter content and proceed to then accept or reject it. They are also participants in communities through which preferences are formed, identities constituted, and meanings made collectively [20]. Rather than containing our inquiry into AI slop to whether it satisfies individual preferences, we can broaden to ask what kinds of communities, relationships, and conditions of meaning-making it assembles or forecloses, and what kinds of impacts it has on them.

¹For an example, see Mark Zuckerberg’s note titled *Building Personal Superintelligence*, in which he describes the goal of AI as enabling individuals “to direct it towards what they value in their own lives”.

Authors’ Contact Information: Sachita Nishal, nishal@u.northwestern.edu, Northwestern University, Evanston, USA; Marijn Sax, m.sax@uva.nl, University of Amsterdam, Amsterdam, Netherlands; Kimon Kieslich, k.kieslich@uva.nl, University of Amsterdam, Amsterdam, Netherlands and University of Hohenheim, Stuttgart, Germany.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, or to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 3068-8590/2026/6-ART

<https://doi.org/10.1145/3828758>

Because AI-generated slop has the potential to radically change—or: pollute?—the online environments where people post, share, and consume all kinds of information, we believe the terms on which we study and evaluate it matter greatly. The rest of this work offers our sociotechnical critiques and research agenda for studying AI slop, through which we hope to broaden the criteria through which the value of AI slop is evaluated.

2 The Arguments for Slop and Its Internal Inconsistencies

Kommers et al. define AI slop with three prototypical properties: (1) *Superficial Competence*. AI slop demonstrates a kind of competence that would take real expertise or skill to achieve without AI; this is belied by a lack of underlying substance, craft, or communicative intent. (2) *Asymmetric Effort*. AI slop is generated with a facile prompt and requires little or no effort of the kind that would be necessary to create such an output without AI. (3) *Mass Producibility*. AI slop is designed to work within a digital ecosystem of production and distribution to reach a mass audience.

From this, their argument for the “social function of slop” goes like this: people have weirdly idiosyncratic preferences they want to satisfy when searching for content. As a collective, humanity can produce astronomical amounts of content, meaning that there is no supply-side problem in terms of *quantity*. There is, however, a supply-side problem in terms of *quality*: with such a plethora of people holding so many different idiosyncratic preferences, producing content that meets each person’s specific desires is difficult at scale. This is where their core example comes in:

“For example, many children would find it entertaining to watch a video of their father dancing wildly in his underwear. The entertainment value comes mainly from the fact that it is one’s own father—not someone else’s. Crucially, there is an insuperable cost to generating this content under usual circumstances: getting your dad to do the dance. AI Slop can eliminate this bottleneck.”

This example is meant to show that AI slop can satisfy particular idiosyncratic preferences, and hence offers a social function. The video of the dancing father would *only* be interesting to the child in question, and without generative AI, the video probably would never exist.

Second, the argument for the aesthetic value of AI slop takes the form of a hypothetical: what if AI slop did, somehow, turn out to have aesthetic value because there are cyclical trends to be discovered on the demarcation between ‘low’ art and ‘real’ art? They refer to the history of kitsch, which was originally evaluated as cheap, formulaic, mechanical content, but was later—at least by some—reevaluated “as having obvious cultural value” [16]. The argument, as we understand it, is that if AI slop is sufficiently similar to ‘low’ culture, and if ‘low’ culture has been rehabilitated in the past, then AI slop may also be rehabilitated in the future.

Each of these arguments presents internal inconsistencies in light of how AI slop is defined, meaning that either the definition does not properly capture the phenomenon of AI slop, or the arguments do not engage with the phenomenon of AI slop. We break them down below.

The example of the kid generating a video of their father describes content created for private consumption based on a specific, idiosyncratic preference, but this runs counter to the definition of AI slop and its *mass producibility*. They argue that a key feature of AI slop is that it is typically mass-produced *with the intent of indiscriminately flooding social media feeds*², but this aspect of scale is entirely missing from the core example, which is meant to be an intuition pump for the article.

A second problem with the core example vis-à-vis the definition is the role of ‘communicative intent,’ given the *superficial competence* property of AI slop. But the example where the child has a specific wish to see their own father perform particular movements clearly does not display a lack of communicative intent. Quite to the contrary, the child wishes to express a very particular humorous message with the imagined scenario involving a person they know very well, which is precisely what makes the generated video all the more funny to the child.

²Consider, in fact, the tutorials widely available on TikTok on how to produce and monetize AI slop videos [17].

The example does not show the impersonal, mindless, intentionality-lacking characteristics which make AI slop “lack communicative intent.”

The authors also claim that AI slop serves a *social* function, but their core argument and example deal with the purely private consumption and preference satisfaction of one individual. The only way for Kommers et al.’s argument from the purely private to the social to work, is to argue that social utility is, by definition, nothing more than the sum of individual utility understood as simple preference satisfaction for individuals detached from their community and society. To be sure, such an argument, if made at all, would be very nihilistic.

Toward the argument for aesthetic value of AI slop, the analogy to ‘low’ culture is not sufficient to establish the aesthetic value of AI slop. For their indirect argument to work, the ‘low’ art they reference should be sufficiently similar to AI slop; i.e., the ‘low’ art should also be completely devoid of communicative intent and substance because if it is not, it also does not have explanatory value or justificatory power for AI slop. No such argument is provided, and in Section 4, we explain why this analogy might not work, through a discussion of the conditions that made the reappraisal of ‘low’ culture possible.

3 What Is Authentic Demand?

We now turn to the external critique, starting with the discussion on *demand*. Kommers et al. write that “People want more content than humans can supply”. The authors take the demand, that allegedly happens to exist, as natural starting point of the analysis. But demand doesn’t emerge naturally, it is created. Big Tech companies design their (manipulative [18, 19]) interfaces to boost engagement in an attempt to manufacture ongoing demand for their services [8, 14, 23]. Some also point to “a program of systematic desocialization” [13], where users are trained to relate to others as content to be consumed constantly, rather than people to connect with. AI slop partly serves the need for constant consumption that this system creates, and calling that a social function unfairly legitimizes what is better understood as a harmful restructuring of how people relate to culture and to each other.

So, even if it were true that people want more content than humans can supply, there is an important normative question begging to be asked: Should demand that is manufactured for the sole purpose of promoting the interests of Big Tech be treated as equally important as *authentic* demand? Moreover, by treating the preferences of users as an unproblematic starting point one also makes end-users responsible for how they are using platforms, and, in the end, how they produce and consume AI slop. This is a dangerously one-sided argument that ignores the strategies of platform to *create* this behavior; and potential demand in AI slop. In sum, the phenomenon of demand cannot be treated as an unproblematic starting point of AI slop analyses, but should itself be part of critical analyses.

4 What Is the Aesthetic Value of AI Slop?

Kommers et al. draw on the history of kitsch to argue that AI slop, like other ‘low’ cultural forms initially dismissed by critics, may turn out to have aesthetic value. However, this analogy only works if kitsch (or any ‘low’ culture) and AI slop are similar not just in how they are produced but in how they function culturally. In fact, the observation that critics have sometimes been wrong about ‘low’ culture would also support rehabilitating phenomena like spam and clickbait. The argument needs a positive account of why AI slop belongs to the class of things that were rehabilitated rather than the class that were not.

As Greenberg [12] argues, kitsch is parasitic on ‘high’ culture to derive its forms and standards. However, it is this relational quality that made kitsch’s later reappraisal possible. Călinescu [3] documents that critics dismissed kitsch on the grounds that its value seemed *merely* social—audiences consumed kitsch to chase social status, rather than for intrinsic aesthetic merit—but this dismissal undervalued the social dimensions of kitsch, i.e., that it spoke to the social aspirations of its audiences and generated shared interpretive frameworks and discourses around it [10].

AI slop, as defined, has a veneer of quality without aspiration toward it, meaning that the conditions which made the reappraisal of kitsch possible are undermined by slop. First, kitsch was in dialogue with recognizable aesthetic standards, where slop is not, since it lacks communicative intent by definition. Second, kitsch circulated as shared cultural objects, but hyper-personalized slop (e.g., the dancing father video) is designed for individual consumption, which may thwart the formation of shared interpretive frameworks. Third, kitsch accumulated its cultural significance through sustained engagement from consumers and critics over time, but the zero-effort, mass-produced nature of AI slop work against this. We can see these conditions at work when we consider other ‘low’ cultures now viewed more favorably by critics. Fanfiction attracts scholarly attention because the conditions for reappraisal are present. Communities are built around shared aesthetic frameworks (genre conventions, shared canon, common tropes), sustained through reciprocal exchange, like gifting and feedback [4, 9].

There is one more important aspect about the analogy to ‘low’ culture, which is that while kitsch was parasitic on high culture, it did not threaten to make cultural artifacts harder to create or sustain. AI slop at scale, as Simpson and Semaan [21] document, degrades the infrastructures that creative communities depend on, and can crowd out human creative work [22]. Our argument here is that the analogy to ‘low’ culture is not sufficient to establish the aesthetic value of AI slop, and in fact, a more urgent question pertains to how AI slop is impacting conditions (e.g., the infrastructures, communities, shared frameworks) through which aesthetic value is generated and sustained by people.

5 The (Hidden) Costs of AI Slop

According to the authors, one core argument for the appeal of AI slop is that “it offers a near-zero cost path to achieving one’s goals.” Much like their arguments based on the phenomenon of demand, here the authors also treat the phenomenon of cost as an unproblematic notion. They only focus on costs *for the individual prompting* and only consider cost in a narrow economic sense. We present four considerations on the *social* costs. These critical considerations also apply to generative AI more generally, which is unsurprising given the close connection between the phenomena of generative AI and AI slop. We do, however, think that there is a specific AI slop angle to all considerations. Built into the concept of AI slop is an element of quantity; it can be generated easily, quickly, in enormous quantities. Its ability to *flood* online spaces helps explain specific risks and harms, making it a case of a difference in degree becoming a difference in kind. So, even though these considerations are also relevant to discussions on generative AI, they are especially relevant for the discussion of AI slop.

Power. Both Generative AI companies and social media platforms stimulate and incentivize the production and circulation of AI slop. As users, both as consumers and producers of AI slop, keep engaging with Big Tech companies, they help consolidate their market power [2]. And nowadays with market power comes real political power, as the Trump regime has shown [5]. Participating in the AI slop economy thus also has potential political consequences.

Environment. Creating AI slop is compute intensive. And as the authors explain, AI slop is especially attractive because it can be used to satisfy infinitely many idiosyncratic personal preferences with highly personalized videos. So, promoting the social value of AI slop comes with an environmental footprint [7, 15]. Ultimately, the normative question arises if the mass-production of AI slop justifies an increasing amount of environmental harm.

Data Privacy. AI slop can also come with privacy costs. In the authors’ own example, a video is generated of an existing father in his underwear. It is telling that the example does not even consider the question whether this father wants this to happen to him. Only the immediate preference satisfaction of the prompter is considered. The example clearly shows issues around *consent* of people who are shown in AI slop output [11].

(Stochastic) Manipulation. The ability to generate enormous amounts of highly specific content at almost no cost *to the prompter* comes with a potentially massive *collective* cost: AI slop for online campaigns that can

manipulate political sentiments and possibly even elections [6]. Even if individual pieces of slop do not manipulate individual users, AI slop's ability to flood online spaces may have nontrivial impacts on collective behavior [1].

6 Conclusion

Kommers et al. [16] argue that AI slop matters. We agree that AI slop is a phenomenon that deserves serious scrutiny, but in this short paper we explained why AI slop matters for reasons different than those presented by Kommers and colleagues. Where their argument depends on several problematically uncritical assumptions vis-à-vis the demand for – and costs and consumption of – AI slop, we argue that research on AI slop must engage with ethical and social science literature to substantiate claims about its societal function and aesthetic value. We call AI slop researchers to adjust or extend their research focus to the following research areas:

Consumption of AI slop. Research on AI slop should engage with – normatively and empirically – what constitutes AI slop demand. Social science researchers, for instance, could conduct qualitative focus groups or interview studies to sketch out rationales of why people are (not) consuming AI slop. Quantitative survey studies could provide statistical evidence in showcasing general consumption patterns and describe conditions under which AI slop is (not) consumed. These factors can include, among others, AI literacy, acts of resistance, authenticity perceptions, and risk-benefit trade-off perceptions. Normative and critical studies could inform theory building and specifically interrogate questions of public values and authenticity in AI slop consumption.

AI slop and art. While we discussed arguments against the definition of aesthetic value brought forward by Kommers et al., it remains an open question in how far AI slop can create any form of aesthetic value. This line of research can be further explored by cultural and humanities scholars as well as artists. We propose that research should look into conditions under which collective sense-making and cultural values can be created around, or with the help of AI slop.

Regulatory interventions towards AI slop. AI slop comes with a plethora of (hidden) societal costs. If we want to design effective regulatory interventions to curtail the harms of AI slop, those societal costs should be part of the regulatory discourse. Academic research on AI slop that builds on critical perspectives on, for instance, political economy, environmentalism, data privacy, and manipulation can inform that regulatory discourse. A richer understanding of the societal costs associated with AI slop will also help regulators explore the potential of important recent legislation like the European Union's AI Act, Digital Services Act, or the General Data Protection Regulation to address AI slop harms.

Author Contributions

All authors contributed equally to this research.

Declaration of Interest

The authors declare no conflict of interest.

Funding

Marijn Sax was funded by the Dutch Research Council, grant number VI.Veni.221F.006. Kimon Kieslich was funded by the German Academic Exchange Service (DAAD) with funds from the Federal Ministry of Research, Technology and Space (BMFTR).

References

- [1] Claire Benn and Seth Lazar. 2022. What's wrong with automated influence. *Canadian Journal of Philosophy* 52, 1 (2022), 125–148.
- [2] Kean Birch and 'Damola Adediji. 2025. Undermining competition, undermining markets? Implications of Big Tech and digital personal data for competition policy. *Big Data & Society* 12, 1 (2025), 20539517241311584.

- [3] Matei Călinescu. 1987. *Five Faces of Modernity: Modernism, Avant-Garde, Decadence, Kitsch, Postmodernism*. Duke University Press, Durham.
- [4] Ruijia Cheng and Jenna Frens. 2022. Feedback Exchange and Online Affinity: A Case Study of Online Fanfiction Writers. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2 (Nov. 2022), 402:1–402:29. doi:10.1145/3555127
- [5] Mark Coeckelbergh. 2026. Technofascism: AI, Big Tech, and the new authoritarianism. *AI & SOCIETY* (2026), 1–14.
- [6] European Commission. 2024. Commission opens formal proceedings against TikTok on election risks under the Digital Services Act. *press release* (2024). https://ec.europa.eu/commission/presscorner/detail/en/ip_24_6487
- [7] Kate Crawford. 2021. *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale University Press.
- [8] Pablo González de la Torre, Marta Pérez-Verdugo, and Xabier E Barandiaran. 2025. Attention is all they need: cognitive science and the (techno) political economy of attention in humans and machines. *AI & SOCIETY* (2025), 1–17.
- [9] Casey Fiesler. 2021. Gift Logic: Labors of Love Flourish Online under Fandom’s Social Norms.
- [10] Clifford Geertz. 2009. *The Interpretation of Cultures: Selected Essays* (repr. ed.). Basic Books, New York, NY.
- [11] Abenezer Golda, Kidus Mekonen, Amit Pandey, Anushka Singh, Vikas Hassija, Vinay Chamola, and Biplab Sikdar. 2024. Privacy and security concerns in generative AI: A comprehensive survey. *Ieee Access* 12 (2024), 48126–48144.
- [12] Clement Greenberg. 1939. Avant-Garde and Kitsch. *Partisan Review* 6, 5 (1939), 34–49.
- [13] John Herrman. 2026. Instagram Reels and the New Era of Desocialized Media. *Intelligencer* (Jan. 2026).
- [14] Marcello Ienca. 2023. On artificial intelligence and manipulation. *Topoi* 42, 3 (2023), 833–842.
- [15] Nidhal Jegham, Marwan Abdelatti, Chan Young Koh, Lassad Elmoubarki, and Abdeltawab Hendawi. 2025. How hungry is ai? benchmarking energy, water, and carbon footprint of llm inference. *arXiv preprint arXiv:2505.09598* (2025).
- [16] Cody Kommers, Eamon Duede, Julia Gordon, Ari Holtzman, Tess McNulty, Spencer Stewart, Lindsay Thomas, Richard Jean So, and Hoyt Long. 2025. Why Slop Matters. *ACM AI Letters* (2025).
- [17] pham. 2026. How I use AI to make \$10,000 a month... Here’s how. TikTok. <https://vm.tiktok.com/ZGdu6go2D/>
- [18] Marijn Sax. 2021. *Between Empowerment and Manipulation: The Ethics and Regulation of For-Profit Health Apps*. Kluwer Law International BV.
- [19] Marijn Sax. 2021. Optimization of what? For-profit health apps as manipulative digital environments: M. Sax. *Ethics and Information Technology* 23, 3 (2021), 345–361.
- [20] Amartya Sen. 1999. *Development as Freedom*. Oxford University Press, Oxford New york.
- [21] Ellen Simpson and Bryan Semaan. 2025. The Enshittification of the Creative Internet. *Proc. ACM Hum.-Comput. Interact.* 9, 7 (Oct. 2025), CSCW489:1–CSCW489:24. doi:10.1145/3757670
- [22] Charlie Warzel. 2025. A Tool That Crushes Creativity. *The Atlantic* (Oct. 2025).
- [23] Shoshana Zuboff. 2023. The age of surveillance capitalism. In *Social theory re-wired*. Routledge, 203–213.