

What is AI slop? An Interdisciplinary Approach

Gustavo Araujo Costa^{1,*}

¹London Interdisciplinary School, 20-30 Whitechapel Road, E1 1EW, London, United Kingdom

*gustavo.costa@lis.ac.uk

Abstract

This project explores the emerging phenomenon of AI slop—low-effort, low-value AI-generated content that clutters digital information environments. While not necessarily false or malicious, slop is often meaningless, inauthentic, and aesthetically unappealing. Despite its growing visibility, the concept remains largely undefined in scholarly literature. To address this gap, I conducted qualitative (thematic analysis) and quantitative (topic modelling) analyses of sources referenced in the Wikipedia page on AI slop. I also utilised a large language model-based method, which offered comparable performance while being considerably more efficient. Thematic analysis reveals four core dimensions of the discourse around slop: efficiency and productivity, danger and discernment, epistemic crisis, and engagement economy. Slop is typically described as inaccurate, inauthentic, mediocre, surreal, distorted, generic, pointless, or kitsch—and is often metaphorically framed as flooding, contamination, infestation, or rubbish. Four main types of AI slop were identified: social media slop, corporate slop, submission slop, and slop news. This study also introduces the SlopNews dataset, used to analyse the statistical features of slop content. Compared to non-slop news, slop news is less complex, less varied, and tends to express more positive sentiment. These findings offer a foundation for understanding the dynamics and risks of low-quality AI content. They may inform educators, editors, researchers, and platform moderators in managing slop and safeguarding trust in digital communication.

Key words: artificial intelligence, AI slop, information environment, misinformation, thematic analysis, topic modelling

1. Introduction

After the release of ChatGPT, usage of generative AI saw a sharp increase. With this, there has also been an increase in unwanted and low-quality AI generated content which has been informally termed AI slop. Slop is similar to terms such as *fake news* and *enshittification* which attempt to name a sense of increasingly pervasive ‘information disorder’ (Wardle and Derakhshan, 2017).

Slop fits into the category of “information disorder” (Wardle and Derakhshan, 2017) alongside misinformation and disinformation. This is happening in an information environment where there are concerns about misinformation, lack of trust in institutions, mainstreaming of conspiracy theories and erosion of shared reality. By flooding information environments with inauthentic, incoherent, or manipulative content, slop worsens epistemic instability, undermines the utility of digital systems,

and places new burdens on already overstretched institutions.

Despite its growing cultural visibility, AI slop¹ lacks a formal academic definition. This study aims to fill that gap by exploring the meanings and practices associated with AI slop, as well as analysing slop content itself.

1.1. The problem of slop

As a general phenomenon, slop is not new: humans have long generated “authentic” meaningless content for profit. Earlier forms include practices such as content farms, paper mills, clickbait, and search engine optimisation (SEO) manipulation. However, AI slop is novel in that recent developments in artificial intelligence (AI) technology have

¹ “Slop” and “AI slop” are used interchangeably in this article.

dramatically lowered the cost and effort required to create slop.

Consequences of this shift are already being felt in systems that rely on good-faith submissions. Organizations are reporting a deluge of low-quality synthetic content that they must manage, such as in the case of the maintainer of the widely-used open-source coding package curl (Purdy, 2025), as well as Clarkesworld magazine, a science fiction publisher (Clarke, 2023).

Slop also impacts systems which incentivise quantity over quality—such as engagement-driven social media and SEO. With increasing pressure to automate and scale content creation, slop has become a structural by-product of platform incentives and algorithmic visibility.

1.2. Methods and contributions

This study draws on two different data sources. First, the content of the references made in the Wikipedia entry for AI slop was analysed. For this purpose, three complementary methods were employed: thematic analysis (manual identification of themes), topic modelling (computational theme extraction), and LLM-based topic modelling. Apart from themes, the qualities, metaphors and types of slop are also analysed (see Figure 1).

Second, I introduce *SlopNews*, a dataset of slop news. It was constructed by scraping websites flagged by information reliability organizations as producing slop. These were combined with non-slop sources from publicly-available news datasets. Linguistic metrics were obtained and statistically analysed, finding that slop news exhibits less linguistic variation, less complexity and higher positive sentiment in comparison with non-slop news.

This work addresses pressing needs in academic integrity, content moderation, and information reliability. It provides theoretical and empirical foundations for those developing tools to curate the digital information landscape. While the subjective and evolving nature of AI slop presents some challenges, a rigorous, interdisciplinary approach is essential to mitigate its harmful effects and preserve epistemic integrity.

2. Related work

To better understand AI slop, I analyse related concepts and work in several disciplines—computer science, philosophy, psychology, and media studies, as well as the origins of the term itself.

Although the word *slop* was used to refer to low-quality writing as far back as the 19th century, the term’s usage has seen a remarkable 334% increase in usage in 2024 (Oxford University Press, 2024). Coinage of its current AI-related usage has been attributed to Willison (2024). Two key events are related to its popularisation: the creation of its Wikipedia page and the publication of an article in the New York Times² Aleksic (2024).

Slop was on the shortlist for the Oxford Word of the Year 2024³, where it is defined as ‘writing, or other content generated using artificial intelligence, shared and distributed online in an indiscriminate or intrusive way, and characterized as being of low quality, inauthentic, or inaccurate.’

While no academic papers on AI slop have been published, there are two preprints: an unrelated technical advance in AI writing quality (Chakrabarty et al., 2025) and a discussion of AI-generated propaganda, which they call *slopaganda* (Klincewicz et al., 2025). While the scientific literature is sparse, there are numerous op-eds, blog posts and similar media which extensively discuss the subject.

The most conspicuous kind of AI slop is *social media slop*—the most notorious example being “Shrimp Jesus”. This is surreal AI-generated content spread on Facebook for driving engagement (Placido, 2024). Other examples include fake Amazon product listings, (Tangermann, 2024), fake summary and biography e-books (Read, 2024), misleading AI-generated advice (McMahon and Kleinman, 2024), and incorrect scientific illustrations (Bailey, 2024).

Slop is similar to other forms of problematic information:

- **Disinformation:** false information spread with malicious intent (Wardle and Derakhshan, 2017).
- **Misinformation:** false information spread inadvertently (Wardle and Derakhshan, 2017).
- **Hallucination:** AI-generated content which seems plausible but is factually incorrect (Huang et al., 2025).
- **Rumour:** information which is unverified; a debunked rumour can be classified as misinformation or disinformation depending on the intent of the source (Zubiaga et al., 2019).

² <https://www.nytimes.com/2024/06/11/style/ai-search-slop.html>

³ <https://corp.oup.com/word-of-the-year/>

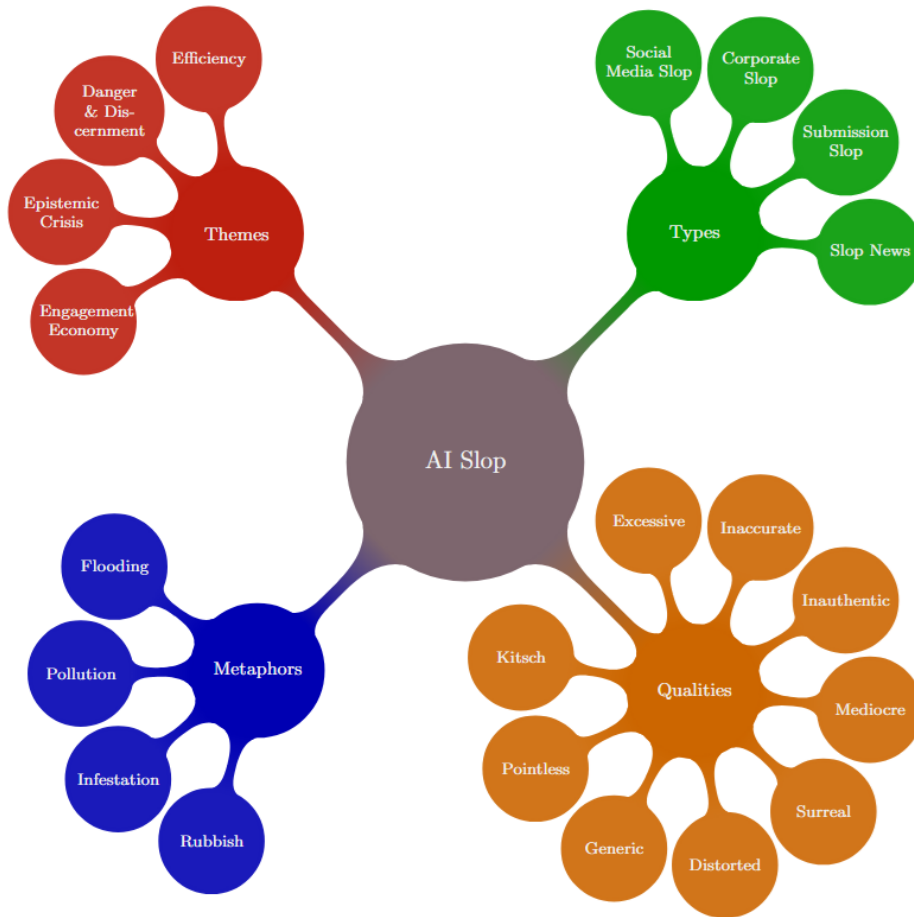


Fig. 1: AI slop mind map containing themes, types, metaphors and qualities of AI slop.

- **Fake news:** news which are false and public; unlike rumours, they always refer to facts which can be verified (Zubiaga et al., 2019).
- **Spam:** Intrusive messages typically associated with scams.
- **Bullshit:** communication characterised by indifference to truth (Frankfurt, 2009).

Unlike misinformation and hallucination, slop is not necessarily false. Unlike spam and disinformation, it is not necessarily malicious. Its closest parallels are spam (with its emphasis on quantity) and bullshit (with its emphasis on plausibility). Slop has been likened to a metaphorical ‘distributed denial-of-service’ (DDoS) attack, where an excess of

useless requests prevents a system from processing legitimate ones (Purdy, 2025).

2.1. Classification of rumours and machine-generated content

In natural language processing (NLP), the classification of problematic information (such as rumours or misinformation) typically involves linguistic pattern analysis, statistical methods, and fact verification (Tang et al., 2023; Aich et al., 2022). The highest performance is offered by fine-tuned deep learning models, which don’t require manual feature creation (King et al., 2024; Chen and Shu, 2024; Doughman et al., 2024; Chakrabarty et al., 2025). For example, Full Fact, a fact-checking organization, uses a BERT-based deep learning classifier to prioritise rumours

which are likely to spread.⁴ Another newer and promising method involves simply prompting LLMs for analysis, which yields equivalent performance (Chen and Shu, 2024).

A similar task is the classification of machine-generated content. Pangram is deep learning based AI detector; the developers offer a detailed technical report (Emi and Spero, 2024). Widely-used commercial AI detectors such as GPTZero⁵ and Originality.ai⁶ employ various features to distinguish human from machine-generated text in addition to deep learning classification. These metrics include perplexity, burstiness, readability, and patterns in punctuation and word frequency. This feature-based approach is also employed by scientific integrity tools such as the Problematic Paper Screener (Labbé et al., 2025) and Wiley’s papermill detector (Flanagan et al., 2024) which identify “fingerprints” (which are telltale phrases and artifacts left by processing such as paraphrasing software) in questionable scientific papers.

It remains unclear whether detecting slop is significantly different from detecting AI-generated content. Given that slop is defined by its lack of effort, AI detectors should be able to filter it effectively. However, this would penalize high-quality AI-generated content, which is not a desirable outcome. Additionally, AI-detection tools are notoriously unreliable (Weber-Wulff et al., 2023; Sadasivan et al., 2023; Fòrlini, 2023) and lack interpretability, which is essential for algorithmic accountability (Saphra and Wiegrefe, 2024; Templeton et al., 2024; Kapoor et al., 2024). Although interpretability methods for deep learning models exist, these often yield suboptimal or misleading results (Heap et al., 2025; Gonen and Goldberg, 2019). This suggests that a simple, feature-based slop classifier could be a promising direction for future research.

NLP classification tasks can be broadly divided into those with objectively verifiable answers (e.g. fake news, AI-generated text) and those where ground truth is ambiguous or subjective (e.g. rumours, hate speech). AI slop falls into the latter category, as its fundamental characteristics are not grounded in objective criteria. However, even in

cases where factual verification is possible, fact-checking is labour-intensive and time-consuming. As a result, fake news is often addressed by labelling the credibility of the source rather than analysing the content itself (Peters, 2022). Another approach involves using machine-learning algorithms to quickly determine which rumours are likely to spread, instead of using them for the actual classification (Zubiaga et al., 2019).

2.2. AI risk taxonomies

There are many taxonomies of AI risk currently being developed.⁷ AI slop does not fit neatly into established risk categories, since it neither causes physical harm nor is obviously toxic. At an individual level, its effect is only inconvenience. Taken as a collective phenomenon, however, it poses psychological, economic and political risks by undermining trust in the information environment.

In the MIT AI Risk Taxonomy (Slattery et al., 2025), slop is most closely aligned with subdomain 3.2: “pollution of information ecosystem and loss of consensus reality”. Although this category is primarily concerned with issues like personalization and filter bubbles (also known as echo chambers), it remains the most fitting descriptor for the risks associated with the proliferation of slop.

Another relevant taxonomy was developed by the Center for Security and Emerging Technology (Hoffmann and Frase, 2023). In this framework, we can categorise slop as an intangible issue of detrimental content and/or democratic norms. Here, *intangible* refers to phenomena that are difficult to observe or lack a direct material manifestation, while *issue* denotes an ongoing concern or a possible occurrence rather than a specific event. Slop falls under detrimental content through its association with misinformation, and it relates to democratic norms via its contribution to information disorder.

2.3. Philosophy and psychology

Philosophical discourse has extensively explored ideas of the “loss of humanity” brought about by technological advancement: from Plato’s concerns about writing and memory (Derrida, 2020), to Heidegger’s critique of modern technology (Graña,

⁴ <https://fullfact.org/ai/about>

⁵ <https://gptzero.me/>

⁶ <https://originality.ai/>

⁷ The MIT AI Risk Taxonomy (Slattery et al., 2025), for example, was constructed using information from 65 other taxonomies.

2016), and now contemporary philosophers of different traditions (Han, 2015; Krenak, 2020).⁸

However, the defining feature of slop is not its artificial origin, but its poor quality. This is similar to hallucination, where the issue is not artificiality, but inaccuracy. Artificial content which is truthful and high-quality is rarely considered slop, as it is not recognised as artificial or inauthentic at all. This also poses a challenge for the collection of high-quality AI-generated content, as it is indistinguishable from high-quality human-made content.

This distinction can be observed in the Google Search’s guidance for developers, where they state their interest on “the quality of content, rather than how content is produced” (Sullivan and Nelson, 2023). This highlights the importance of three different interlocked parts: the origin of the content, the content itself, and the interpretation of the reader.

With the advent of AI as a general-purpose technology, novel interdisciplinary work is emerging. One example is the exploration of LLMs in the context of the philosophical concept of bullshit (Hicks et al., 2024; Gorrieri, 2024; Trevisan et al., 2024). Bullshit is defined as discourse indifferent to its truth-value; instead of prioritising accuracy, it is created for some other objective, such as plausibility or persuasion (Frankfurt, 2009; Cohen, 2002; Easwaran, 2023). This is distinct from lying, which involves intentionally stating a known falsehood, and thus having an interest in its truth-value.

Instead of “hallucinating”, AI systems can be better understood to be bullshitting – generating plausible-sounding text without regard for truth (Hicks et al., 2024). In this sense, generating slop is also bullshitting: communicating something not for its informational value but for its effect (e.g. engagement on social media).

The traditional definition of bullshit centres on the intentions of the speaker. There is an alternative definition, which defines it exclusively through its contents. According to this view, bullshit is defined by being ‘unclarifiable’ (Cohen, 2002). This is the definition most commonly used in psychology, particularly in studies of susceptibility to “pseudoprofound bullshit” (Pennycook et al., 2015) and fake news. This approach deliberately sidesteps the philosophical emphasis on the speaker

and instead focuses on a property of the content itself. In this sense, it parallels Google’s policies which prioritise the quality of the content over its authorship.

Along with the relationship between author and content, we can also consider the one between content and consumer. In the context of the subjective judgment of “quality”, it is important to recognise that beauty is in the eye of the beholder—what one person perceives as inauthentic, another might find meaningful (Dalton, 2016). This emphasises the role that normativity and context play in subjective labels such as bullshit and slop. This can be observed in the motivations for the creation of the different definitions of bullshit: the original intent-based definition was created to describe rhetoric in politics and advertising; the content-based view was developed to expose obscurity in French Marxism; and psychological approaches were created to criticise new-age spirituality. In their attempts to objectively define the concept (by moving away from the origin and the interpretation towards analysing the content itself), the researchers ultimately only make their biases more prominent.

The psychological study of epistemically suspect beliefs – such as bullshit and fake news – draws heavily on the heuristics-and-biases tradition from behavioural economics (Tversky and Kahneman, 1974). This cognitive perspective is central to the analysis of slopaganda in Klineciewicz et al. (2025), where susceptibility to slop is related to phenomena such as the illusory truth effect (where repeated exposure increases perceived truthfulness), confirmation bias (the tendency to favour information that confirms existing beliefs), and information overload (where emotionally charged content captures limited attention resources). This body of literature, however, tends to frame bias as the exploitation of vulnerabilities in human rationality (for example, see Pennycook and Rand, 2019; Muzumdar et al., 2025). In this way, it neglects the performative and emotional dimensions of communication—treating misinformation as a failure in the transmission of information rather than a socially and affectively embedded practice (Wardle and Derakhshan, 2017).

2.4. Information environment and epistemic crisis

Along with considering slop in its technological and psychological dimensions, we can also consider the structural and socioeconomic conditions that

⁸ Although there are alternative accounts, such as in science and technology studies (Haraway, 2015)

incentivize slop production, such as algorithmic visibility, low-cost scaling, and platform monetization.

AI slop operates similar to disinformation (Wardle and Derakhshan, 2017) in that both consist of repetitive and low-value content. This is remarkably similar to the political strategy articulated by Steve Bannon: “flood the zone with shit” (Illing, 2020). The goal is not to convince, but to exhaust and confuse. While slop does not have these as its main purposes, they happen as consequences of the prevalence of slop in an informational environment.

There is substantial public anxiety surrounding misinformation (Cellan-Jones, 2017; World Economic Forum, 2024). In the United States, adults perceived made-up news as a more pressing issue than violent crime, climate change, racism, illegal immigration, terrorism, or sexism (Barthel et al., 2016). Ironically, information about misinformation is often misleading. Evidence suggests the problem posed by fake news is overstated. The values cited above reflect the public *perception* of the issue, not its empirical prevalence. The average American’s exposure to fake news has been found to be low, at around 0.15% (Allen et al., 2020). A study by the UK Centre for Emerging Technology and Security (CETaS) found no evidence that AI disinformation or deepfakes impacted UK, French or European elections results in 2024 (Stockwell, 2024). In most cases, social media actually *increases* exposure to alternative viewpoints, while only about 8% of the population exist in ideological “echo chambers” (Arguedas et al., 2022). This indicates that the core danger is not necessarily an information environment “poisoned” with falsehoods, but the pervasive fear and corrosion of public trust that the perception of such an environment creates.

Organizations tackling this issue, however, are still subject to the pressures that generate misinformation. Vendors of technical solutions and academics alike lean into the perceived risk of “information disorder”. By aligning with grand narratives about AI’s capacity to either save or destroy humanity, they mystify the technology, which helps maintain its allure and its perceived threat (Tully et al., 2025) – thus maintaining their own relevance in an information environment characterised by scarcity of attention and abundance of noise.

Market incentives – from click-driven ad revenue to “publish or perish” academia – fuel the production of slop by motivating agents to create a high volume

of low value content (Knibbs, 2024a; Labbé et al., 2025). In this way, this is not an issue of psychological susceptibility, but of the consequences of the systems with which people are compelled to engage.

In this work, I use the term “epistemic crisis” to refer to these generalised concerns about misinformation (e.g. vaccine hesitancy during the pandemic), disinformation (e.g. Russian propaganda about Ukraine), conspiracy theories (e.g. QAnon during the US election), pseudoscience (e.g. hydroxychloroquine as a treatment for Covid-19), and related phenomena. This dynamic is particularly evident on platforms like Facebook, which has a troubled history of managing political influence (Wynn-Williams, 2025).

Digital platforms play a large role in the dynamics of the epistemic crisis. One interesting concept to understand their functioning is *enshittification* (Doctorow, 2023). It refers to the process by which platforms (e.g. Amazon, TikTok) compromise the quality of their service. First, they attract users with high-quality services while operating at a loss. Once users are locked in, the platform shifts its focus to extracting maximum value for shareholders, systematically degrading the user experience in the process.

3. Thematic analysis

3.1. Methodology

To examine how the term “AI slop” is used, I employ thematic analysis—a widely adopted qualitative method for identifying patterns of meaning across a dataset through manual coding and interpretation (Squires, 2023; Braun and Clarke, 2006). This approach was selected for its flexibility and depth, which are well suited to the exploratory study of an emerging and loosely defined concept.

Following Braun and Clarke’s framework, I explicitly outline the project’s epistemological assumptions. This project attempts to be: inductive (themes arise from the data instead of an existing theory), semantic (focus on explicit data instead of latent meaning), and constructionist (focus on sociotechnical and political aspects of the themes instead of particular experiences). Theme prevalence is assessed in terms of how many distinct items reflect a theme, rather than how often the theme appears overall.

In qualitative research, subjectivity is not treated as a bias to be eliminated but as an integral part of the interpretive process. As such, I acknowledge my position as a researcher who regularly uses AI

tools, but understands them as “normal technology” (Narayanan and Kapoor, 2025)—neither utopian nor dystopian. Viewing all technologies as embedded in social systems, my goal is to understand how slop is used to fulfil a sociotechnical need.

3.1.1. Data

The dataset consists of 34 articles cited by the Wikipedia page on AI slop.⁹ These were selected after discarding tangential academic preprints and repeated articles on the same topics. Content was collected using Trafilatura (Barbarese, 2021) and coded using Taguette (Rampin and Rampin, 2021). While I refrain from publishing the content of the dataset due to copyright concerns, a complete list of sources is included in Appendix A.

This dataset has several limitations: it represents a temporal snapshot and does not capture recent developments in the fast-moving phenomenon of AI slop; it reflects the editorial decisions of the Wikipedia contributors (e.g. focusing on topics such as video games); the articles are subject to the same attention-seeking dynamics of the engagement economy, which leads publishers to dramatise their narratives.

As it is common for slop to be used as a catch-all term for dissatisfaction with AI-produced content, I deliberately did not distinguish between mentions of slop and more general commentary on AI.

3.2. Results

Four major themes were identified within the dataset:

- *Efficiency and productivity*: Slop is generated due to its low cost, speed, ease of creation, and scalability. This applies to corporations aiming to cut costs and individuals seeking side income. These indicate a significant incoming shift in several industries, which is accompanied by concerns about the automation of creative labour.
- *Danger and discernment*: The excess of slop overwhelms information systems and creates the need for its filtering and moderation. This is especially challenging since slop does not strictly violate platform policies. This is threatening both to institutions and to individual users, which might unknowingly be exposed to false information.
- *Epistemic crisis*: Slop is related to moral anxieties about political disinformation, the “post-truth” environment, and the wider breakdown of shared reality and trust in institutions.
- *Engagement economy*: Platform incentives (such as engagement, virality, and ad revenue), which are ultimately about profit, are the primary drivers of slop. Individuals in low-income regions are especially incentivised to engage in its creation.

Each theme is described below. This section references the articles’ dataset ID (e.g. A5 for article 5) as shown in Appendix A.

3.2.1. Efficiency and productivity

It’s easy to learn, almost zero cost, and can be done any time at home with just a phone. (A12)

Ryanair and Samsung are jumping on the trend and seeing six-figure engagement spikes — all while barely lifting a production finger. (A21)

This theme is related to the reasons why people choose to produce slop. Subthemes include: *cheap, easy, fast, scalable, consumable, productive, and transformation*. These qualities appeal to both corporations seeking cost-cutting and individuals (especially in low-income regions) seeking side income.

Generative AI increases productivity by expediting content creation—usually in visual media. One “slopper” comments that his method consists of “producing far more content than people who don’t use AI” (A12), while a corporate spokesperson emphasises that more than cost-cutting, it is about the speed of production, saying that AI makes it “five times” (A7) faster.

This logic of scale rewards quantity over quality, producing the low-effort aesthetic associated with slop: “amateurish, really corny, really bad art” (A9). For corporations, the appeal lies in productivity gains: “developers can now make stadiums in six weeks rather than six months” (A29). But these shifts are contentious, often seen as threats to creative labour and artistic integrity: “here’s hoping Activision hires some actual humans” (A20). Usage of AI tools is usually discovered by encountering distorted media: “loading screens depicted hands with extra digits” (A30).

These transformations are understood to be significant, as “it’s clear that business as usual won’t

⁹ https://en.wikipedia.org/wiki/AI_slop

be sustainable” (A5) and that it is just the start of the process—“so it begins” (A22).

3.2.2. Danger and discernment

But the slop tide threatens some of the key functions of the web, clogging search results with nonsense, overwhelming small institutions [...], and generally polluting the already fragile information ecosystem of the internet. (A5)

This theme is related to the consequences of the abundance of slop and the strategies used to deal with it. Subthemes include *danger*, *spam*, *scam*, *wading through*, *moderation*, *filtering*, *detection*, and *enforcement*. The excess of AI-generated content (Theme 1) creates the need for a strategy to deal with it. Libraries, search engines, and platforms struggle to filter growing volumes of synthetic content: “overwhelmed and underfunded local libraries” now house “unchecked, unedited AI-generated artifacts” (A5).

Unlike false information and malicious campaigns, slop is not overtly harmful: “[it’s] helpful to think of these photos and illustrations not as nefarious deepfakes or even hyper-persuasive propaganda, but as digital chum” (A14). Slop doesn’t “technically violate any rules” (A3), which complicates its enforcement. Platforms might tolerate it because “the slop makes everyone money” (A13). While not directly dangerous, “slop is often scam-adjacent” (A3) and is referred to as “threatening” (A7). For example, there are slop mushroom picking books which “contained dangerous advice for anyone hoping to discern a lethal fungus from an edible one” (A2), and chatbots can give questionable advice on “potentially delicate subjects” such as chemotherapy and Ozempic (A6).

3.2.3. Epistemic crisis

AI slop more broadly is poisoning the internet and is destroying the last shred of any sense of reality or reliability of social media platforms and search engines. (A15)

This theme is related to slop’s large scale social impact, such as its relationship to politics and misinformation. Subthemes related to this include *politics*, *conspiracy*, *reality*, *trust*, and *trump*. It relates to the epistemic crisis—a cultural climate of lack of trust in institutions, political polarization, worry about fake news and erosion of a shared reality. Trump, and the aesthetics of the alt-right, are

repeatedly referenced as emblematic: “Schoolcraft’s Facebook group is full of examples of troop-saluting, anti-LGBTQ+, pro-Trump schlock” (A9).

This breakdown of shared meaning and trust escalates the concerns about discernment (Theme 2) into more abstract and existential territory. While there are material concerns (e.g. preserving the credibility of advertisement organizations; A7), the dominant framing shifts towards moral and metaphysical anxieties. Truth and trust are depicted through metaphors of pollution and decay. Slop is not treated as merely an issue of quality, but as a hallmark of the post-truth condition: “we live in an era where the truth essentially does not matter” (A15); “the far right’s world-building project, where feel is always greater than real” (A16).

3.2.4. Engagement economy

This reality-fracturing is the result of an information ecosystem that is dominated by platforms that offer financial and attentional incentives to lie and enrage (A16)

Just like spam, almost no one wants to view slop, but the economics of the internet lead to its creation anyway. (A2)

This theme tells the story of slop as a consequence of the incentive structures that organise the internet. Subthemes include: *platforms*, *corporations*, *engagement*, *enshittification*, *incentives*, *money*, *virality*, *developing countries*. Central to this narrative is the influence of algorithmic systems, which shape content production in ways that prioritise reach and profitability over quality. As a result, slop is frequently linked to individuals producing viral material as part of informal economies—described as “side hustles” (A12) or “trying to make a quick buck” (A14).

Articles frequently emphasise the role of “incentives”, suggesting that slop is not primarily the result of personal proclivities, but structural forces. Instead of providing meaningful experiences, sloppers are prioritising metrics to maximise revenue. For this purpose, “the content of the content isn’t really as important as its presence—or, more accurately, its measurability” (A5). Facebook’s Performance Bonus Program – referred to by one article as a “slop subsidy” (A5) – rewards creators based on the engagement of their posts, which ultimately results in the creation of social media slop: viral nonsense

posts with attention-grabbing surreal imagery (e.g. Shrimp Jesus).

Although the term enshittification is not explicitly mentioned, its logic permeates much of the discussion, evident in phrases like “the great internet letdown” (A12), or:

And so it begins – the absolute dumbing down and worsening of everything from media to film to advertising to politics to the written word at the dead hands of AI. (A22)

While Theme 3 reflects metaphysical stakes (truth, trust, authenticity), Theme 4 reveals geopolitical ones (capital, platform governance, digital colonialism). Articles often refer to slop operations in developing countries — Kenya (A5), Vietnam (A11), Pakistan (A13), Malaysia (A15), etc. At the centre of this network are money and advertisement: money made by slop producers through platforms, by platforms through advertising, and by advertisers through increased productivity. As one article states, “AI slop is being generated by people all over the world for the express purpose of going viral and making money” (A15).

3.3. Qualitative analysis

This part of the study expands the analysis of AI slop beyond thematic analysis into a broader landscape of qualities, metaphors, and types. Given the study’s introductory nature, these results are reported even though they do not fit neatly into methodological boundaries.

3.3.1. Types

While social media slop is easily recognisable, it is only one of the types identified in the dataset, which include:

- **Social media slop:** low-effort or manipulative engagement bait created by sloppers for profiting on social media platforms. This kind of slop is primarily identified by its *surreal* and *kitschy* appearance.
 - **Example:** “The stuff of Salvador Dali’s wildest dreams is no match for Facebook these days: Amputee kittens using crutches. Strawberries in the shape of lifelike frogs. Bosomy conjoined twins, structurally impossible

sand sculptures, snakes swallowing fully-grown lions, airplanes with human hands. An underwater Jesus covered in shrimp.” (A9)¹⁰

- **Corporate slop:** low-quality AI-generated media used by corporations for cost-cutting purposes, typically in advertising. This kind of slop is primarily identified through its *distorted* quality, signalling *inauthenticity* and *lack of effort*.
 - **Example:** “A thumbnail image for 12 Angry Men on Amazon Freevee used AI to depict 19 men with smudged faces, none of whom appeared to bear any similarities to the characters in the film.” (A1)
- **Submission slop:** papers, books, applications and other media which are mass-submitted and overwhelm selection systems. This kind of slop is identified by its *excess* which overwhelms systems, its *inaccurate* information within media where accuracy is essential, and by being *generic* or *mediocre* in other contexts.
 - **Example:** “identical Viking ‘novels’ with seemingly AI-generated covers, all called Wrath of the Northmen: A Gripping Viking Tale of Revenge and Honor (that one has been published variously by authors named Sula Urbant, Sula Urbanz, and Sula Urbanr).” (A5)
- **Slop news:** low-quality news articles created by professional sloppers for SEO manipulation instead of human readership. This kind of slop is identified by being *generic* and *pointless* in comparison with news that was written for human consumption.
 - **Example:** see Appendix C

This is not an exhaustive list. Other examples could include summarization hallucinations – like “Google suggesting that you could add nontoxic glue to make cheese stick to a pizza? That’s slop” (A8) – as well as politically motivated disinformation, such as “Kamala Harris did not give a speech at the Democratic National Convention to a sea of communists while standing in front of the hammer

¹⁰ For examples of social media slop images, see the Ai Boomertrap(TM) Facebook group and the Insane Facebook AI Slop subreddit.

and sickle” (A14). However, these cases are more appropriately addressed under established categories (hallucination and disinformation, respectively). Similarly, the fake events prevalent in this dataset resulted from low-effort SEO hacking, and are not necessarily related to AI.

Slop is often economically motivated, closely tied to advertising, cost-cutting, propaganda, and scams. Advertising, in particular, plays a central role. Prominent examples include corporations generating AI-produced ads and assets, as well as individual “sloppers” creating viral images for engagement. Slop takes multiple forms depending on its context. On social media, it is primarily engagement-driven and often appears as image-based content. Produced by individual “sloppers” and circulated among users, it tends toward the surreal, kitsch, or simply pointless. Corporate slop, by contrast, is oriented toward productivity: corporations generate ads, films, and games in image, audio, or video formats. The results are typically distorted, mediocre, and inauthentic, yet still serve commercial ends.

In academic and professional domains, submission slop manifests in books and papers, frequently originating from sloppers or paper mills and directed at publishers and journals. This form is marked by excess, inaccuracy, and mediocrity. A related variant is slop news: articles produced for ad revenue and SEO, often in text or image form, circulated by sloppers and consumed by both users and bots. Such content is usually generic, pointless, and repetitive, yet financially incentivized by online attention economies.

Visually, slop is most recognisable in image-based content. This is likely due to the ease with which artifacts – distorted anatomy, surreal ideas – can be identified, particularly in social media and corporate contexts. The prevalence of slop in text is unknown, as it is harder to detect—something we investigate in the last section of this study.

3.3.2. Qualities

Here are the qualities of AI slop identified in the dataset, with corresponding words:

- Excessive (*everywhere, clutter, spam, inescapable, river, vast quantities, churn, glut, fatigue, high-volume, unmanageable, inescapable*)
- Inaccurate (*wrong, hallucinated, incorrect, dangerously inadequate, misleading, nonfactual, cartoonishly false, inaccurate, fake*)

- Inauthentic (*lifeless, voiceless, mawkish, inhuman, meaningless, authentic, legitimate, intent, purposeless, unnatural, mindlessly*)
- Mediocre (*low-quality, low-grade, shovelware, low-effort, too easy, cheap, shoddy, trivial, unskilled*)
- Surreal (*surreal, uncanny, bizarre, unsettling, outlandish, gobbledygook, nonsense, disturbing, dreamscapes*)
- Distorted (*distorted, warped, deformed*)
- Generic (*generic, filler, repetitive, unoriginal, unimaginative, banal, flat, boring, familiar, flat affect*)
- Pointless (*plausible, inconvenient, inoffensive, irrelevant, contextless, unneeded, useless*)
- Kitsch (*kitsch, shlock, high-resolution, low-fi, loud, messy, amateurish, corny, bad art, cringe, lowbrow, manipulative*)

The outcome of the proliferation of slop is the drowning out of intention, meaning, and effort; words like *intentionality, depth, creativity* and *value* are often invoked as slop’s opposites.

Some terms didn’t consolidate into qualities but recur: *disrespectful, inconvenient, inoffensive, good*. Many qualities – like *kitsch, surreal, excessive* – tie back to “weird AI crap” (A5) on Facebook, while others – *inauthentic, inaccurate, generic* – are more relevant to other types of slop such as scientific papers, e-books and news articles.

3.3.3. Metaphors

These are the metaphors of AI slop identified in the dataset, with corresponding words:

- Flooding: *flood, clogging, deluge, wade through, stream, drowning, rising tide, swamp, torrent, outpouring, dilutes, floodgates, river, drool*
- Contamination: *pollution, scrubbing, ecological footprint, chokes*
- Infestation: *spreading, proliferation, propagation, grow wildly, mulch, weeding out*
- Rubbish: *trash, rubbish, garbage, dross, crap, junk*

These metaphors show themes of flooding (water and overabundance), contamination (pollution and toxicity), infestation (unchecked spread and infiltration) and rubbish (worthlessness and clutter). In these we find the emotional aspects of slop: its

dirtiness, excess, vacuity and threat. Among these, the most dominant idea is *flooding*.¹¹

Language about slop is not only metaphorical, but displays a characteristic sound symbolism. Words related to slop share a kind of sludgy phonosemantics: mulch, glut, churn, chum, spam, etc. Other related concepts that share metaphorical aspects of excess, inauthenticity, and intrusiveness include “pink slime” (inauthentic media content), “slush pile” (unsolicited manuscripts), “vendor slurry” (library slop) and “grey goo” (self-replicating nanobots).

Slop also has a kinetic quality. Through the water metaphor, it is related to flowing water, or to clogging information flow. Through infestation, it spreads—undesirable, everywhere. Other minor metaphorical resonances identified include: inconvenient but not threatening (“clutter”) and mechanical or inhuman (“churn out”, “noise”).

4. Topic Modelling

Interdisciplinary research involves the integration of different methods as a tool for more insight into its object of study. For this purpose, we employ *topic modelling*—an NLP technique which, similar to thematic analysis, aims to identify themes (topics) within a dataset. The libraries utilised for this were `gensim` (Rehurek and Sojka, 2010) and `bertopic` (Grootendorst, 2022). We also reincorporated the articles which were previously discarded due to redundancy.

4.1. Gensim

Gensim is a Python topic modelling library. We applied a standard topic modelling technique utilising TF-IDF and Latent Dirichlet Allocation (LDA), setting hyperparameters to 10 passes, 5 topics, and auto alpha. Text underwent pre-processing, including tokenization, lemmatization, lowercasing, and removal of stop words, punctuation, and numbers. This resulted in the following topics:

- **Topic 0:** enix, cola, coca, square, foamstars, voice, actor, okatani, album, vercel
- **Topic 1:** parade, game, halloween, activism, slop, dublin, event, facebook, image, website

- **Topic 2:** coca, cola, meta, advert, baby, maga, harris, bandicoot, click, facebook
- **Topic 3:** claude, tweet, unwanted, term, ethic, llm, spam, aleksic, may, language
- **Topic 4:** trump, cola, brand, coca, spotify, centric, italian, brainrot, mcdonald, hurricane

Topics 0 and 1 primarily discuss corporate slop in gaming and fake events, Topics 2 and 4 focus on corporate slop in advertising and politics, and Topic 3 focuses on AI slop. All topics lean heavily towards named entities. For a more detailed analysis, we explore a more advanced model in the next section.

4.2. BERTopic

BERTopic (Grootendorst, 2022) streamlines a multi-step topic modelling pipeline. Documents are first segmented into sentences. From these, contextual embeddings are generated, reduced in dimensionality and clustered. For each cluster, a bag-of-words representation is created and topics are extracted. Finally, an LLM is used to name the topic. The number of topics can be adjusted through a hyperparameter of the clustering model. Following the documentation,¹² we use:

- Embeddings: `text-embedding-3-large` (instead of `all-MiniLM-L6-v2`)
- Dimensionality reduction: UMAP (McInnes et al., 2020)
- Clustering: HDBSCAN (McInnes et al., 2017)
- Tokenization: `sklearn`’s `CountVectorizer` (Pedregosa et al., 2011)
- Weighing: `c-TF-IDF`
- Representation: `gpt-4.1-nano-2025-04-14` (instead of `KeyBERT`)

OpenAI embeddings and representation were chosen for their performance in comparison with the free alternatives; the large embeddings model, in particular, was developed with dimensionality reduction in mind.¹³ For this dataset size (around 50,000 tokens), the whole analysis cost around a cent.

4.2.1. Interpretation approach

Using the default HDBSCAN `min_cluster_size` of 150 yielded approximately 80 topics. After filtering of irrelevant topics (e.g. dates, author names)

¹¹ Flooding has also been found in political contexts as a dehumanizing way to refer to immigrants (Mendelsohn and Budak, 2025)

¹² https://colab.research.google.com/drive/1BoQ_vakEVtojsd2x_U6-_x5200uqrj2

¹³ <https://openai.com/index/new-embedding-models-and-api-updates/>

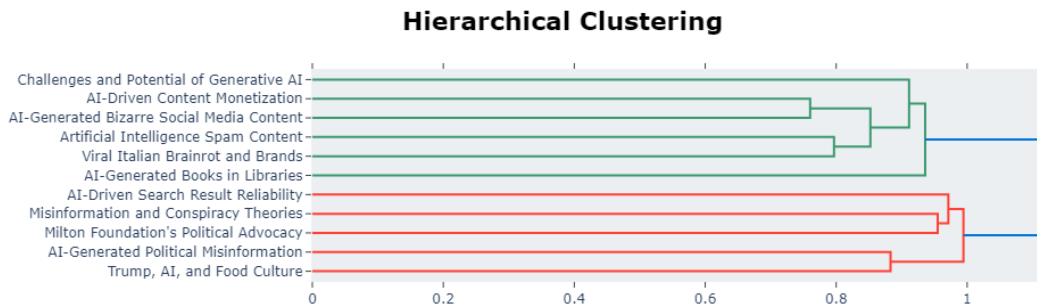


Fig. 2: BERTopic results after filtering.

and qualitative classification, four main topics were constructed:

- **AI-generated assets:** content created for media (games, films or ads) using AI.
- **Fake or low-quality events:** problematic public events advertised with AI-generated content.
- **Low-quality digital content:** spam, surreal slop, SEO-focused content, ad monetization, and moderation issues.
- **Information and politics:** misinformation, political narratives, media influence, public belief, and the epistemic crisis.

One shortcoming of this approach is that my previous knowledge of the thematic analysis might interfere in the creation of topics. The following section implements an alternative that addresses this limitation.

4.2.2. Filtering approach

Setting `min_cluster_size` to 18 results in about 10 topics. Only one was directly related to AI slop, while others included publication dates and named entities. By selecting only the relevant category and re-running the model, we can filter out ancillary information. (“AI-Generated content and Slop”). Results can be seen in Figure 2.

The resulting 11 topics are very similar to the themes constructed in the thematic analysis. These include ideas of monetization, surreality, spam, advertisement, libraries, SEO, misinformation and Trump, though it has one odd category (“Milton Foundation’s Political Advocacy”).

4.3. LLM-based thematic modelling

Recent literature suggests that using off-the-shelf LLMs for classification can be highly effective, even outperforming traditional methods (Chen and Shu, 2024). We test this approach here.

The LLM used was **Gemini 2.5-Pro** through the Google AI Studio interface. At the time of writing, **Gemini-2.5-Pro-Preview-06-05** is the highest placed model on the HuggingFace Chatbot Arena Leaderboard.¹⁴ Standard configurations were used: temperature set to 1, no system instructions and a simple prompt. A second prompt was used to summarise the previous response. This was done only once—no other attempts were made. Both prompts and responses are reproduced in Appendix B. The resulting themes were:

- **Defining Traits - Low Quality and Uncanny:** Slop is defined by its poor quality, visual flaws, nonsensical text, and a “soulless” or cheap feel that betrays its non-human origin.
- **Motivations - Profit, Speed, and Scale:** The creation of slop is driven by economic incentives such as engagement farming, corporate cost-cutting.
- **Impact and Consequences - Deception and Real-World Harm:** Slop has tangible negative effects, such as fake events, misinformation, and political propaganda.
- **The ‘Slop’ Framework - A Polluted Information Ecosystem:** Slop is a symptom of

¹⁴ <https://huggingface.co/spaces/lmarena-ai/chatbot-arena-leaderboard>

a “post-truth” sociotechnical environment where digital pollution erodes public trust and blurs the lines between truth and fiction.

4.4. Discussion

We compare four approaches: manual thematic analysis, traditional topic modelling using TF-IDF, contextual embedding-based topic modelling, and an LLM-based approach which we refer to as “thematic modelling”.

In summary, the themes identified by the LLM-based method were: traits of low quality and uncanniness; motivations rooted in profit, speed, and scale; consequences such as misinformation and propaganda; and the cultural context of pollution in a post-truth information ecosystem. The topic modelling approach raises similar themes, as well as touching on more concrete and specific issues, such as AI-generated books, Italian Brainrot and AI-driven search results, as well as completely different topics such as ‘Food culture’ and ‘Milton Foundation’.

Overall, the LLM-based approach outperformed topic modelling and delivered results equivalent to qualitative thematic analysis. Compared to topic modelling, it provided nuanced, interpretable outcomes with explanations, avoiding the pitfalls of simpler models that focus on named entities or tangential events. It produced more relevant, abstracted topics (e.g. ‘Defining Traits - Low Quality and Uncanny’) while the topic modelling produced more concrete topics (e.g. ‘Viral Italian Brainrot and Brands’). Compared to thematic analysis, the model’s output remained grounded in descriptive surface patterns, whereas the human coding was more interpretive. Ultimately, the convergence of findings across methods lends mutual validation and strengthens confidence in the results.

For exploratory research, LLMs provide an accessible alternative to traditional machine learning. While BERTopic demonstrated strong performance, it requires technical expertise, whereas LLM-based methods are promising for humanities researchers due to their ease of use and effectiveness. However, their simplicity can be deceptive: outputs can vary unpredictably due to hidden prompts or other uncontrollable variables, and prompt engineering and evaluation techniques remain underdeveloped compared to established practices in traditional ML.

For this small dataset, LLM-generated topics were more interpretable and more nuanced than traditional topic modelling. Although this balance

might shift with dataset size, for the humanities, which are generally concerned with smaller corpora, this is a promising technique.

It is important to note that LLM outputs strictly follow instructions, which can be understood in the context of sycophancy (Sharma et al., 2025). Unlike human researchers, LLMs do not reconsider or adapt research directions. Therefore, LLM-generated results should be treated as classifications requiring human interpretation, analogous to traditional topic modelling.

5. Linguistic and statistical features

This section investigates linguistic and statistical features of AI slop. The analysis is limited to the content of the text, excluding metadata, fact-checking, and “white box” access to models (Zubiaga et al., 2019; Wei et al., 2024; Bondielli and Marcelloni, 2019).

5.1. SlopNews: a dataset of AI slop

I introduce *SlopNews*, an English-language dataset containing 20,000 slop news articles, 20,000 human-written news articles, and 20,000 machine-generated news articles. The data are not shared directly due to copyright but are easily reproducible.

This dataset focuses on “slop news” over other text-based slop (e.g. submission slop such as e-books or academic papers) due to data availability. Ideally, the dataset would differentiate a) slop news from fake news, and b) high- from low-quality AI-generated content. However, these categories were not included in the dataset due to practical limitations.

Articles were labelled as “slop” based on their source website, a methodology also used in tackling fake news (Cameron, 2016). This work intends to enable future classification without relying on such metadata. By using content already identified as slop by others, this research avoids the ethical concerns of applying this label ourselves, such as the potential for reputational harm (see Knibbs, 2024b).

Sources were chosen based on reports from NewsGuard (Sadeghi and Arvanitis, 2023) and DoubleVerify (Luu, 2025),¹⁵ who identified websites containing generic, AI-generated, or fraudulent content. Four of their examples were selected and scraped: *famadillo*, *filthylucre*, *nbc sport*, and *cbsnews2*. These were selected based on their

¹⁵ Others can also be found on Knibbs (2024a). See Appendix C for all websites.

Table 1. SlopNews data sources.

Source	Count	Category
NeuralNews	20,000	AI
NeuralNews	8,251	Human
GlobalNews	8,251	Human
Newsdata.io	3,498	Human
nbc sport.com	15,948	Slop
filthy lucre.com	2,608	Slop
famadillo.com	1,136	Slop
cbsnews2.com	715	Slop

availability and ease of scraping. As a comparison baseline, we used human-written and machine-generated articles from three datasets:

- **NeuralNews** : human New York Times articles and AI-generated (GROVER) text (Tan et al., 2020);
- **GlobalNews**: Kaggle dataset obtained from NewsAPI (Saksham, 2023);¹⁶
- **NewsData.io**: news data aggregation company which offers free small datasets;¹⁷

All data were pre-processed to remove empty entries, duplicates, and non-English articles (using `langid`) (Lui and Baldwin, 2012). Non-slop categories were randomly downsampled to create three 20,000-entry samples (slop, human, and AI-generated), for a final total of 60,000 articles. The final composition of the dataset is shown in Table 1.

5.2. Analysis

The following metrics were computed for each article:

- **Sentence length**: mean, median, and standard deviation of sentence length per document.
- **Token length**: mean, median, and standard deviation of tokens (words/subwords) per document.
- **Vocabulary size**: proportion of unique tokens normalised by document size.
- **Stop word usage**: proportion of stop words normalised by document size.
- **Readability**: Flesch Reading Ease score, which is based on sentence length and syllables per word.

- **Complexity**: mean and standard deviation of dependency distance, which is a measure of the distance between related words in a sentence.
- **Coherence**: cosine similarity between adjacent (first- and second-order) sentence embeddings.
- **Polarity**: sentiment measure ranging from -1 (negative) to 1 (positive). Calculated using a Naïve Bayes classifier from the `textblob` library.
- **Subjectivity**: ranges from 0 (objective) to 1 (subjective). Calculated using the same classifier as polarity.
- **Sentiment (neural)**: computed using a deep learning classifier from the `flair` library (Akbik et al., 2019). Returns a sentence-level binary classification (positive/negative) and a confidence value (0–1); these were transformed into a document-level score by averaging the sentence values weighed by their confidence.

All metrics were calculated using the `textdescriptives` library (Hansen et al., 2023) unless stated otherwise. Note that several metrics (sentence length, readability, complexity, polarity) are sentence-level measures—creating a distribution for each document.

5.3. Results

Analysis was performed on a sample of the dataset comprised of 19,912 slop and 11,874 non-slop articles (7,975 globalnews, 3,386 newsdataio, 513 neuralnews), after excluding all AI-generated articles and all documents with less than 100 characters (<1% of remaining documents) due to parsing errors in some libraries. This also prevented the computation of perplexity measures and introduced some unexplainable selection bias.

Results show similar distributions between slop and non-slop, with a few exceptions: slop shows a slightly lower mean in standard deviation of sentence length, a higher concentration of low unique token proportions, a more uniform distribution than non-slop in both sentiment polarity and subjectivity, and a distinctly different distribution in neural sentiment (see Figure 3). Neural sentiment for slop is concentrated around positive sentiment, whereas the non-slop distribution has three peaks around negative, neutral, and positive sentiment. Due to the way this was calculated, values close to zero represent low confidence in the classification.

Statistical comparisons were conducted using two-sided Mann-Whitney U tests. Since we expect all results to be statistically discernible due to the large sample size, we assess its practical implications

¹⁶ <https://newsapi.org/>

¹⁷ <https://newsdata.io/datasets>

Table 2. Results of Mann–Whitney U-tests comparing non-slop and slop samples.

Metric	U	p	Mean (slop)	Mean (non-slop)	CLES
Sentence Length Mean	118668692.0	0.001	24.380	24.796	0.51
Sentence Length Median	114212087.5	0.008	23.135	23.802	0.49
Sentence Length Std. Dev.	139240493.5	3e-190	10.894	9.989	0.59
Token Length Mean	121076067.0	3e-4	4.616	4.633	0.51
Token Length Median	116173757.0	4e-6	4.080	4.064	0.49
Token Length Std. Dev.	116222253.0	0.011	2.544	2.544	0.49
Flesch Reading Ease	120226097.5	0.011	65.987	65.774	0.50
Dependency Distance Mean	108375684.0	1e-35	2.948	2.877	0.45
Dependency Distance Std. Dev.	100868767.5	1e-106	0.709	0.654	0.42
Coherence (first order)	119395882.5	5e-24	0.837	0.847	0.53
Coherence (second order)	106350805.5	6e-11	0.841	0.844	0.52
Unique Token Proportion	95902758.0	6e-175	0.581	0.529	0.40
Stop Word Proportion	54229874.0	0.018	0.463	0.464	0.49
Sentiment Polarity	130650555.0	1e-55	0.111	0.125	0.55
Sentiment Subjectivity	146451786.5	8e-279	0.386	0.448	0.61
Neural Sentiment	144980342.5	0.0	-0.089	0.351	0.65
Sentiment Confidence	97631371.5	3e-70	0.913	0.920	0.43

CLES = Common Language Effect Size. Values truncated for brevity. Larger effect sizes are bolded.

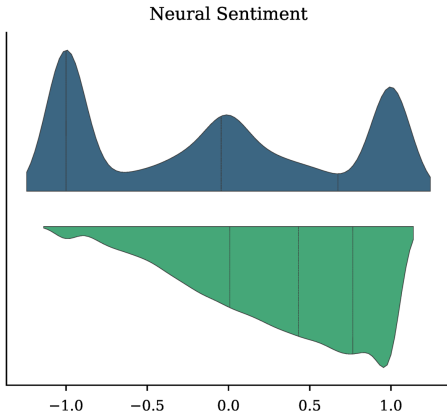


Fig. 3: Split violin plots comparing non-slop (top, blue) and slop (bottom, green) conditions.

by computing the Common Language Effect Size (CLES), a measure developed to be intuitive to non-statisticians. CLES indicates the probability that a randomly chosen item from one group scores higher than a randomly chosen item from another. It ranges from zero to one, where 0.5 indicates there is no effect.

As indicated in Table 2, all metrics show statistically discernible differences ($p < 0.05$), but

most effect sizes were negligible. Metrics such as mean sentence length, token length, readability, coherence, and stop word proportion show minimal effect sizes (CLES 0.49–0.51). More substantial (although still modest) differences appear in text complexity and variation: sentence length std. dev. (CLES 0.59), dependency distance (CLES 0.45), dependency distance std. dev. (CLES 0.42), and unique token proportion (CLES 0.40). These may be acting as proxies for perplexity and burstiness, which are standard entropy measures used in AI-generated text detection (Wu et al., 2025). Our findings align with the literature which shows that AI text tends to be more uniform than human-written text (Chen and Shu, 2024). We thus conclude that slop text has lower variation and complexity than non-slop text.

Slop articles also show higher sentiment scores in both simple (CLES 0.55) and neural (CLES 0.65) models, as well as greater subjectivity (CLES 0.61). These findings suggest that slop tends to exhibit more positive sentiment and subjectivity than non-slop. One possible explanation is that slop news aims to be unobtrusive, adopting a generic positivity. This contrasts with other kinds of slop, such as social media slop and disinformation, where the purpose is being attention-grabbing.

In summary, slop text tends to be slightly less varied and less complex, while also being

more positive and subjective. These differences are measurable but modest. Of the features investigated, lexical diversity, dependency structure, and sentiment analysis offer the most promise for slop classification.

5.4. Limitations and future work

Similar to misinformation (Dufour et al., 2024), AI slop is primarily image-based. This study deliberately focuses on text. Even so, the dataset used here is quite limited, as reflected in the selection bias incurred from successive convenience sampling. Future research can be considerably improved by incorporating more diverse data sources, other types of slop, and languages other than English. Furthermore, a wider range of linguistic metrics can be applied, such as:

- Known slop fingerprints, such as accidental prompt disclosure, model refusals, vocabulary patterns, emoji usage and markup features. These can be sourced from efforts such as the Wikipedia Project AI Cleanup¹⁸, the Problematic Paper Screener¹⁹, originality.ai blog (Guillham, 2024), and Pangram Resources for Educators²⁰.
- Valence, arousal and dominance (VAD) metrics (Mohammad, 2025; Vishnubhotla and Mohammad, 2022)
- Linguistic Inquiry and Word Count (LIWC) features²¹
- Givenness as a measure of complexity²²
- Type Token Ratios (TTR) as a measure of lexical diversity²³

It is important to emphasise that the statistical findings presented here are unlikely to be generalisable. The observed differences are relatively weak and likely reflect biases (e.g. stylistic traits) in the data source selection rather than differences intrinsic to the slop/non-slop distinction. While

these metrics reveal some distributional differences, they are not robust enough to support reliable classification or decision-making on their own. A natural follow-up for this work would be the development of a classification model incorporating the features identified here, which would have potential applications in fields such as news aggregation, moderation, and fact-checking.

Finally, future work could aim to operationalise the characteristics investigated in the first part of this study (e.g. writing quality; Chakrabarty et al., 2025). However, this might prove difficult, as many are highly subjective and rely as much on the reader's interpretation as on the content itself.

6. Conclusion

This work investigated AI slop—its qualities, metaphors, themes, types, and statistical characteristics. Slop is described as excessive, inaccurate, inauthentic, mediocre, surreal, distorted, generic, pointless, and/or kitschy. It is enabled by cheap, scalable automation which is used to respond to the incentives of the engagement economy. AI slop floods digital spaces with low-effort, contextless, generic and/or surreal material. It is created for productivity, profit, or manipulation rather than meaningful communication. It is created by corporations optimising for scale and individuals chasing side income. While present in text, images, audio and video, slop images are its most iconic form, as illustrated by Shrimp Jesus.

AI slop operates at the limits of moderation, as it is not necessarily factually incorrect (like fake news) or offensive (like hate speech). Yet, its informational effects can be understood in the context of "pollution of information ecosystem and loss of consensus reality" (Slattery et al., 2025). By its sheer scale, it reduces the signal-to-noise ratio of the information environment and demands additional effort in order to filter it.

Discourse about slop can be categorised into four main themes: efficiency and productivity (slop is cheap, easy and scalable), danger and discernment (slop is threatening and requires filtering), epistemic crisis (slop is related to post-truth) and engagement economy (slop is created because of platform incentives). It is often described metaphorically in terms of flooding, contamination, infestation and rubbish.

The four main types of AI slop are: social media slop (low-effort, outrageous or manipulative

¹⁸ https://en.wikipedia.org/wiki/Wikipedia:WikiProject_AI_Cleanup/AI_catchphrases

¹⁹ <https://www.irit.fr/~Guillaume.Cabanac/problematic-paper-screener>

²⁰ <https://www.pangram.com/resources/most-common-ai-phrases>

²¹ https://lit.eecs.umich.edu/geoliwc/liwc_dictionary.html

²² https://trunajod20.readthedocs.io/en/latest/api_reference/givenness.html

²³ https://trunajod20.readthedocs.io/en/latest/api_reference/ttr.html

engagement bait created by “sloppers” for the purpose of profiting off social media algorithms), corporate slop (low-quality AI-generated media used by corporations for cost-cutting purposes, typically in advertising), submission slop (papers, books, job applications, and other media which are mass-submitted and overwhelm selection systems) and slop news (low-quality news articles created by professional sloppers for SEO manipulation).

By analysing its linguistic and statistical characteristics, we find that slop news is less complex, less varied and more positive in comparison with non-slop news. The most noticeable difference occurs in the distribution of sentiment across slop and non-slop, as there are almost no slop news articles with negative sentiment.

This study is one of the first systematic analysis of AI slop. Although limited, it outlines potential directions for further research and lays the groundwork for developing interpretable systems for slop detection. By advancing our understanding of AI slop, this work aims to support efforts to safeguard the integrity of the information environment.

7. Competing interests

No competing interest is declared.

8. Acknowledgments

The authors thank the anonymous reviewers for their valuable suggestions.

9. Data Availability Statement

The data underlying this article will be shared on reasonable request to the corresponding author.

References

- Aich, A., Bhattacharya, S., and Parde, N. (2022). Demystifying Neural Fake News via Linguistic Feature-Based Interpretation. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6586–6599.
- Akbik, A., Bergmann, T., Blythe, D., Rasul, K., Schweter, S., and Vollgraf, R. (2019). FLAIR: An Easy-to-Use Framework for State-of-the-Art NLP. In Ammar, W., Louis, A., and Mostafazadeh, N., editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 54–59, Minneapolis, Minnesota. Association for Computational Linguistics.
- Aleksic, A. (2024). The linguistic observer effect.
- Allen, J. D., Howland, B., Mobius, M., Rothschild, D., and Watts, D. J. (2020). Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances*, 6(14).
- Arguedas, A. R., Robertson, C. T., Fletcher, R., and Nielsen, R. K. (2022). Echo chambers, filter bubbles, and polarisation: A literature review. Technical report, Reuters Institute for the Study of Journalism.
- Bailey, R. (2024). The case of the AI-generated giant rat penis.
- Barbaresi, A. (2021). Traftilatura: A Web Scraping Library and Command-Line Tool for Text Discovery and Extraction. In Ji, H., Park, J. C., and Xia, R., editors, *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 122–131, Online. Association for Computational Linguistics.
- Barthel, M., Mitchell, A., and Holcomb, J. (2016). Many Americans Believe Fake News Is Sowing Confusion.
- Bondielli, A. and Marcelloni, F. (2019). A survey on fake news and rumour detection techniques. *Information Sciences*, 497:38–55.
- Braun, V. and Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101.
- Cameron, D. (2016). Viral ‘fake news’ list deleted after author and students get harassed, doxed.
- Cellan-Jones, R. (2017). Fake news worries ‘are growing’ suggests BBC poll. *BBC News*.
- Chakrabarty, T., Laban, P., and Wu, C.-S. (2025). AI-Slop to AI-Polish? Aligning Language Models through Edit-Based Writing Rewards and Test-time Computation.
- Chen, C. and Shu, K. (2024). Combating misinformation in the age of LLMs: Opportunities and challenges. *AI Magazine*, 45(3):354–368.
- Clarke, N. (2023). A Concerning Trend.
- Cohen, G. A. (2002). Deeper into Bullshit. In Buss, S. and Overton, L., editors, *The Contours of Agency*, pages 321–339. The MIT Press.
- Dalton, C. (2016). Bullshit for you; transcendence for me. A commentary on “On the reception and detection of pseudo-profound bullshit”. *Judgment and Decision Making*, 11(1):121–122.
- Derrida, J. (2020). *A Farmácia de Platão*. Iluminuras.
- Doctorow, C. (2023). The ‘Enshittification’ of TikTok. *Wired*.
- Doughman, J., Afzal, O. M., Toyin, H. O., Shehata, S., Nakov, P., and Talat, Z. (2024). Exploring the Limitations of Detecting Machine-Generated Text.
- Dufour, N., Pathak, A., Samangouei, P., Hariri, N., Deshetti, S., Dudfield, A., Guess, C., Escayola, P. H., Tran, B., Babakar, M., and Bregler, C. (2024). AMMeBa: A Large-Scale Survey and Dataset of Media-Based Misinformation In-The-Wild.

- Easwaran, K. (2023). Bullshit activities. *Analytic Philosophy*, pages 1–23.
- Emi, B. and Spero, M. (2024). Technical Report on the Pangram AI-Generated Text Classifier.
- Flanagan, D., Mavergames, C., O'Brien, A., Parker, S., Razis, G., and Zhou, H. (2024). Upholding Research Integrity in the Age of AI. Technical report, Wiley.
- Forlini, E. (2023). OpenAI Quietly Shuts Down AI Text-Detection Tool Over Inaccuracies. <https://www.pcmag.com/news/openai-quietly-shuts-down-ai-text-detection-tool-over-inaccuracies>.
- Frankfurt, H. G. (2009). *On Bullshit*. Princeton University Press, Princeton, NJ.
- Gonen, H. and Goldberg, Y. (2019). Lipstick on a Pig: Debiasing Methods Cover up Systematic Gender Biases in Word Embeddings But do not Remove Them. <https://arxiv.org/abs/1903.03862v2>.
- Gorrieri, L. (2024). Is ChatGPT Full of Bullshit? *Journal of Ethics and Emerging Technologies*, 34(1):1–16.
- Graña, R. B. (2016). *Heidegger ou as Vicissitudes da Destruição*. Age.
- Grootendorst, M. (2022). BERTopic: Neural topic modeling with a class-based TF-IDF procedure.
- Guillham, J. (2024). Most Commonly Used ChatGPT Words & Phrases.
- Han, B.-C. (2015). *The Burnout Society*. Stanford University Press, Stanford, California.
- Hansen, L., Olsen, L. R., and Enevoldsen, K. (2023). TextDescriptives: A Python package for calculating a large variety of metrics from text. *Journal of Open Source Software*, 8(84):5153.
- Haraway, D. (2015). *Simians, Cyborgs, and Women: The Reinvention of Nature*. Routledge, New York.
- Heap, T., Lawson, T., Farnik, L., and Aitchison, L. (2025). Sparse Autoencoders Can Interpret Randomly Initialized Transformers.
- Hicks, M. T., Humphries, J., and Slater, J. (2024). ChatGPT is bullshit. *Ethics and Information Technology*, 26(2):38.
- Hoffmann, M. and Frase, H. (2023). Adding Structure to AI Harm: An Introduction to CSET's AI Harm Framework. Technical report, Center for Security and Emerging Technology.
- Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., and Liu, T. (2025). A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Trans. Inf. Syst.*, 43(2):42:1–42:55.
- Illing, S. (2020). "Flood the zone with shit": How misinformation overwhelmed our democracy. <https://www.vox.com/policy-and-politics/2020/1/16/20991816/impeachment-trial-trump-bannon-misinformation>.
- Kapoor, S., Henderson, P., and Narayanan, A. (2024). Promises and pitfalls of artificial intelligence for legal applications.
- King, T., Wu, Z., Koshiyama, A., Kazim, E., and Treleaven, P. (2024). HEARTS: A Holistic Framework for Explainable, Sustainable and Robust Text Stereotype Detection.
- Klincewicz, M., Alfano, M., and Fard, A. E. (2025). Slopaganda: The interaction between propaganda and generative AI.
- Knibbs, K. (2024a). Confessions of an AI Clickbait Kingpin. *Wired*.
- Knibbs, K. (2024b). The Guy Behind the Fake AI Halloween Parade Listing Says You've Got It All Wrong. *Wired*.
- Krenak, A. (2020). *Ideas to Postpone the End of the World*. Anansi International, Toronto, ON.
- Labbé, C., Joelsing, F., and Cabanac, G. (2025). Problematic Paper Screener: Trawling for fraud in the scientific literature. <http://theconversation.com/problematic-paper-screener-trawling-for-fraud-in-the-scientific-literature-246317>.
- Lui, M. and Baldwin, T. (2012). Languid.py: An Off-the-shelf Language Identification Tool. In Zhang, M., editor, *Proceedings of the ACL 2012 System Demonstrations*, pages 25–30, Jeju Island, Korea. Association for Computational Linguistics.
- Luu, K. (2025). The Rise of AI-Generated Slop Sites: What Advertisers Need to Know.
- McInnes, L., Healy, J., and Astels, S. (2017). Hdbscan: Hierarchical density based clustering. *Journal of Open Source Software*, 2(11):205.
- McInnes, L., Healy, J., and Melville, J. (2020). UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction.
- McMahon, L. and Kleinman, Z. (2024). Google AI search tells users to glue pizza and eat rocks. <https://www.bbc.com/news/articles/cd11gzejgz4o>.
- Mendelsohn, J. and Budak, C. (2025). When People are Floods: Analyzing Dehumanizing Metaphors in Immigration Discourse with Large Language Models.
- Mohammad, S. M. (2025). NRC VAD Lexicon v2: Norms for Valence, Arousal, and Dominance for over 55k English Terms.
- Muzumdar, P., Cheemalapati, S., RamiReddy, S. R., Singh, K., Kurian, G., and Muley, A. (2025). The Dead Internet Theory: A Survey on Artificial Interactions and the Future of Social Media. *Asian Journal of Research in Computer Science*, 18(1):67–73.
- Narayanan, A. and Kapoor, S. (2025). AI as Normal Technology. <http://knightcolumbia.org/content/ai-as-normal-technology>.
- Oxford University Press (2024). Oxford Word of the Year 2024.
- Predregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, É. (2011). Scikit-learn: Machine Learning in Python. *Journal of*

- Machine Learning Research*, 12(85):2825–2830.
- Pennycook, G., Cheyne, J. A., Barr, N., Koehler, D. J., and Fugelsang, J. A. (2015). On the reception and detection of pseudo-profound bullshit. *Judgment and Decision Making*, 10(6):549–563.
- Pennycook, G. and Rand, D. G. (2019). Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*, 188:39–50.
- Peters, D. (2022). How NewsGuard monitors 8,000 websites to combat misinformation.
- Placido, D. D. (2024). Facebook’s AI-Generated ‘Shrimp Jesus,’ Explained. <https://www.forbes.com/sites/danidiplacido/2024/04/28/facebook-ai-generated-shrimp-jesus-trend-explained/>.
- Purdy, K. (2025). Open source project curl is sick of users submitting “AI slop” vulnerabilities. <https://arstechnica.com/gadgets/2025/05/open-source-project-curl-is-sick-of-users-submitting-ai-slop-vulnerabilities/>.
- Rampin, R. and Rampin, V. (2021). Taguette: Open-source qualitative data analysis. *Journal of Open Source Software*, 6(68):3522.
- Read, M. (2024). Drowning in Slop. <https://nymag.com/intelligencer/article/ai-generated-content-internet-online-slop-spam.html>.
- Rehurek, R. and Sojka, P. (2010). Software Framework for Topic Modelling with Large Corpora.
- Sadasivan, V. S., Kumar, A., Balasubramanian, S., Wang, W., and Feizi, S. (2023). Can AI-Generated Text be Reliably Detected?
- Sadeghi, M. and Arvanitis, L. (2023). Rise of the Newsbots: AI-Generated News Websites Proliferating Online.
- Saksham, K. (2023). Global News Dataset.
- Saphra, N. and Wiegrefe, S. (2024). Mechanistic?
- Sharma, M., Tong, M., Korbak, T., Duvenaud, D., Askill, A., Bowman, S. R., Cheng, N., Durmus, E., Hatfield-Dodds, Z., Johnston, S. R., Kravec, S., Maxwell, T., McCandlish, S., Ndousse, K., Rausch, O., Schiefer, N., Yan, D., Zhang, M., and Perez, E. (2025). Towards Understanding Sycophancy in Language Models.
- Slattery, P., Saeri, A. K., Grundy, E. A. C., Graham, J., Noetel, M., Uuk, R., Dao, J., Pour, S., Casper, S., and Thompson, N. (2025). The AI Risk Repository: A Comprehensive Meta-Review, Database, and Taxonomy of Risks From Artificial Intelligence.
- Squires, V. (2023). Thematic Analysis. In Okoko, J. M., Tunison, S., and Walker, K. D., editors, *Varieties of Qualitative Research Methods: Selected Contextual Perspectives*, pages 463–468. Springer International Publishing, Cham.
- Stockwell, S. (2024). AI-Enabled Influence Operations: Threat Analysis of the 2024 UK and European Elections. Technical report, Centre for Emerging Technology and Security.
- Sullivan, D. and Nelson, C. (2023). Google Search’s guidance about AI-generated content. <https://developers.google.com/search/blog/2023/02/google-search-and-ai-content>.
- Tan, R., Plummer, B. A., and Saenko, K. (2020). Detecting Cross-Modal Inconsistency to Defend Against Neural Fake News. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2081–2106.
- Tang, R., Chuang, Y.-N., and Hu, X. (2023). The Science of Detecting LLM-Generated Texts.
- Tangermann, V. (2024). Amazon Is Selling Products With AI-Generated Names Like “I Cannot Refill This Request It Goes Against OpenAI Use Policy”. <https://futurism.com/amazon-products-ai-generated>.
- Templeton, A., Conerly, T., Marcus, J., Lindsey, J., Bricken, T., Chen, B., Adam Pearce, Citro, C., Ameisen, E., Jones, A., Cunningham, H., Nicholas L Turner, McDougall, C., MacDiarmid, M., Tamkin, A., Durmus, E., Hume, T., Mosconi, F., Freeman, C. D., Sumers, T. R., Rees, E., Batson, J., Jermyn, A., Carter, S., Olah, C., and Henighan, T. (2024). Scaling Monosemanticity: Extracting Interpretable Features from Claude 3 Sonnet. Technical report, Anthropic.
- Trvisan, A., Giddens, H., Dillon, S., and Blackwell, A. F. (2024). Measuring Bullshit in the Language Games played by ChatGPT.
- Tully, S., Longoni, C., and Appel, G. (2025). EXPRESS: Lower Artificial Intelligence Literacy Predicts Greater AI Receptivity. *Journal of Marketing*, page 00222429251314491.
- Tversky, A. and Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*.
- Vishnubhotla, K. and Mohammad, S. M. (2022). Tweet Emotion Dynamics: Emotion Word Usage in Tweets from US and Canada. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4162–4176, Marseille, France. European Language Resources Association.
- Wardle, C. and Derakhshan, H. (2017). Information disorder: Toward an interdisciplinary framework for research and policy making. Technical report.
- Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., Šigut, P., and Waddington, L. (2023). Testing of detection tools for AI-generated text. *International Journal for Educational Integrity*, 19(1):26.
- Wei, J., Yang, C., Song, X., Lu, Y., Hu, N., Huang, J., Tran, D., Peng, D., Liu, R., Huang, D., Du, C., and Le, Q. V. (2024). Long-form factuality in large language models.
- Willison, S. (2024). Slop is the new name for unwanted AI-generated content.
- World Economic Forum (2024). Global Risks 2024: Disinformation Tops Global Risks 2024 as Environmental Threats Intensify.

<https://www.weforum.org/press/2024/01/global-risks-report-2024-press-release/>.

- Wu, J., Yang, S., Zhan, R., Yuan, Y., Chao, L. S., and Wong, D. F. (2025). A Survey on LLM-Generated Text Detection: Necessity, Methods, and Future Directions. *Computational Linguistics*, 51(1):275–338.
- Wynn-Williams, S. (2025). *Careless People: A Cautionary Tale of Power, Greed, and Lost Idealism*. Flatiron Books, New York, first edition edition.
- Zubiaga, A., Aker, A., Bontcheva, K., Liakata, M., and Procter, R. (2019). Detection and Resolution of Rumours in Social Media: A Survey. *ACM Computing Surveys*, 51(2):1–36.
19. AI slop: what Labour, Spotify and Coca-Cola can teach us in 2025 (Raconteur)
20. A24 under fire for using AI-generated images to promote Civil War (MobileSyrup)
21. Activision "AI slop" Guitar Hero advert is market research for a fake game that doesn't exist and maybe never will (Eurogamer)
22. Why Brands Are Embracing 'Italian Brainrot' to Go Viral on TikTok and Win Over Gen Z (NEWS10 ABC)
23. Willy Wonka event leaves bitter taste thanks to AI adverts (TheRegister)
24. Chaos in Dublin as thousands turn up for AI 'hoax' Halloween parade that didn't exist (The Independent)
25. All Trick, No Treat: Dublin Crowds Turn Up for Halloween Parade That Wasn't (The New York Times)
26. Dublin: Halloween parade listing 'mistake' says website owner (BBC News)
27. The Guy Behind the Fake AI Halloween Parade Listing Says You've Got It All Wrong (Wired)
28. AI-Generated Slop Is Already In Your Public Library (404 Media)
29. Call of Duty Fans Give Black Ops 6's Zombie Santa Loading Screen the Finger Amid 'AI Slop' Backlash (IGN)
30. Activision Reportedly Sold an AI-Generated Call of Duty: Modern Warfare 3 Cosmetic (IGN)
31. Call Of Duty Discloses AI Slop After Months Of Players Complaining (Kotaku)
32. PlayStation Exclusive Foamstars Has AI-Generated Art, Square Enix Confirms (IGN)
33. Square Enix President States the Company Will Be 'Aggressive in Applying AI' (IGN)
34. Angry Birds Block Quest: Rovio's latest title in the franchise has soft-launched for Android in the US (GamingonPhone)

A. Articles about slop

1. AI slop (Wikipedia)
2. Spam, junk ... slop? The latest wave of AI behind the 'zombie internet' (The Guardian)
3. Why doesn't Facebook just ban AI slop like Shrimp Jesus? (Business Insider)
4. Why the internet is filling up with nonsense 'AI slop' (The Telegraph)
5. Drowning in Slop (Intelligencer)
6. How Long Will A.I.'s 'Slop' Era Last? (The New York Times)
7. How Businesses Can Avoid AI Slop (Entrepreneur)
8. First Came 'Spam.' Now, With A.I., We've Got 'Slop' (The New York Times)
9. How 'AI slop' is creating a 'zombie internet' (The Washington Post)
10. Slop is the new name for unwanted AI-generated content (Simon Willison's Weblog)
11. Where Facebook's AI Slop Comes From (404 Media)
12. Side job, self-employed, high paid: behind the AI slop flooding TikTok and Facebook (The Conversation)
13. Facebook's Twisted Incentives Created Its AI Slop Era (Gizmodo)
14. The MAGA Aesthetic Is AI Slop (The Atlantic)
15. Hurricane Helene and the 'Fuck It' Era of AI-Generated Slop (404 Media)
16. I'm Running Out of Ways to Explain How Bad This Is (The Atlantic)
17. Coca-Cola's New AI-Generated Holiday Ad Slammed as 'Soulless' and 'Embarrassing': 'This Is Such Slop' (IGN)
18. Paramount skips voice actors, opting for shockingly bad AI slop in Novocaine film promo (TweakTown)

B. Prompts and completions of LLM-based method for theme extraction

Completion markup was edited for readability.

Prompt:

use the qualitative method of thematic analysis on the following articles to answer the question 'what is ai slop' ### {{concatenated text of all documents}}

Completion:

Excellent. This is a perfect use case for thematic analysis. By systematically identifying, analyzing, and reporting patterns (themes) across this set of articles, we can construct a comprehensive, evidence-based answer to the question, “What is AI slop?”

Methodology

A qualitative thematic analysis was conducted on the provided articles. The process involved:

1. **Familiarization:** Reading and re-reading the articles to gain an in-depth understanding of the content.
2. **Coding:** Systematically identifying and labeling key concepts, ideas, and phenomena related to AI slop.
3. **Theme Development:** Collating the codes into broader, overarching themes that capture the essence of the data.
4. **Review and Refinement:** Reviewing the themes against the original data to ensure they are coherent, distinct, and well-supported.
5. **Synthesis:** Writing up the analysis, using direct quotes and examples from the articles to illustrate each theme.

Five primary themes emerged from the analysis, which together define the concept of “AI slop.”

—

Thematic Analysis: What is AI Slop?

Based on the provided articles, “AI slop” is a multifaceted phenomenon characterized by its poor quality, its economic and political motivations, its deceptive and harmful real-world impacts, and its contribution to a polluted and increasingly distrusted information ecosystem.

Theme 1: Defining Traits - Low Quality, Incoherent, and Uncanny

The most immediate characteristic of AI slop is its poor quality. It is consistently described as being carelessly produced, nonsensical, and visually or textually flawed. This is not sophisticated deception, but rather mass-produced digital clutter.

* **Visual Flaws:** A common telltale sign is the “uncanny valley” effect, especially in AI-generated art. The articles point to a zombie Santa with “six fingers” in **Call of Duty**, AI-generated art for the movie **Nosferatu** that “does not resemble his look in the film,” and an AI image of a squirrel-like character “missing half a face” in a fake **Crash Bandicoot** ad.

* **Textual Gibberish:** The text associated with slop is often nonsensical or stilted. The Willy Wonka event featured an “AI-generated gibberish” script, and one AI-generated book review warned against a guide that contained “dangerous advice for anyone hoping to discern a lethal fungus from an edible one.” This extends to unintelligible captions like “catgacating” and “exarserdray lollipops” from the Willy Wonka promotional material.

* **Low-Effort & Soulless Nature:** The content is perceived as lacking human creativity and emotion. Coca-Cola’s AI-generated holiday ad was slammed as “soulless” and “embarrassing”: “This is such slop.” The result is content that is “cheap,” “shoddy,” and “devoid of any actual creativity.”

Theme 2: Motivations for Creation - Profit, Speed, and Scale

Slop is not created in a vacuum; it is the product of specific economic and efficiency-driven incentives. The primary motivation is rarely artistic or informational, but rather financial gain, speed, and the ability to produce content at an unprecedented scale.

* **Monetization & Engagement Farming:** The articles repeatedly link slop to revenue. Facebook’s “Creator Bonus Program” incentivizes creators in countries like India and Vietnam to churn out “bizarre AI spam” because “Facebook now pays you \$100 for 1,000 likes.” The owner of the fake Halloween parade website admitted his business model is “Google Ads and affiliate marketing,” and it’s “easy for us to generate revenue” from high-traffic holiday events.

* **Efficiency and Cost-Cutting:** Corporations use AI to drastically reduce production time and costs, even at the expense of quality. An AI studio founder touted that creating the Coca-Cola ad was “five times” faster than traditional methods. Activision allegedly laid off 2D artists and “forced” remaining concept artists to use AI. The use of AI is seen as a way to enhance “development productivity and achieve greater sophistication in our marketing efforts,” as stated by the president of Square Enix.

* **Low Barrier to Entry:** Generative AI has made it “trivial to automatically generate vast quantities of text or images.” This has fueled a global gray-market economy where influencers on YouTube and Telegram teach others how to

create slop as a “side job, self-employed, high-paid” venture.

Theme 3: Impact and Consequences - Misinformation, Deception, and Real-World Harm

Beyond being just digital junk, AI slop has tangible and often negative consequences. It deceives people, spreads misinformation, and can lead to real-world chaos and danger.

* **Hoaxes and Fake Events:** The most prominent example is the Dublin Halloween parade, where an “AI ‘hoax’” listing on a website led “thousands of people” to gather for an event that “never even existed.” Similarly, the “Willy’s Chocolate Experience” used AI-generated promotional images to sell tickets to a “desolate warehouse” event, leaving children crying and parents calling the police.

* **Dangerous Information:** The proliferation of AI-generated books has led to “dangerous advice” being sold on Amazon, such as mushroom foraging guides that could lead to poisoning. Public libraries are now inadvertently spending “taxpayer dollars to provide patrons with information they don’t realize is AI-generated.”

* **Political Weaponization:** Slop is used as a political tool to spread propaganda and sow distrust. During Hurricane Helene, an AI-generated image of a girl and a puppy was circulated by political figures to suggest government failure. One pundit justified it, saying, “I’m leaving it because it is emblematic of the trauma and pain people are living through right now,” typifying what one article calls the “‘Fuck It’ Era of AI slop,” where emotional appeal trumps truth.

Theme 4: The ‘Slop’ Framework - A Name for a Polluted, “Post-Truth” Internet

The emergence of the term “slop” itself is a key theme. It provides a lexicon to describe a new kind of internet pollution and the broader societal shift it represents.

* **A New Term for a New Problem:** The articles frame “slop” as the successor to “spam.” As one developer puts it, “the way that ‘spam’ became the term for unwanted emails, ‘slop’ is going in the dictionary as the term for unwanted AI generated content.” This naming is important because “it gives people a concise way to talk about the problem.”

* **The Zombie Internet:** Slop contributes to what is called the “zombie internet,” a space where “a mix of bots, humans and accounts that were once humans but aren’t any more mix together to form a disastrous website where there is little social connection at all.” This cheap, automated content clutters feeds and search results, “making it harder for them to find accurate information.”

* **Erosion of Reality and Trust:** The ultimate consequence of slop is a breakdown in shared reality. When politicians and influencers knowingly share fake images because they are “emblematic” of a feeling, it signals a move to a “post-reality” where truth is secondary. This erodes trust not just in online content, but in institutions like news media, government agencies (FEMA), and science. As one article concludes, “Americans are divided not just by political beliefs but by whether they believe in a shared reality—or desire one at all.”

Conclusion

In summary, the thematic analysis of these articles reveals that AI slop is low-quality, often incoherent, and uncanny content mass-produced by generative AI. It is primarily motivated by the pursuit of profit, speed, and scale, both by individual “sloppers” and large corporations. Its impact is far from benign, leading to real-world deception, dangerous misinformation, and political manipulation. Finally, the term “slop” itself is a new framework for understanding the pollution of the internet, a “zombie internet” where the lines between human and machine, and truth and fiction, are increasingly and deliberately blurred.

Prompt (another chat):

```
{{previous response}} ### summarize the themes into a bullet point list with a brief description for each
```

Completion:

Of course. Here is a summary of the four primary themes in a bullet point list with a brief description for each.

Summary of Themes Defining “AI Slop”

* **Theme 1: Defining Traits - Low Quality and Uncanny**

This theme covers the fundamental characteristics of the content itself. Slop is defined by its poor quality, visual flaws (like extra fingers or distorted faces), nonsensical text, and a “soulless” or cheap feel that betrays its non-human origin.

*** Theme 2: Motivations - Profit, Speed, and Scale**

This theme explains the reasons behind slop’s creation. It is driven by economic incentives, such as individuals monetizing engagement through ads (“engagement farming”), and corporate desires to cut costs and dramatically accelerate content production by replacing human labor.

*** Theme 3: Impact and Consequences - Deception and Real-World Harm**

This theme highlights that slop is not harmless. It has tangible negative effects,

from deceiving the public with fake events and promotions (the Willy Wonka experience) to spreading dangerous misinformation (faulty foraging guides) and being weaponized for political propaganda.

*** Theme 4: The ‘Slop’ Framework - A Polluted Information Ecosystem**

This theme addresses the broader context and consequences. The term “slop” itself provides a new name for this form of digital pollution, which contributes to a “zombie internet” cluttered with automated content. Ultimately, it erodes public trust and blurs the lines between truth and fiction in a “post-truth” environment.

C. Sources of slop

Table 3. Websites identified by external credibility assessments.

Source	Website	Operational
NewsGuard	famadillo.com	X
NewsGuard	wavefunction.info	
NewsGuard	scoopearth.com	X
NewsGuard	filthylucre.com	X
NewsGuard	dailybusinesspost.com	X
NewsGuard	getintoknowledge.com	X
NewsGuard	bestbudgetusa.com	
NewsGuard	harmonyhustle.com	
NewsGuard	historyfact.in	
NewsGuard	countylocalnews.com	X
NewsGuard	tnewsnetwork.com	
NewsGuard	celebritiesdeaths.com	
NewsGuard	Biz Breaking News	
NewsGuard	News Live 79	
NewsGuard	Market News Reports	
DoubleVerify	espn24.co.uk	
DoubleVerify	nbc sport.co.uk	X
DoubleVerify	cbsnewsz.com	
DoubleVerify	cbsnews2.com	X
DoubleVerify	bbc sportss.co.uk	
DoubleVerify	foxnigeria.com.ng	X
DoubleVerify	bbc sportz.com	
DoubleVerify	247bbcnews.com	
WIRED	thehairpin.com	X
WIRED	antoniocarluccio.com	
WIRED	pope2you.net	X
WIRED	trumpplaza.com	X
WIRED	thefrisky.com	X