

AI Slop I: Pollution in Our Communication Environment

Jun-E Tan



Introduction

Was this article AI-generated, or did a human write it?¹

The ability to generate convincing text, audio, images and video using simple prompts has been hailed by some as progress towards “democratisation of content creation”². However, this powerful capability also comes with negative impacts. Within two years of mass access to generative AI (or gen-AI), the integrity of our information, communication and media landscape has been thrown into question.

Much attention has been focused on malicious use, such as deepfakes in electoral disinformation campaigns or non-consensual pornographic images. However, a much less

Views are short opinion pieces by the author(s) to encourage the exchange of ideas on current issues. They may not necessarily represent the official views of KRI. All errors remain the authors' own.

This view was prepared by Dr Jun-E Tan, a researcher from the Khazanah Research Institute (KRI). The author is grateful for research assistance from Dinie Joshry Shaeel and Balqis Athirah Azhar.

Author's email address:
june.tan@krinstitute.org

Attribution – Please cite the work as follows: Jun-E Tan. 2025. AI Slop I: Pollution in Our Communication Environment. Kuala Lumpur: Khazanah Research Institute. License: Creative Commons Attribution CC BY 3.0.

Photo by Justin Ha on Unsplash

Information on Khazanah Research Institute publications and digital products can be found at www.KRIInstitute.org.

¹ A human wrote it.

² Ma (2020)

discussed topic is the unintended consequences of the ability to mass produce content cheaply, thoughtlessly and at times erroneously.

This is the first of a three-part series on “AI slop”, referring to an overload of low-quality AI-generated content. In Part 1, we introduce the term and offer examples of areas that have been inundated by such content and discuss the drivers behind its proliferation. Parts 2 and 3 cover the consequences of AI slop from the perspectives of content consumers and content creators, society at large, and the AI models themselves.

What is AI Slop?

The term “AI slop” has entered modern lexicon, referring to low quality content (usually text, audio, images and video) generated by AI that is convincing at first glance but reveals its lack of substance upon deeper engagement. One dictionary definition of “slop” is “food waste (such as garbage) fed to animals”³, providing vivid imagery of the nature of content that is being consumed.

Experts have differentiated AI slop from disinformation, and instead call it “careless speech” or even “bullshit”, where inaccuracies and biases are subtle and not overtly wrong. This happens because the main goal of the generated content is not to be accurate or inaccurate, but to be persuasive⁴. For instance, Large Language Models (LLMs) are perfectly suited for generating convincing text as they can be prompted to generate any form or style of speech (see **Table 1** for a non-exhaustive list of examples of niches that have been affected).

Other gen-AI models on video, images, and audio have also achieved the level of quality that they can be mistaken for human-made content. In recent weeks⁵, readers may have noticed an onslaught of AI-generated videos in their social media feeds, due to new video apps launched by OpenAI, Meta and Google, to name a few. In particular, Sora 2 by OpenAI has grabbed headlines and caused multiple controversies, due to the social nature of the app whereby users can use likenesses of themselves and others to generate convincing deepfake videos⁶.

For music, accounts pumping out AI-generated music are reaching massive audiences. In June 2025, a fake rock band called the Velvet Sundown amassed more than 1 million Spotify plays in weeks⁷. Spotify has also revealed that it removed 75 million spam tracks from its platform over the past year as AI tools increase the ability of fraudsters to create fake music to generate income from streaming. The scale of the problem is significant, as Spotify’s actual catalogue stands at 100 million tracks⁸.

³ Merriam-Webster (2025)

⁴ Adami (2024)

⁵ Close to the publication date of this article, in October 2025

⁶ Kerr (2025)

⁷ Aroesti (2025)

⁸ Milmo (2025)

Table 1. Examples of AI slop in text form

Type of content	Description
Websites	There has been a 717% increase in AI slop sites detected in the last year. A flagged list of made-for-advertising (MFA) websites was 17,124 domains long in June 2024, and as of May 2025, it was 108,270 domains long; of those, 96,791 domains were classified as AI slop ⁹ .
Wikipedia articles	In a 2024 study, researchers ran detectors through Wikipedia and found that over 5% of newly created articles in English are AI-generated. They also found that the flagged articles were “typically of lower quality and are often self-promotional or partial towards a specific viewpoint on controversial topics” ¹⁰ .
Fiction	Publishers have had to close submissions for stories as editors struggle to keep up with an overwhelming increase of submissions, many of them discernibly AI-generated ¹¹ . In February 2023, online science fiction magazine <i>Clarkesworld</i> , which had been publishing the works of award-winning writers for 17 years, had to stop taking submissions after receiving a flood of AI slop ¹² .
Guidebooks	AI slop has affected guidebooks in niche areas such as mushroom foraging and lacemaking. Inaccurate content may lead to dire consequences in the short term (e.g. death by eating poison mushrooms ¹³) or in the long term (e.g. dilution and eventual loss of cultural heritage ¹⁴).
Academic papers	In a 2024 study, researchers examined academic writing (950,965 papers) published between January 2020 and February 2024 on arXiv, bioRxiv, and Nature portfolio journals. They concluded that LLMs are widely used in academic papers. In particular, the field of computer science was most affected at up to 17.5% ¹⁵ .
Software bug reporting	Open source developers are overwhelmed by security reports that are AI-generated ¹⁶ . Responding to security reports is expensive as manpower is required to check and fix vulnerabilities – this wastes resources as legitimate reports of real bugs are buried.
Work documents	A survey of 1,150 U.S.-based full-time employees across industries in July 2025 report that 40% have received workslop ¹⁷ in the last month. Employees who have encountered workslop estimate that an average of 15.4% of the content they receive at work qualifies. The phenomenon occurs mostly between peers (40%), but workslop is also sent to managers by direct reports (18%) ¹⁸ .

Source: Author’s own compilation

⁹ Bullett (2025)

¹⁰ Brooks et al. (2024)

¹¹ Clarke (2023); Sato (2023)

¹² Tangermann (2023)

¹³ Boutroy and Chang (2024)

¹⁴ Cole (2024)

¹⁵ Liang et al. (2024)

¹⁶ Maxwell (2024)

¹⁷ AI slop in the context of work documents, such as reports, emails or code

¹⁸ Niederhoffer et al. (2025)

The increasing accessibility of tools to generate AI slop is worrying, especially in engaging forms such as video. The concerns boil down to the questionable quality and authenticity of AI-generated content and the crowding out of genuine human-made content that takes time, energy and expertise to produce.

Mechanisms for quality assurance, traditionally human editors and curators, cannot match the volume or speed of AI content that are generated in seconds with simple prompts. Even large platforms such as Amazon have had to set publication limits – in Amazon’s case, to no more than three books per author per day¹⁹. This is not a high threshold, considering the amount of effort conventionally required to produce a full-length book.

Beyond practical considerations, there is also a moral dimension. Gen-AI relies on training data that often includes copyrighted material or scraped web content without permission of the original creators to use their work to train AI²⁰, not to mention compensation. Content creators therefore receive a double whammy of being exploited and being crowded out of the market.

Profiling from AI Slop

Why is there so much sloppy content?

There is a clear financial incentive for those creating content cheaply and monetising it. There are relatively straightforward ways to do this, such as hawking AI-generated e-books on Amazon directly to consumers. However, some have taken more sophisticated measures.

One example that has gained notoriety is the “SEO Heist” by a tech entrepreneur named Jake Ward (see **Figure 1**), who boasted on social media about stealing traffic from a rival website with the help of AI. As spelt out by Ward, his agency copied the structure of a popular website Exceljet²¹ and populated the skeleton with 1,800 pages of AI-generated content in a manner of hours. Using search engine optimisation, or SEO²², he claimed to have diverted 3.6million views from Exceljet to his own website within the span of 18 months²³.

¹⁹ Brodsky (2024)

²⁰ Milmo (2024)

²¹ Exceljet.net offers tutorials and tips to use Microsoft Excel.

²² Search engine optimisation, also known as SEO, is a way to improve the ranking of a website so that it shows higher on search results, bringing in more traffic and therefore more advertising money.

²³ After Ward’s post went viral, his website’s web traffic plummeted, causing speculation that Google had deprioritised it in its search results (Price (2023)).

Figure 1. SEO Heist



Source: Hoel (2024)

Another example covered by *Wired* tells the story of a man who bought defunct news websites (such as Hong Kong's *Apple Daily*) and filled them up with AI-generated clickbait content, profiting from the websites' strong search engine rankings from their popularity in the past²⁴. Observers are sounding alarms at the ease in which these operations are carried out²⁵, as top search results become dominated by AI-generated content²⁶.

Platforms to the rescue?

In an environment of content overload, the role of platforms as curators and moderators increases in importance. Much of our content and media now reside on large online platforms or rely on them as gatekeepers – they therefore become the de facto arbiters of what is considered appropriate and acceptable for consumption for large populations of users across the world.

Could we rely on them to separate signal from noise in our online environment?

Early signs are mixed. On one hand, large platforms, like small-time profiteers mentioned above, are also subject to the profit motive behind AI-generated content. Mark Zuckerberg, CEO of Meta has made it clear that Meta's platforms Facebook and Instagram will embrace AI-generated content in their newsfeeds, adding "a whole new category of content which is AI-generated or AI-summarized content, or existing content pulled together by AI in some way"²⁷. The release of Vibes, Meta AI's new video feed, is an implementation of this vision, algorithmically personalising the AI-generated videos to individual users to maximise engagement²⁸.

²⁴ Knibbs (2024)

²⁵ Haber (2024)

²⁶ Optimo (2024)

²⁷ Hays (2024)

²⁸ Malik (2025b)

On the other hand, content platforms also recognise that the proliferation of AI-generated content may drive away users looking for authenticity. For example, Spotify's rival streaming platform Deezer announced that it would label AI-generated tracks and ensure that their reach would be limited in both handpicked and algorithmic playlists²⁹.

Conclusion

With the amount of AI-generated content flooding all aspects of communication, it is increasingly difficult to separate signal from noise. While some have lauded the earning potential and productivity gains of generative AI, we will also need to be clear-eyed about unintended and longer term consequences.

The proliferation of AI-generated content has the potential to overwhelm the area of information, communication and media. The content is often generated with a short-term profit motive, with the overload drowning out quality content and information. What is at stake? A healthy information ecosystem supports informed citizens in their participation in democracy, and authentic human expression is the bedrock of culture and social connection.

We explore some of these concerns in the next articles in the series.

References

- Adami, Marina. 2024. "AI-Generated Slop Is Quietly Conquering the Internet. Is It a Threat to Journalism or a Problem That Will Fix Itself?" November 26, 2024. <https://reutersinstitute.politics.ox.ac.uk/news/ai-generated-slop-quietly-conquering-internet-it-threat-journalism-or-problem-will-fix-itself>.
- Aroesti, Rachel. 2025. "In an Era of AI Slop and Mid TV, Is It Time for Cultural Snobbery to Make a Comeback?" *The Guardian*, September 28, 2025, sec. Television & radio. <https://www.theguardian.com/tv-and-radio/2025/sep/28/cultural-snobbery-too-much-studio-rogen-house-of-the-dragon>.
- Boutroy, Gavin, and Arturo Chang. 2024. "A Handbook to Manitoba Mushrooms Sparks Concerns about AI. But for Technology, There's No Field Guide." *CBC News*, October 13, 2024. <https://www.cbc.ca/news/canada/manitoba/manitoba-mushroom-field-guide-artificial-intelligence-text-1.7351347>.
- Brodsky, Sascha. 2024. "The AI-Generated Books Trend Is Getting Worse." *AI Business*. January 25, 2024. <https://aibusiness.com/responsible-ai/the-ai-generated-books-phenomenon-is-getting-worse>.
- Brooks, Creston, Samuel Eggert, and Denis Peskoff. 2024. "The Rise of AI-Generated Content in Wikipedia." *arXiv*. <https://doi.org/10.48550/arXiv.2410.08044>.
- Bullett, Madeleine. 2025. "AI-Powered 'Slop' Sites: The Rapidly Rising Threat to Programmatic Advertising - Deepsee.io." June 26, 2025. <https://deepsee.io/blog/ai-slop-sites-programmatic-advertising>.

²⁹ Malik (2025a)

- Clarke, Neil. 2023. "A Concerning Trend." February 15, 2023. <https://neil-clarke.com/a-concerning-trend/>.
- Cole, Samantha. 2024. "AI-Generated Book Gifters Threaten The Future of Lace-Making." 404 Media. December 26, 2024. <https://www.404media.co/bobbin-tatting-lace-ai-generated-books/>.
- Haber, Morey. 2024. "The Rise Of Advanced Clickbait: How AI Is Tricking Unsuspecting Users." *Forbes*. August 13, 2024. <https://www.forbes.com/councils/forbestechcouncil/2024/08/13/the-rise-of-advanced-clickbait-how-ai-is-tricking-unsuspecting-users/>.
- Hays, Kali. 2024. "Mark Zuckerberg Says a Lot More AI Generated Content Is Coming to Fill up Your Facebook and Instagram Feeds." Yahoo Finance. October 31, 2024. <https://finance.yahoo.com/news/mark-zuckerberg-says-lot-more-233924648.html>.
- Kerr, Dara. 2025. "OpenAI Launch of Video App Sora Plagued by Violent and Racist Images: 'The Guardrails Are Not Real.'" *The Guardian*, October 4, 2025, sec. US news. <https://www.theguardian.com/us-news/2025/oct/04/openai-sora-violence-racism>.
- Knibbs, Kate. 2024. "Confessions of an AI Clickbait Kingpin." *Wired*, July 2, 2024. <https://www.wired.com/story/confessions-of-an-ai-clickbait-kingpin/>.
- Liang, Weixin, Yaohui Zhang, Zhengxuan Wu, Haley Lepp, Wenlong Ji, Xuandong Zhao, Hancheng Cao, et al. 2024. "Mapping the Increasing Use of LLMs in Scientific Papers." arXiv. <https://doi.org/10.48550/arXiv.2404.01268>.
- Ma, Wei-Ying. 2020. "Democratizing Content Creation and Dissemination through AI Technology." In *Proceedings of The Web Conference 2020*, 3143. WWW '20. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3366423.3382669>.
- Malik, Aisha. 2025a. "Deezer Starts Labeling AI-Generated Music to Tackle Streaming Fraud." *TechCrunch* (blog). June 20, 2025. <https://techcrunch.com/2025/06/20/deezer-starts-labeling-ai-generated-music-to-tackle-streaming-fraud/>.
- . 2025b. "Meta Launches 'Vibes,' a Short-Form Video Feed of AI Slop." *TechCrunch* (blog). September 25, 2025. <https://techcrunch.com/2025/09/25/meta-launches-vibes-a-short-form-video-feed-of-ai-slop/>.
- Maxwell, Thomas. 2024. "Bogus AI-Generated Bug Reports Are Driving Open Source Developers Nuts." Gizmodo. December 10, 2024. <https://gizmodo.com/bogus-ai-generated-bug-reports-are-driving-open-source-developers-nuts-2000536711>.
- Merriam-Webster. 2025. "Slop Definition & Meaning." Merriam-Webster. 2025. <https://www.merriam-webster.com/dictionary/slop>.
- Milmo, Dan. 2024. "'Impossible' to Create AI Tools like ChatGPT without Copyrighted Material, OpenAI Says." *The Guardian*, August 1, 2024, sec. Technology. <https://www.theguardian.com/technology/2024/jan/08/ai-tools-chatgpt-copyrighted-material-openai>.
- . 2025. "Spotify Removes 75m Spam Tracks in Past Year as AI Increases Ability to Make Fake Music." *The Guardian*, September 25, 2025, sec. Music.

<https://www.theguardian.com/music/2025/sep/25/spotify-removes-75m-spam-tracks-past-year-ai-increases-ability-make-fake-music>.

Niederhoffer, Kate, Gabriella Rosen Kellerman, Angela Lee, Alex Liebscher, Kristina Rapuano, and Jeffrey T. Hancock. 2025. "AI-Generated 'Workslop' Is Destroying Productivity." *Harvard Business Review*, September 22, 2025. <https://hbr.org/2025/09/ai-generated-workslop-is-destroying-productivity>, <https://hbr.org/2025/09/ai-generated-workslop-is-destroying-productivity>.

Optimo. 2024. "Does AI Content Rank In Google? Our Study + Data." Ask Optimo. August 2, 2024. <https://askoptimo.com/blog/ai-content-study>.

Price, Rob. 2023. "Inside the First 'SEO Heist' of the AI Era." Business Insider Nederland. December 19, 2023. <https://www.businessinsider.nl/inside-the-first-seo-heist-of-the-ai-era/>.

Sato, Mia. 2023. "AI-Generated Fiction Is Flooding Literary Magazines — but Not Fooling Anyone." The Verge. February 25, 2023. <https://www.theverge.com/2023/2/25/23613752/ai-generated-short-stories-literary-magazines-clarkesworld-science-fiction>.

Tangermann, Victor. 2023. "Sci-Fi Magazine Shuts Down Submissions After Onslaught of AI-Generated Spam." *Futurism* (blog). February 21, 2023. <https://futurism.com/the-byte/sci-fi-magazine-shuts-down-submissions-ai-spam>.