Is computational creativity flourishing on the dead internet?

Terence Broad
Creative Computing Institute
University of the Arts London
United Kingdom
t.broad@arts.ac.uk

Abstract

The dead internet theory is a conspiracy theory that states that all interactions and posts on social media are no longer being made by real people, but rather by autonomous bots. While the theory is obviously not true, an increasing amount of posts on social media have been made by bots optimised to gain followers and drive engagement on social media platforms. This paper looks at the recent phenomenon of these bots, analysing their behaviour through the lens of computational creativity to investigate the question: is computational creativity flourishing on the dead internet?

Introduction

The dead internet theory is a conspiracy theory that emerged in the late 2010's or early 2020's that states that large parts of the internet, in particular on social media are no longer occupied by humans and human generated content, but rather posts by AI-driven bots that are designed to control or influence human behaviour (IlluminatiPirate 2021).

Whist the theory emerges from the fringes of the internet, stemming in conspiratorial thinking as a way of explaining broad-based changes to society from nefarious actors, many commentators have observed that there is a grain of truth to the theory (Tiffany 2021). With the emergence and widespread adoption of generative deep learning, which is able to generate media in many domains to a plausible human level of fidelity, the possibility of bots acting autonomously or semi-autonomously on social media in order to garner influence is no longer a speculative possibility, but a real phenomenon on many social media platforms.

Engagement driven social media bots, or *AI influencers* (Walter 2024), that are optimised to maximise common social media metrics now exist on many social media platforms such as X (formerly Twitter), Facebook, Instagram, Reddit and TikTok. A recent report investigating web bots and internet traffic estimates that nearly 50% of web traffic is now driven by bots (Imperva 2024).

Bots on social media for the purposes of generating art, poetry or other kinds of content that are designed explicitly as computationally creative agents or artistic experiments in themselves are not new (Veale and Cook 2018). What is new about the phenomenon of these *AI influencers* is that they are not explicitly listed as bots, but instead posing as real people on social media or accounts that aggregate and share the creative work of real people. It is likely these bots are being developed by spammers and scammers to create accounts that have a large amount of engagement so that they traffic off the platforms to content-farm sites where advertising revenue can be generated (Koebler 2024a).

This paper will investigate the recent emergence of bots on various social media platforms, that appear to be explicitly optimised to drive engagement, and are acting in an at least somewhat autonomous fashion. This paper will analyse the output and behaviour of these models through the lens of computational creativity to determine if these bots are acting as autonomously computationally creative systems in the wild.

Social Media and Content Farms

Social media platforms, such as Facebook, Instagram, Reddit and TikTok rely on content generated by users as the means of engagement on their platforms and to leverage network effects through an 'architecture of participation' (O'Reilly 2005). The business model of these platforms is primarily to serve advertisements alongside user generated content.

Since the invention of the Facebook 'News Feed' in 2006, content is usually presented to users on these sites on a homepage, showing them the content from accounts that they follow (Arrington 2006). For many years the presentation of the feed has not been given in a chronological manner, but recommendation algorithms are used to to personalise the experience for users and will prioritise some content over other content in an algorithmic fashion. Since 2022, Facebook's Feed algorithm has been pushing content onto feeds from accounts that users do not follow to compete with TikTok's popular 'For You Page' (Heath 2022) which employs collaborative filtering to personalise the feeds of users to show them content from creators they primarily do not follow (Boeker and Urman 2022).

Content farms are websites or media creation organisations that make low effort and low quality content that makes money from advertisements, traditional online content farms would have paid freelancers to write articles that would get traffic from organic search, news aggregators and social media traffic (Bakker 2012). With the decline in revenues from

organic search and the changes made to search engines and social media platforms that discouraged users to leave their platforms and prioritised rankings for established organisations, content farming has moved to social media platforms themselves in search of revenue through content creator payments, advertisement revenue, and through promoting spam content.

AI-Powered Content Farming

Algorithmic content farms have existed long before the recent wave of generative AI tools. Content farms that make animated musical videos on the platform YouTube that take advantage of the young children's passive engagement with YouTube and it's auto-play featured that is powered recommendation algorithms, have been widely observed and criticised (Bridle 2017). The new wave of content farming is making use of generative AI to produce spam content and is being actively promoted by the platforms themselves into unsuspecting users feeds (DiResta and Goldstein 2024). These bots accounts appear to be making use primarily of modern text-to-image generation using techniques such as latent diffusion (Rombach et al. 2022) and large language models for text generation such as the Generative Pre-Training (GPT) class of models (Radford et al. 2018) to automate text generation on the click-farm websites.

An example of an AI-powered content farm is the social media account *Inspiring Designs* and associated website (inspiringdesigns.net), active on the Facebook, Instagram and TikTok platforms, with hundreds of thousands of followers on each platform. The posts and articles document new fictional product categories, such as 'power-tool toilets' (see Figure 1) 'cowboy-hat showers' and 'pickup-truck strollers'. The articles are very likely written by large language models and are written in a highly consistent style and structure. The goal of this site appears to be in order to host ads and generate revenue from affiliate links to the e-commerce website Amazon for products that bear little resemblance to the ones described in the articles.

Another example of a content farm that makes use of generative AI is the TikTok account *Globetrots*. Globetrots does not fully use generative AI, but rather amalgamates content from satellite imagery, publicly available information based on national and international statistical rankings (i.e. 'top 5 most dangerous motorways in the UK', or 'top 10 Biggest IKEA stores in the USA'). This account make use of text-to-speech voiceovers which are available to all users in Tik-Tok and widely used by AI spam content producers (Koebler 2024b). These videos are sometimes set to musical melodies and 'sung' by AI in various different genres.[1]

'Shrimp Jesus' and Combinatorial Creativity

A series of images that quickly rose to prominence and raised broader awareness to emergence of bot accounts on social media was the now infamous 'Shrimp Jesus' (Figure 2). This became a widely shared and discussed example of

---

[1]Such as this pop-punk musical countdown of the top 10 worst schools in the UK ranked by Ofsted reports https://vm.tiktok.com/ZGeVG15kd/.

Inspiring Designs

These Power Tool Toilets Will Drill Up Excitement in Any Handy Home

💬 0 Comments

Figure 1: Screenshot of the headline of an article from the website Inspiring Designs, detailing the fictional product category of 'power-tool toilets'. Available at https://inspiringdesigns.net/power-tool-toilets/

how many accounts on social media were using automated methods to combine concepts to make images and drive engagement on posts. The nonsensical combinations of concepts, combined with the mechanical regularity with which images are posted, are the main explanatory factors that have convinced many people that pages were being run in a partly or wholly autonomous fashion (Koebler 2024a).

Many people assume that these images were not just generated by a person entering text prompts into a text-to-image generator, but that some form of automated process was instrumental in creating the text prompts used to generate these images. This is evidenced by the mechanical regularity with which these accounts are posting new images and continuously trying new combinations of visual concepts, regardless of how nonsensical they may be. If these accounts are in fact driven by automated algorithms, then whatever underlying process that is generating these images is clearly an example of combinatorial creativity (Boden 2004) or bisociation (Koestler 1964), where more than one concept is brought together to make a new concept.

Religious Pareidolia

Fascination with religious pareidolia (the perception of religious iconography in otherwise random of ambiguous patterns) has long existed (Obadia 2018). Tabloid newspapers are regularly writing articles about people seeing images of Jesus in toast (Willis 2014) and various other objects.

Many bot driven accounts make use of the phenomenon of pareidolia of religious iconography to make posts that drive engagement. Figure 3 shows examples of two posts of photorealistic imagery that have the pareidolia effect of images of Jesus. These are likely done using the method of using

(a) (b)

Figure 2: Images posted on the (now defunct) *Love God &God Love You* (a) and the *Love Father &Mother Bless You* (b) Facebook pages. Both posts are AI generated images depicting 'Shrimp Jesus' – the figure of Jesus in part by the anatomical features of Shrimps underwater, posted with the caption *'Made with my own hands! Thanks to everyone who appreciate this'*.

(a) (b)

Figure 3: Images posted to the Facebook that make use religious pareidolia to drive engagement to their posts, both images show images of the face of Jesus appearing in photo-realistic imagery. (a) Image posted to the page *Cute Babies* showing two girls standing in water holding hands making the figure of Jesus. (b) Image posted to the page *Interesting stories* of a tropical beach setting with palm trees making out the figure of Jesus. Both images are posted with the caption 'Close your eyes 70% and see magic'.

a mask pattern to guide a text-to-image diffusion model that was originally developed to make photorealistic images that could act as scannable QR codes (Fu 2023).

Most of the comments on these religious posts are people replying with the Christian religious statement 'Amen'. These accounts that share religious imagery are ultimately taking advantage of people faith and their willingness to performatively engage with these posts and comment on them to demonstrate and reaffirm their faith on social media. These images are created with apparently little regard for their actual content or any sensitivity or moderation of imagery that some might consider blasphemous (such as the infamous 'Shrimp Jesus' images (Fig. 2)).

Engagement Hacking as Extrinsic Motivation

The intention behind the development of these bot accounts is likely not for them to be explicitly creative agents, but to maximise engagement on social media platforms in the most efficient and inexpensive way possible. The speed with with generative deep learning can produce realistic media, in what can be considered a mass produced fashion (Smith and Cook 2023), means that this technology has now become the cheapest and fastest way to generate content and drive engagement. Templates for posts and combinations of different concepts can be iterated on extremely quickly, with near instantaneous feedback with likes, comments and shares from other social media users.

Social media engagement metrics are clearly acting as a means of extrinsic motivation for the agents posting on these accounts in some way. Figure 4 depicts images depicting fictitious people in distressing situations who are celebrating their 'birthday'. These images are posted with captions like 'Happy Birthday To Me, but I haven't received any blessings yet' and are clearly designed in order to deceive users on the platform into liking the posts and commenting

birthday wishes for these fictitious people. The images in these posts often contain people with physical ailments and disabilities in order to elicit empathy from users, these ailments include prosthetic limbs (Fig. 4a), malnourishment (Fig. 4b), missing limbs (Fig. 4c), or respiratory problems (Fig. 4d). Alongside this these images will combine nationalistic, religious and sporting symbolism's to further drive engagement to these posts.

To what extent these accounts are using social media engagement metrics as direct form of extrinsic motivation in a closed autonomous loop, or whether this is mediated by humans managing these bots is difficult to determine. But it is clear that social media engagement and recommendation engine optimisation is the primary goal for these accounts. The quality, coherence and suitability of the content being posted is of a secondary concern and appears to have little human moderation.

Framing of Authorship

Many of these accounts on Facebook and other social media platforms frame themselves as aggregators of content, which are sharing (or reposting) the work of individuals, in the vein of a 'meme page' or 'meme aggregator' (Țăran 2014). The pages themselves are not framing themselves as the creators of the works, but simply the aggregators of content.

Often times on these pages the work is framed as having a human creator that is present in the generated image. Figure 5 shows several examples of these images which with the images and associated captions are clear examples of deceptive framing (Cook et al. 2019) being utilised to drive engagement. These images are not of real people but deepfake human avatars that are generated as part of the image, presumably included in the prompt of a text to image generator.

By including images of the supposed human creators in the images, these bot accounts are seeking to enhance the

(a)   (b)

(c)   (d)

Figure 4: AI generated images depicting fictitious people on their 'birthday' next to birthday cakes posted on Facebook. (a) Image posted to the page *Love dogs Cats Hub* of an injured American soldier with a prosthetic leg next to an image of Jesus holding a birthday cake with a sign saying 'Today is my Birthday'. (b) Image posted to the page *Mast* of a small child in a football top holding a birthday cake and a sign saying 'Happy birthday to me'. (c) Image posted to the page *Mast* of a small child riding a horse next to a birthday cake, both of which are both missing their lower limbs on their arms and front legs. (d) Image posted to the page *Social help* of two young girls with oxygen masks lying in shallow water with a sign saying 'Today is my birthday'.

persuasiveness of human authorship. This is akin to creating works under a pseudonym, but here it is a deepfake human avatar sitting besides the work. The real-world settings is likely designed to enhance supposed authenticity. In most of these posts, there is little other effort is made to frame the works beside a click-bait style caption.

Computational Creativity in the Wild

With the expense and difficulty in moderating social media platforms, and the difficult in detecting AI generated content, this phenomena is unlikely to disappear anytime soon. Bots are already generating content on social media in at least a semi-autonomous fashion. The computational creativity research community will need to start taking this phenomena seriously. By examining these examples in a critical fashion, this will further our understanding how models of computational creativity are being utilised by nefarious actors to generate spam and and scam content. These bots are utilising creative processes for ulterior motives with little consideration for the suitability of the content being produced. This could be considered a form of dark creativity (Cropley et al. 2010) or malevolent creativity (Cropley, Kaufman, and Cropley 2013).

If we are to understand these systems we will need to start analysing them from an external perspective to try to infer how they operate, as the likelihood is that many of the creators of these bot accounts will not be forthcoming about this. Platform restrictions permitting, the behaviour of these bot accounts could be analysed in a systematic way to try and reverse engineer their behaviour, such as prompts used and evolution over time. Leakage of text from prompts into writing on the images may also be a means to investigate the instructions given to these models. In addition, qualitative approaches could also be taken to better understand the behaviour of these bot accounts.

Conclusion

This paper suggests that models of computational creativity, whether intentionally or not, are being employed by nefarious actors to maximise engagement on social media in order to generate spam, to underpin content farming operations, and to generate engagement on accounts that can later be used for promoting scams. This is a clear delineation from, and emergence of computational creativity from the traditional worlds of academic research and creative practice, where computationally creative systems have previously been developed and disseminated. The goal of these bots are likely not necessarily to be creative agents per se, but simply a means to an end for some ulterior profit making motive.

How much of the behaviour of these bots is fully automated, and how much creativity autonomy in turn they have (Berns et al. 2021) is not currently possible to determine. Evaluating the extent of autonomy in these systems requires further investigation. However, the nature of the generations, both in their highly abnormal combinations of concepts, their seemingly regular and mechanical behaviour, and the perceived lack of content moderation of items being posted seem to suggest that these bots are acting in at the very least a semi-autonomous fashion.

As more and more of the internet is being populated by content generated by generative AI and automated bots, moving us closer towards a 'dead internet', then the prospect of these agents acting in creative ways, whether intentionally or not should be taken seriously. Studying the behaviour of these bots through the lens of computational creativity will help us understand the impact that these bots are having on both cultural production and the broader media ecology.

References

Arrington, M. 2006. New facebook redesign more than aesthetic. *Tech Crunch*.

Bakker, P. 2012. Aggregation, content farms and huffinization: The rise of low-pay and no-pay journalism. *Journalism practice* 6(5-6).

Berns, S.; Broad, T.; Guckelsberger, C.; and Colton, S. 2021. Automating generative deep learning for artistic pur-

Figure 5: AI generated images posted on Facebook that include fictitious human creators in the generated images. (a) Image posted on the page *Interesting stories* with the caption 'I did it myself, I'm waiting for your feedback'. (b) Image posted on the page *Interesting stories* with the caption: 'Made it with my own hands, but no one appreciates this'. (c) Image posted on the page *Interesting stories* with the caption: 'This granny wants to be rated for her beautiful crochet work'. (d) Image posted on the page *Life is beautiful* with the caption: 'I am taking part in a competition and this is my first work, please appreciate it'. (e) Image posted on the page *Life is beautiful* with the caption: 'My son painted his first picture I am very proud of him'.

poses: Challenges and opportunities. In *International Conference on Computational Creativity*.

Boden, M. A. 2004. *The creative mind: Myths and mechanisms*. Routledge.

Boeker, M., and Urman, A. 2022. An empirical investigation of personalization factors on tiktok. In *Proceedings of the ACM web conference 2022*.

Bridle, J. 2017. Something is wrong on the internet. https://medium.com/@jamesbridle/something-is-wrong-on-the-internet-c39c471271d2. Accessed: 2024-05-31.

Cook, M.; Colton, S.; Pease, A.; and Llano, M. T. 2019. Framing in computational creativity-a survey and taxonomy. In *International Conference on Computational Creativity*.

Cropley, D. H.; Cropley, A. J.; Kaufman, J. C.; and Runco, M. A. 2010. *The dark side of creativity*. Cambridge university press.

Cropley, D. H.; Kaufman, J. C.; and Cropley, A. J. 2013. Understanding malevolent creativity. In *Handbook of research on creativity*. Edward Elgar Publishing. 185–195.

DiResta, R., and Goldstein, J. A. 2024. How spammers and scammers leverage ai-generated images on facebook for audience growth. *arXiv preprint arXiv:2403.12838*.

Fu, A. 2023. Stylistic qr code with stable diffusion. https://antfu.me/posts/ai-qrcode. Accessed: 2024-05-31.

Heath, A. 2022. Facebook is changing its algorithm to take on tiktok, leaked memo reveals. *The Verge*.

IlluminatiPirate. 2021. Dead internet theory: Most of the internet is fake. https://forum.agoraroad.com/index.php?threads/dead-internet-theory-most-of-the-internet-is-fake.3011/. Accessed: 2024-03-22.

Imperva. 2024. 2024 Bad Bot Report.

Koebler, J. 2024a. Facebook's algorithm is boosting ai spam that links to ai-generated, ad-laden click farms. *404 Media*.

Koebler, J. 2024b. Inside the world of tiktok spammers and the ai tools that enable them. *404 Media*.

Koestler, A. 1964. The act of creation.

Obadia, L. 2018. Urban pareidolia: Fleeting but hypermodern signs of the sacred? *Bulletin for the Study of Religion47* (1).

O'Reilly, T. 2005. What is web 2.0. https://www.oreilly.com/pub/a/web2/archive/what-is-web-20.html. Accessed: 2024-03-22.

Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. 2018. Improving language understanding by generative pre-training.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.

Smith, A., and Cook, M. 2023. AI-generated imagery: A new era for the 'readymade'. In *SIGGRAPH Asia 2023 Art Papers*. ACM.

Ţăran, F. 2014. The structure and dynamics of meme aggregators. *Studia Universitatis Babes-Bolyai-Ephemerides* 59(1).

Tiffany, K. 2021. Maybe you missed it, but the internet 'died' five years ago. *The Atlantic*.

Veale, T., and Cook, M. 2018. *Twitterbots: Making machines that make meaning*. MIT Press.

Walter, Y. 2024. Artificial influencers and the dead internet theory. *AI & SOCIETY*.

Willis, A. 2014. Second coming of christ: Slice of toast bears face of jesus himself. *The Metro*.