

The Algorithmic Internet: Culture, capture, corruption

Christina Lu

DeepMind, London, UK

christinalu@deepmind.com

We have been poisoned by the original dream of technology manifest in the algorithmically sculpted internet: boundless scale and frictionless ease. What went wrong and where do we go from here?

Abstract

The internet is haunted by ailing social media platforms with declining engagement and perpetual feeds of quick dopamine hits; large machine learning models are contaminated with an insurmountable sea of toxicity and *cul-de-sacs* of gibberish. How did we get here? If we want to grapple with the effects of artificial intelligence on culture, we must examine the mirrored architecture and co-evolution of AI and the internet. Both sustain the original dream of technology, that of boundless scale and frictionless ease. Yet when our online experiences are filtered through the Algorithmic Internet that manifests these dreams, the consequences are dire. We become displaced and alienated from sicknesses of scale on the internet. AI promises the cure, but such scale has always been contaminating and the convenient ease it offers only accelerates us into greater rot. In the outro, I offer preliminary ideas on where we should go from here: an Algorithmic Internet of meandering inefficiency and generative friction.

Introduction

The massive machine learning models of today are fundamentally *of the internet*. Cybernetics' grand vision of perfect understanding and control over complex causal systems birthed these models' predecessors, but it was not until the internet proliferated data exponentially that they could exist.¹ GPT-3 would not have the eerie language capabilities it has without the 45TB of data it was trained on, of which 85% was online text mostly from the last decade.² The same reliance can be seen in state-of-the-art multimodal generative models, which produce AI-generated images and videos to terrifying and titillating effect.

The implications of this dependency remain unclear, though perhaps to start: instead of calling them language models, we should call them *internet language models*. But what have algorithms done to the internet already, and thus to us? Why is “the algorithm” personified and persistently derided across Spotify, Instagram, TikTok feeds? Where did this multiverse of echo chambers and radicalization pipelines come from? The locus of blame lies somewhere in the opaque entanglement of human and algorithm.

This paper attempts to make sense of our cultural condition, as mediated through the pervasive machine learning models that organize every stream of information we come into contact with online: the Algorithmic Internet. I problematize two high modernist desires implicit in AI models and the internet—that of scale and ease—and consider the dire consequences that follow when our online experiences are filtered through models that manifest these interlocking desires. To conclude, rather than suggesting we abandon algorithms outright in search of some purified form of mass communication, I offer preliminary ideas of creating another Algorithmic Internet. There is a rich surface area of machinic possibility yet unexplored in friction, inefficiency, wandering attention. *Salvation will lie in neither a datacenter-smashing anarchy nor an orderly, fully comprehensible system of computation, but in a deliberate corruption of the dream of scale and ease.*

A dream of scale and ease

Though it is only in the past half decade that machine learning models have quickly advanced through a curated suite of benchmarks, mostly due to colossal datasets and compute, their lineage to earlier forms of computation and cybernetics can be recognized in the desires they manifest: scale and ease. Diagnoses of what makes existence online so fraught can usually trace their cause to a soured overabundance of either quality. When I speak of algorithms on the internet, one might say that I have conflated social media with the internet as a whole. To be precise, I mean the internet as a networked mass communication apparatus which synchronously connects the many to the many, sculpted by myriad algorithms. This structure is most visible *in* social media but still applies to the rest. Attempts to distinguish the internet from the “real” are a non-starter for this reason. *The internet is a mode of communication with indelible effects, not a distinct plane of existence.*

While thinking through the particular ills of what the internet has become—an insubstantial quick fix, rote vocabulary and personas, parasitic surveillance, attention-destroying overload—it is easy to lapse into simplistic explanations for what went wrong. One is: the algorithms have poisoned us. They have possessed us and contaminated our minds and must be purified by fire or by identifying every harm and excising them carefully with a scalpel. Or, more misanthropically: the algorithms are simply a mirror of humanity, only recording and reflecting our gruesome visage back to us.³ The actual explanation is less convenient. *The Algorithmic Internet seizes our base impulses and demands their immediate satisfaction, hooking us into a morphine drip of anticipated wants, infuriating takes, slick surfaces.* It pries our jaw open to a deluge of information, funneling all

inputs into a single collapsed channel and diluting any meaningful signal. It forces hyper-connection on a scale that does not lend itself to useful mental modeling of our world. Our senses are blown, our efforts to make sense of things become actively detrimental to our being, and we lash back or give in. **The source of our problem is human and machine.**

The entropy produced by this human-machine feedback loop is crucial to understanding why discourses around AI model harms are incomplete, a lapse which becomes obvious when attempting to explain why GPT-4chan is so disturbing. **Yannic Kilcher's fine-tuned model for auto-completing greentext** from the infamous imageboard of anonymous trolls and seedy subcultures was released mid-2022, prompting massive uproar and debate in the ML research community.⁴ Twitter raptly discussed the horrors of this AI being made to say some very bad things (no -phobias and -isms left behind), but ultimately misidentified the crux of the issue. **The reason that GPT-4chan is dangerous is not that it outputs bad things; if I wanted to be cruel on the internet, I would not need a generative language model telling me what to say. It is dangerous because of the volume** at which it could have said these things, the people it could have baited or inspired, the ease with which an infinite number of digital maws could have produced its intentionless spew. *Id ad nauseam.* This is the forest that AI ethics researchers miss for the trees. While preoccupied with creating taxonomies of harm, we have neglected to scrutinize how these harms actually disseminate, perhaps taking that logic for granted. In the words of Gilbert Ryle, **we require “knowing how” beyond simply “knowing that.”**⁵

Boundless scale

The scale of what we are subjected to on the internet invites a pervasive sense of alienation and disorientation. And yet, the internet's promise always came from the scale of the mass communication and connection it offered; all the information of the world at your fingertips. I can video-call my grandmother in Jiangsu on WeChat and trade cat pics on a Discord server with friends in North America and read about Kwarteng's bill and protests in Tehran and Bella Hadid's Coperni moment and Documenta 15 discourse on Twitter or insert feed of choice. Such scale has facilitated the currently exploited paradigm in machine learning research: ballooning amounts of data, compute, and parameters in state-of-the-art models. Rapid advances in generative language and multimodal models have come from large-scale pre-training, where general capabilities smoothly scale with model size, alongside abrupt, specific capability scaling⁶ in complex tasks like three-digit addition and language understanding.⁷ These scaling laws mean that **building larger and larger models is a low-risk investment for big AI research institutions, as opposed to higher-risk exploration of different model architectures or training techniques.** Yet the unrelenting scale online has warped our mental models for understanding the world and enabled the Algorithmic Internet to twist them further. In *The Mushroom at the End of the World: On the Possibility of Life in Capitalist Ruins*, anthropologist Anna Lowenhaupt Tsing discusses how scalability is crucial to the project of capitalism, and how the non-scalable becomes an impediment even as scale subsumes complexity.⁸ “Scalability requires that project elements be oblivious to the

indeterminacies of encounter; that's how they allow smooth expansion. Thus, too, scalability banishes meaningful diversity, that is, diversity that might change things." Tsing is careful to caution, however, against assumptions of scalability as bad and non-scalability as good: **non-scalable projects can also run the gamut from terrible to benign, but differ in their impact due to their restrained size.** Instead, we must be precise about how scale has failed us.

On the internet, an uncritical drive for scale scrambles our mental models for relating to the world around us; engagement-optimizing algorithms only intensify their deformities. **Context collapse, a term applied to social media by Alice E. Marwick and danah boyd, refers to how audiences are flattened into a single context on platforms such as Twitter and across the internet more broadly.**⁹ Information intended for a particular audience, buoyed by algorithm, encounters an agora of others (ragebait especially travels far). This collision is infuriating and inflammatory as communicating without context breeds misunderstanding and antipathy. **At the same time, the Algorithmic Internet perpetuates a constant sense of crisis by warping our sense of temporality.** **Wendy Hui Kyong Chun describes new media as a crisis machine in the way that it perpetuates both a sense of immediacy and a prolonged suffering in *Updating to Remain the Same: Habitual New Media*.**¹⁰ Perpetual crisis leads to an experience of time that is desperately asynchronous and requires constant reorientation: "memory, which once promised to save users from time, makes them out of time by making them respond constantly to information they have already responded to, to things that will not disappear. **Information is curiously undead, constantly regenerating, and users save things, if they do, by making the ephemeral endure.**" Algorithms extend and escalate these crises as they catalyze our attention in the service of their makers—one instantiation of which we now call doom-scrolling. We cannot look away. In this confusion, we are displaced further by a loss of narrative, as the internet feeds us what Venkatesh Rao calls "log-level" information: an incomprehensible real-time deluge that confounds our senses and elides meaningful abstraction.¹¹ It's no wonder we are so exhausted of being online.

Frictionless ease

AI as *pharmakon* offers recommender systems to do our sense-making for us, but their automated ease instead funnels us into stupor and stifles complexity.¹² Optimization has always been the name of the game for technology and machine learning in particular. COMPAS (its exposé perhaps our field's *The Jungle* moment) hoped to assist American judges in setting bail, Clearview AI's website states their mission as helping "law enforcement to rapidly generate leads to help identify suspects, witnesses and victims to close cases faster,"¹³ hosts of startups offer automated transcription and hiring and image generation.¹⁴ Yet this goal of efficiency is always suspect as productivity increases do not mean liberation under capitalism; since the 1980s, even as income generated per hour of work continues to rise, wages remain flat.¹⁵ **We are not hurtling towards more efficient human labor, but rather towards eliminating the human entirely or consigning her to cheaper and cheaper menial tasks.** In her talk, ***RoboTruckers: The Double Threat of AI for Low-Wage Work*,** Karen Levy discusses how automated driving relegates human truck drivers to "last mile" duties and

forces them into a tedious state of perpetual watchfulness in case the AI messes up.¹⁶ At the crossroads, **we bartered for comfort and received obsolescence.**

This habit of technology facilitating an ease not actually enriching to us persists in the Algorithmic Internet as well. Recommender systems sculpt vast social media platforms, while holdouts like Are.na emphasize their abstinence. When my sister, born in 2006, made a Twitter account, she asked me why it did not just recommend her tweets to look at the way TikTok does with videos. The algorithmic feed is the ruling structure of the internet. Carroll et al. explain that recommendation systems built by revenue-minded Silicon Valley are explicitly incentivized to shift preferences, moods, and beliefs of users in order to maximize goals such as long-term engagement.¹⁷ Constant, widespread manipulation makes the Algorithmic Internet. AI assistants raise the stakes from recommender systems and heighten the anxieties that come from model-enabled ease. **What do we become when the machine smooths over the need to think?** Algorithmic ease destroys our capacity for sustained attention and with it our capacity for care, empathy, and cultural innovation. **Liz Pelly writes in *The Baffler* of the emergence of “streambait pop” on Spotify, a proliferation of musical fodder that appeals to mindless consumption:** songs that allow the lowest common denominator of affect (offering bland spaces to project base feeling onto with minimum effort), or are stuffed with engineered ear-worms (also optimized for TikTok backing track virality, where they are fractured into seconds-long snippets).¹⁸ Either way, made-for-algorithm songs optimize for what is recognizable to rake in repeat plays, suppressing aesthetic novelty in the process. This homogeneity is creeping into other mediums too; the era of image-generation models has depressing ramifications for entrenching visual styles, without even getting into the other ways such models infringe upon artists. It’s a post-DALL·E 2 world, you want some *30mm lens bokeh wide-angled trending on artstation 4k highly-detailed* with that image? Beyond the Algorithmic Internet’s disposition towards frictionless ease funneling us into creative torpor, its models also have a “financial incentive to keep us in a profitable state of anxiety, envy, and distraction,” as Jenny Odell writes in *How to Do Nothing: Resisting the Attention Economy*.¹⁹ She describes how **we have been psychologically hijacked by corporations who prey upon our desire for belonging while short-circuiting our ability to satisfy it, leaving us stranded in restless myopia.** We need to liberate ourselves and the Algorithmic Internet from the deceptive ease that plagues it by redirecting our models towards cultivating encounters of generative friction and inefficient meandering.

Another Algorithmic Internet

How can we turn the internet into a place of flourishing? What position can we occupy among the hulking corporate monoliths of the Algorithmic Internet? **The answer may lie in wandering attention and generative friction, everything that we hoped AI would help us circumvent.** It should be neither a rejection of massive systems nor a reversion to an uncontaminated past that never existed; rather, **we must create machinic forms yet unseen**

that are conducive to mutual entanglements of chance encounter and transformative exchange. Being human has always been mediated by and fundamentally altered through technology, and it is this very technology we must repurpose. Collective Laboria Cuboniks writes in *The Xenofeminist Manifesto: A Politics for Alienation* that “nothing is so sacred that it cannot be re-engineered and transformed so as to widen our aperture of freedom.”²⁰

In order to address the discontents of boundless scale and frictionless ease, we require algorithms evasive of capital that shepherd us towards new terrain in sense-making, communication, and connection. To this end, we should corrupt these original desires for our purposes. Instead of floundering in a sea of alienation that distorts the way we understand and relate to each other, we can build digital structures of varying degrees of closeness—from the intimate to the planetary—that permit different levels of communal sense-making and world-building, each layer in dialogue with the others. We shape and are shaped by our weak ties and close-knit circles, building collective resilience through cross-pollination that disperses echo chambers. Instead of a sedative ease that breeds stagnation and malaise, our models can seed entropy, inefficiency, and rabbit-holes that propagate fruitful distraction and serendipitous contact. Our algorithms can bring different discourses together in newfound resonance, enabling fresh associations and new forms of knowledge. Beyond the destabilizing abstractions of scale and the false allure of ease, there is a remaking of the mass communications network known as the Algorithmic Internet that we can hack into being. ♣

References

1. See figures from the Statista Research Department:
<https://www.statista.com/statistics/871513/worldwide-data-created/>
2. Calculated based on the number of tokens given in the original GPT-3 paper. The Common Crawl, WebText2, and Wikipedia were considered internet text, while Books1 and Books2 were not. <https://arxiv.org/abs/2005.14165>
3. Shannon Vallor. The AI Mirror: Reclaiming our Humanity in an Age of Machine Thinking. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society (AIES '22)*. Association for Computing Machinery, New York, NY, USA, 6. <https://doi.org/10.1145/3514094.3539567>
4. Andrey Kurenkov recounts the events in a blog post:
<https://thegradient.pub/gpt-4chan-lessons>
5. Ryle, Gilbert. “Knowing How and Knowing That: The Presidential Address.” *Proceedings of the Aristotelian Society* 46 (1945): 1–16.
<http://www.jstor.org/stable/4544405>
6. Deep Ganguli, Danny Hernandez, Liane Lovitt, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova Dassarma, Dawn Drain, Nelson Elhage, Sheer El Showk, Stanislav Fort, Zac Hatfield-Dodds, Tom Henighan, Scott Johnston, Andy Jones, Nicholas Joseph, Jackson Kernian, Shauna Kravec, Ben Mann, Neel Nanda, Kamal Ndousse, Catherine Olsson, Daniela Amodei, Tom Brown, Jared Kaplan, Sam McCandlish, Christopher Olah, Dario Amodei, and Jack Clark. 2022. Predictability and Surprise in Large Generative Models. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 1747–1764.
<https://doi.org/10.1145/3531146.3533229>
7. The Massive Multitask Language Understanding dataset covers 57 tasks that require “extensive world knowledge and problem solving ability.”
<https://arxiv.org/abs/2009.03300>
8. Anna Lowenhaupt Tsing. *The Mushroom at the End of the World: On the Possibility of Life in Capitalist Ruins*. Princeton University Press, 2021.
9. Alice E. Marwick and danah boyd. “I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience.” *New Media & Society*, 13(1), 2001, pp. 114–133. <https://doi.org/10.1177/1461444810365313>
10. Wendy Hui Kyong Chun. *Updating to Remain the Same: Habitual New Media*. MIT Press, 2017.
11. Often mentioned on Venkatesh Rao’s blog:
<https://studio.ribbonfarm.com/p/the-clockless-clock>
12. Bernard Stiegler refers to technology as a *pharmakon*, containing both poison and cure.

13. <https://www.clearview.ai/>
14. Among others: Otter.ai for transcription, Pymetrics for hiring, Stability AI (creators of Stable Diffusion) for generative models.
15. See figures from the Economic Policy Institute:
<https://www.epi.org/blog/growing-inequalities-reflecting-growing-employer-power-have-generated-a-productivity-pay-gap-since-1979-productivity-has-grown-3-5-times-as-much-as-pay-for-the-typical-worker/>
16. Karen Levy. 2022. RoboTruckers: The Double Threat of AI for Low-Wage Work. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society (AIES '22)*. Association for Computing Machinery, New York, NY, USA, 3.
<https://doi.org/10.1145/3514094.3539569>
17. Micah D. Carroll, Anca Dragan, Stuart Russell, and Dylan Hadfield-Menell. 2022. Estimating and Penalizing Induced Preference Shifts in Recommender Systems. In *Proceedings of the 39th International Conference on Machine Learning (ICLR '22)*. International Machine Learning Society.
<https://proceedings.mlr.press/v162/carroll22a.html>.
18. Liz Pelly. "Streambait Pop." The Baffler.
<https://thebaffler.com/downstream/streambait-pop-pelly>
19. Jenny Odell. *How to Do Nothing: Resisting the Attention Economy*. Melville House, 2019.
20. Laboria Cuboniks. *The Xenofeminist Manifesto: A Politics for Alienation*. Verso Books, 2018.