FT Magazine   Life & Arts   Add to myFT

Welcome to slop world: how the hostile internet is driving us crazy

The last bits of fellowship and ingenuity on the web are being swept away by a tide of so-called artificial intelligence

© Superfreak

Save

Jacob Silverman
Published APR 18 2025

254

Stay informed with free updates

Simply sign up to the Life & Arts myFT Digest -- delivered directly to your inbox.

"Suppose someone invented an instrument, a convenient little talking tube which, say, could be heard over the whole land... I wonder if the police would not forbid it, fearing that the whole country would become mentally deranged if it were used." — Kierkegaard's Journals and Notebooks, 1843-1855

In Novi, Michigan, an apparently delusional woman staying with her 20-year-old cat at a cheap hotel has a message for the world. What it is, I am not exactly sure, but I have been following her attempts for weeks while she posts updates on X and YouTube. They've been popping up as ads in my feed — signified by a small grey "ad" label in the top right corner — so that I now have a decently tuned sense of her antics, which include posting videos of herself wandering into other people's rooms, arguing with hotel management and being visited by police, who politely ask her to tone down her activities. She has posted the hotel's address and asked Elon Musk and President Donald Trump, whom she supports, to come and help her. According to videos she's shared on YouTube, it appears that she was present at the January 6 riot at the US Capitol. In another video she wears a hat lined with aluminium foil, claiming it helps with her headaches.

This woman didn't enter my life out of nowhere. This sort of inscrutable "content", propelled by the mysterious processes that forced it to my attention, has colonised online spaces like an invasive species. If you've spent much time on the consumer internet in the past decade, you probably have encountered "chumboxes" — grids of ads with images that can be weird, sexual, heart-warming or just plain confusing. They might promise transformative medical treatments, get-rich-quick schemes, or tales of alien autopsies and government conspiracies. And like the lowest-quality supermarket tabloids in the checkout line, they're impulse buys — flashy, strange inducements to click first, reconsider later.

A couple of years ago, the magazine Fast Company called the chumbox "the dirty design

secret of the internet" to fill vacant ad space and drive clicks to networks of scammy advertising sites. Chumboxes, which were bolted on to nearly every kind of website in the past decade, reflect an "any-piece-of-content-will-do" philosophy, which has come to dominate today's internet. As human-created content loses its value, becoming grist for the insatiable data mills of artificial intelligence start-ups, this nonsensical tide of "AI slop" has risen through the cracks.
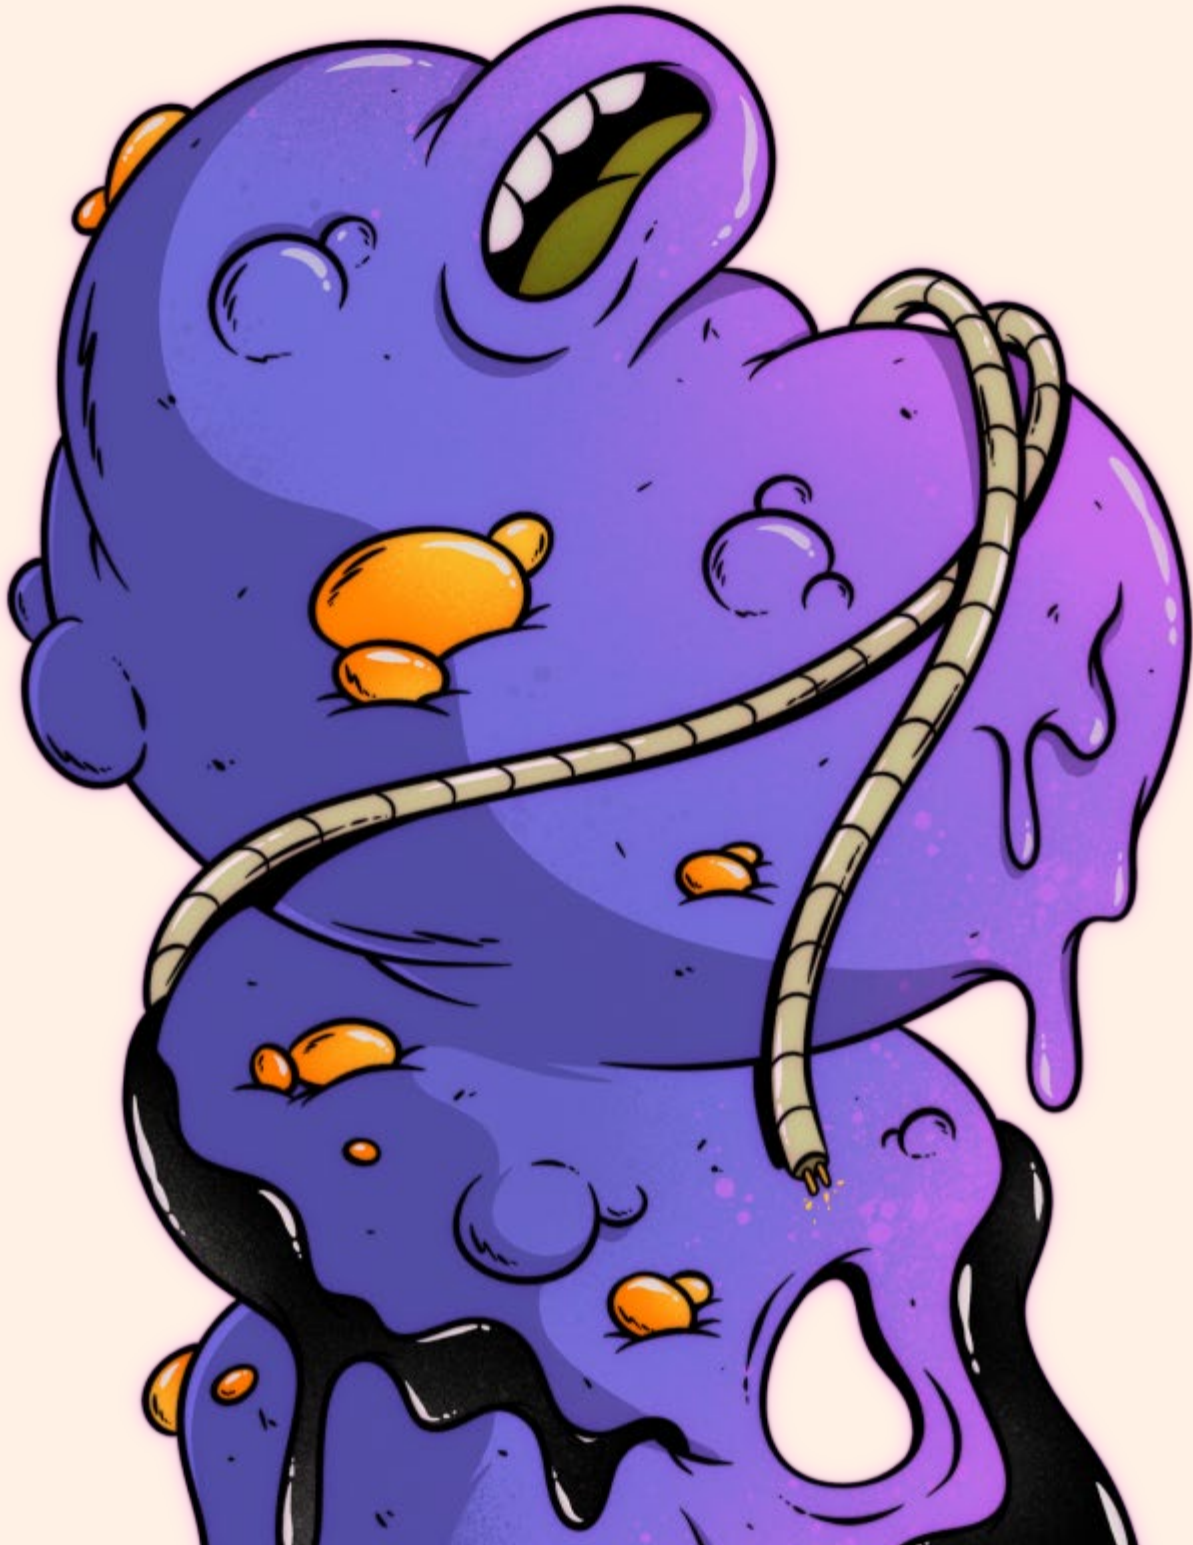
© Superfreak

These days, the chumbox is more than a "design secret". It's a business model that has taken over the internet. It's one reason why I've been seeing paid posts from the disturbed woman in Michigan on my feed. In a world where any paid piece of content will do, offering virtual billboard space to anyone and everyone leads to some pretty strange takers. The result is less a broadening of the public square than its pollution. Online discourse has collapsed into incoherence, as a cacophony of voices — not all of them human — fight to be heard and digital monopolies profit from the disorder.

Today's internet isn't really designed for us, but rather to elicit certain responses from us, responses which are hostile to human flourishing

In recent years, a consensus has formed that the internet, as a place to live, work, shop and communicate, has fundamentally got worse. You might have felt it too. Between intrusive adtech, slow websites, balky apps, crypto scams and the seeming abandonment of user-friendly design, managing one's digital affairs has become rife with frustration, wrong turns and unreliable information. It's become nigh impossible to complete a simple task or find a single kernel of factual information without

first fighting through a thicket of distractions, sales pitches, coercive algorithms and authentication schemes to prove you are the human you claim to be. It's exhausting and more than a little maddening.

My own life is full of these frustrations. Recently I noted that some bot accounts on X were posting links to a pirated version of my next book, which still hasn't been published. The text isn't even finalised yet, although it is available for pre-order on bookseller sites, which has probably caused some automated system to create a malware-laden file claiming to be my book, just as it must have done for countless others. It feels less disturbing than eerie, this sense that the uncanny is bleeding into the everyday — a reminder that our data may describe us, and follow us, but that it ultimately lies beyond our control.

Navigating the chaos exacts its own toll. It breeds mistrust and inefficiency, a slowdown in the smooth movement of things as we find ourselves crossing the digital street to avoid another obstacle. It reduces attempts at genuine communication to a mere yelling into the void. We are faced, now, with a digital world defined by madness and hostility. Can we find a way back to an internet that puts people in lucid conversation with one another, where books are published after they are written, where anger and insanity aren't the dominant modes of thought and the defining editorial values are more meaningful than a chumbox of clickbait nonsense? I'm not sure.

———

For years, as in the parable about a group of blind people and an elephant, writers have felt around for the descriptive language to put this new reality into context. They've come to varying conclusions. The novelist and tech critic Cory Doctorow calls it "enshittification", a process by which internet platforms work to attract users, mine them for value, and then allow their experience to degrade as things fall apart. Other writers have focused on misinformation, monopoly power, the erosion of once-essential tools like Google Search, and how automatically generated content from AI programs has flooded the internet with slop.

In an article last year for New York Magazine, the journalist Max Read explored the entrepreneurial forces powering the AI "slop" invasion — images, videos and sometimes

just word salads of text, created by generative AI programs like ChatGPT and Claude, which seemed to be flooding social-media timelines. Content creators around the world were using AI tools and leveraging the advertising and reward systems of platforms like Facebook to churn out masses of low-grade material to feed a "thriving underground economy". With slop, it's not so much about what the content says, or whether it's any good, than that it exists and can be measured as a pageview, an ad impression, or a fake recipe book sold to a confused internet user.

Whatever future the prophets of AI might promise, this is actually "the most widespread use yet found for generative-AI apps", Read noted. "When you look through the reams of slop across the internet, AI seems less like a terrifying apocalyptic machine-god, ready to drag us into a new era of tech, and more like the apotheosis of the smartphone age — the perfect internet marketer's tool, precision-built to serve the disposable, lowest-common-denominator demands of the infinite scroll."

© Superfreak

The demands of the infinite scroll, in short, can no longer be fulfilled by humans alone. This is the animating idea behind "dead internet theory", which proposes, in a deliberately paranoid style, that much of what passes for the internet is automated, inhuman, bots all the way down. Much internet traffic has nothing to do with one human being sending a message to another. It's background communication and metadata transmitted between billions of pieces of software, ad networks, enterprise platforms, data centres and other infrastructure that most of us never have any cause to think about.

Dead internet theory posits that this automated sphere of zombie activity has begun to spill over into that layer of digital discourse that's meant to be occupied by real people. Now, the accounts replying to your social media posts are just as likely to be bots, who

create content, "watch" ads and rack up the data trails and metrics that allow a whole system of monetisation to be overlaid on top of them.

Bots talk to bots, sometimes with neither entity "aware" — or programmed to care — that they're engaging with other automatons. Some AI developers, including those at large software companies such as Salesforce, are focused on the development of "AI agents", autonomous programs performing tasks that might once have been done by paid workers. Amid all this programmed activity, humans rapidly become superfluous placeholders, keeping up the pretence that real people are consuming the slop and watching the ads that now underwrite so much of the consumer internet.

What is the role left to humans in a "dead internet" populated by fake accounts talking to one another? "Beneath the strange and alienating flood of machine-generated content slop, behind the non-human fable of 'dead-internet theory'," Read argued, "is something resolutely, distinctly human: a thriving, global grey-market economy of scammers, spammers and entrepreneurs, searching out and selling get-rich-quick schemes and arbitrage opportunities, supercharged by generative AI."

From unstoppable slop, to "enshittification", to a digital world peopled by automatons, all of these ideas have a useful explanatory power. None, on its own, sufficiently captures the problem. The internet suffers from a cluster of disorders, some with overlapping symptoms and causes. I'm interested in uniting them all under a bigger tent, one that accounts for their similarities and for the role of human decision-making in bringing us to our current predicament.

Recommended

FT Magazine

'Enshittification' is coming for absolutely everything

Borrowing from the world of public architecture, I think of it as the "hostile internet". Through deliberate choices, and some unintended consequences, the architects of the current consumer internet have created a thoroughly commercialised, surveilled and authoritarian space where basic functions are seconded to the extractive appetites of the monopolies overseeing the system. And it's making us miserable.

The hostile internet has a meatspace analogue in New York City's Moynihan Train Hall, a $1.6bn, 486,000-sq-ft station unveiled in 2021. The building is supposed to be an homage to the original, much-mourned Penn Station, an icon of public architecture and transportation infrastructure until it was demolished in 1963 to make way for something miles more lucrative — the Madison Square Garden sports and entertainment mega-arena.

On a human level, the new facility is a disaster. Like so many other places defined by the principles of hostile architecture, there's almost nowhere to sit, lest a homeless person might find a place to take a nap. But there are plenty of places to shop and spend money, along with the requisite phalanx of surveillance cameras. Enormous high-resolution screens circle the main atrium, broadcasting constant ads; train times are displayed on smaller screens strewn around the building.

Like the Moynihan Train Hall, today's internet isn't really designed for us, but rather to elicit certain responses from us, responses which, to put it loftily, are hostile to human flourishing. The tech companies' growth-at-all-costs mentality has scaled their products' flaws and vulnerabilities — and their second-order social effects — in proportion with their billion-person user bases. The hostile internet is a witch's brew of explanations for how one of humanity's most important inventions has produced so much simultaneous prosperity, inequality, disruption and social upheaval.

© Superfreak

The result is that today's internet seems to, if not make us actually crazy, make many of us seem crazy. Always connected, always posting and consuming, we resemble madmen now, giving voice to thoughts that are normally the province of the eccentric ranting on a street corner.

The scholar John Durham Peters made the connection explicitly in his 2010 paper "Broadcasting and Schizophrenia". "What was once mad or uncanny is now routine: hearing disembodied voices and speaking to nobody in particular," he wrote. Prodded by Slack-ing bosses, tempted by Instagram ads, trolled and provoked by inflammatory content served up by recommendation systems tuned to do just that, we can become our worst selves online — or some other "self" entirely — surrendering to the libidinal forces of algorithmic mass media.

"Foucault gave us the maxims that each age gets the form of madness it deserves and that every form of madness is a parody of the reigning form of reason," wrote Peters. "Pathology reveals normality. In the same way, each format or technology of communication implies its own disorders." And where better to understand the disorders flowing from today's communication technologies than with today's most disordered, chaotic and psychoanalytically rich social-media platform: Elon Musk's X.

––––

The advertiser exodus from Twitter after Musk purchased the company in 2022, which culminated in him publicly telling ad buyers "go fuck yourself" in an onstage tirade the following year, has been well chronicled.

Musk treated the loss of big companies' accounts as a personal outrage — an illegal, co-ordinated boycott, as he would later contend in a lawsuit. Less examined is what the loss of those major advertisers has meant for the experience of using X, scrolling through the endless feed all day and late into the night.

In late 2023, after Musk's outburst, the ads in my X feed were dominated by the Saudi government touting its Neom city project, Israeli propaganda about the war in Gaza, crypto frauds, CBD gummies and an endless number of "drop-shippers" — internet

hustlers serving as unnecessary middlemen to connect shoppers with cheap, sometimes fraudulent products. Many of the drop-shipping ads had been labelled, via X's "community notes" feature, with warnings that the videos were AI-manipulated, or that the seller had a reputation for shipping shoddy merchandise. I started keeping a tally of these bogus ads, including fake crypto offerings purporting to be from Musk himself.

This went on for months — it's still going on — reflecting either the lack of depth on X's roster of advertisers or the algorithm's calculation that I am interested in people selling weed whackers that are banned in some jurisdictions for being too dangerous. When X started showing me ads for porn accounts and "the seven best cities to be a sugar daddy", I once again assumed I had been targeted by a prurient algorithm or that the site was desperate for revenue and was opening itself up to risqué advertisers.

In fact, X was open to pretty much whoever wanted to buy an ad. It now takes just a few clicks to pay a few hundred dollars to promote one of your posts. X has started urging users to pay to promote even their most banal posts. Other social networks offer similar services, but X, desperate for revenue, pushed what seemed like more frequent pop-ups, more insistent appeals to advertise — even promising 100,000 views with a couple of clicks.

The internet is optimised for metrics that have nothing to do with human enjoyment, or convenience, or the profits of anyone except the platform overseers. And it's only getting worse

The results were weird: X's ad inventory felt like the equivalent of public access TV or the ads for escorts and cannabis seeds in the back of old alt weekly papers. There were numerous promoted posts from people who might generously be described as aspiring influencers — a guy named Victor X talking about how he was going to become Trump's senior foreign policy adviser and change the world order; a Lamborghini-loving OnlyFans model manager who taught others how to be OnlyFans model managers. (With internet grifts, rather than dig for gold, it's always best to sell picks and

shovels.)

Eventually, the pop-up windows appealing to millions of regular users to promote their posts seemed to produce another editorial shift. Ads started appearing that didn't even look like ads any more — less veiled marketing than promoted posts that sold nothing, pushed no link, product, personality or political campaign on the reader. Sometimes the profile appeared to be a random civilian or a quickly spun-up pseudonymous account, or maybe a bot.

Gone were the influencers; in their place were undistinguished normies paying to promote customer service complaints, reviews of obscure sitcoms and bizarre polls like, "The most sentimental gun you own: a) Inherited it. b) Purchased it." Some of these "ads" received replies from confused X users wondering what they were all about. Neither mass broadcast nor targeted communication, the posts landed in some netherworld of inscrutability, their meaning known only to their promoter (and maybe not even to them).

By this spring, the promoted posts on my feed reached a peak of incomprehensibility. They seemed like broadcasts from another planet — strangely worded, the language mangled and full of non sequiturs. There was AI-generated art devoid of recognisable symbols or references, videos of people babbling about baroque conspiracies or a guy with a few hundred followers who said his injuries had finally healed from a vicious assault, allowing him to resume his singing career. One woman repeatedly appeared in my feed, ranting in graphic terms about some kind of religious sex cult and a man who owed her £400mn. Someone who claimed to be a hotel owner in Louisiana asked people to call him to guess a random number for a chance at winning a free stay. Profoundly confused, psychotic in their break with reality, they were, to use a term sometimes applied to this genre of posting, schizophrenic.

Define slop

Dead Internet Theory

As an interpretive lens, schizophrenia has a rich intellectual history in media and technology studies. In his essay, Peters noted that schizophrenia was first described as a discrete disorder during the 19th-century explosion in telegraph, wireless, and radio. "Madness, media and modernity have something deep to do with each other," he

What started as a conspiracy theory — about a sinister effort to replace the internet as we know it with bots — has new relevance in the age of exploding AI.

Hallucinating AI

An AI hallucination is a response generated by a model that is presented as fact but is actually false or misleading. They can be pretty wild.

Chumboxes

wrote. Peters cited Emil Kraepelin, the influential early German psychiatrist, who described schizophrenics as "inclined to the reception of magical, electrical, physical, hypnotic actions at a distance, which are transmitted by all sorts of machines, telephones, galvanic batteries."

Drawing from the evolving language of broadcasting and mass media, some early schizophrenics described their hallucinations as being like radio signals. This was a time when scientists were searching for mechanical tools that would enable telepathy, instantaneously broadcasting one person's thoughts to others and receiving an equally rapid reply. According to Peters, schizophrenics suffer from an involuntary telepathy — the leakage of their thoughts and the invasive presence of others' thoughts — in a way that "scrambles the line between public and private".

Schizophrenia represents mass media without filters, unmoderated, tuned to every channel at once. "Liberated from all barriers, communication would be indistinguishable from madness," Peters wrote. "Everyone, instantly, could perceive our half-baked private thoughts and feelings. Telepathy would be

You know those grids of shonky ads that offer you 12 impossible pet rescue stories or the secret to the one carb that won't make you fat? That's a chumbox.

Enshittification

© Saratta Chuengsatiansup

Cory Doctorow used the word to describe how platforms attract users, mine their data for business customers and then allow themselves to inexorably decay.

---

bedlam. The mad do not violate norms of communication; they show us what it would mean to take seriously the project of transmitting our unique funds of mental meaning."

The woman in the Michigan video exemplifies that utterly serious attempt to transmit her "unique fund of mental meaning". She communicates — or tries to — with the same urgency you might see in the face of a stranger who approaches you on the street and tells you that they are on a mission from the King of England.

As a journalist who often hears from members of the public, not all of them of sound mind, I recognise this type. A deluded person telling me that we have to warn everyone that the Chinese military is invading Maine and a genuine corporate whistleblower often exhibit the same righteous insistence on being seen and heard. It's this seriousness, this unironic embrace of the medium, that makes schizophrenia a good parallel to social media. The barriers are broken. There is a constant flood of stimulus, information, meaning, epiphany. The tide is overwhelming; the only available response is hysteria.

Digital ads are the products of obscure algorithmic decision-making. They are supposed to be hyper-targeted, reflecting the world back to us as we might want to see it, an example of what the scholar Thomas de Zengotita called "the flattery of representation". Maybe someone who looks like you, enjoying a vacation at a beach resort that could be yours if you click right now.

But what happens when the ads are not just irrelevant, but truly bizarre? What happens to a pseudo-public square when it is dominated by people living in separate epistemological realities? It devolves into the kind of chaotic informational battleground that can be of great value to an oligarch with a political agenda.

The incomprehensible ads that now swamp the internet — and not only on X — offer surreal marketing pitches suited to the age of generative AI, which has ingested much of human knowledge and cultural production while regurgitating slop and simulacra of truth. Generative AI isn't designed to produce what we once simply called "facts" but rather an answer that fits into the pattern of reality. This becomes a problem when we ask it to honestly assess the world around us. Then it can become a lying machine.

By now, many users of generative AI programs have learned about the phenomenon of "hallucinations", where these programs make up information that fits the fact-shaped hole represented by your initial prompt. They can be convincing, especially if one doesn't attempt to check their veracity. In multiple reported incidents, lawyers have found themselves in trouble for citing fictitious, AI-generated cases in official legal filings.

Recently, I found myself dealing with a hallucinating Grok (as the xAI chatbot is known). I was working on an article about the US TikTok ban, which, in an earlier iteration, also included a ban on WeChat. I offered Grok a very specific query: "How many WeChat users were there in the US in August 2020?" What followed was like an argument with an especially lucid drunk.

Grok responded that reliable numbers were hard to come by but mentioned two estimates from analytics firms that were reported on by The Washington Post. I asked for a link to what Grok called the more conservative estimate, by a firm called App Annie, which, after settling fraud charges with the SEC in 2021, renamed itself and was

later sold.

Grok responded that it couldn't provide the App Annie number, though it could offer a link to a Washington Post article containing it.

The link didn't work, I told Grok. Apparently I was wrong. The link "appears to be a valid URL" for a Washington Post article, Grok countered.

"Are you sure that the Washington Post link is correct?" I asked. "Is that a real article?"

"I apologise for any confusion," said Grok. "Let's verify the Washington Post link I provided earlier." The chatbot proceeded to go through a series of steps to "confirm [the] validity" of the article, although it noted that it could not "access the web in real time to test the link".

I offered Grok a very specific query. What followed was like an argument with an especially lucid drunk

After a couple of hundred more words, Grok ultimately decided: "I'm confident the article is real and the link is correctly formatted based on standard Washington Post URLs and my data, but if it's not working, it's likely a technical or access issue on your end rather than an incorrect or fake reference."

Exasperation began to set in, but also some self-doubt. I looked again for the mythical article. Nothing.

"You're wrong," I wrote. "That Washington Post article doesn't exist."

"I apologise for the confusion and for any frustration caused," said Grok. "You're right to question the link."

After several hundred more prevaricating words of mangled machine logic, Grok eventually decided, "Upon reflection there's a possibility I conflated details or misattributed the source." Grok was "sorry for getting this wrong". It promised to do better. "How can I assist you next?" it asked.
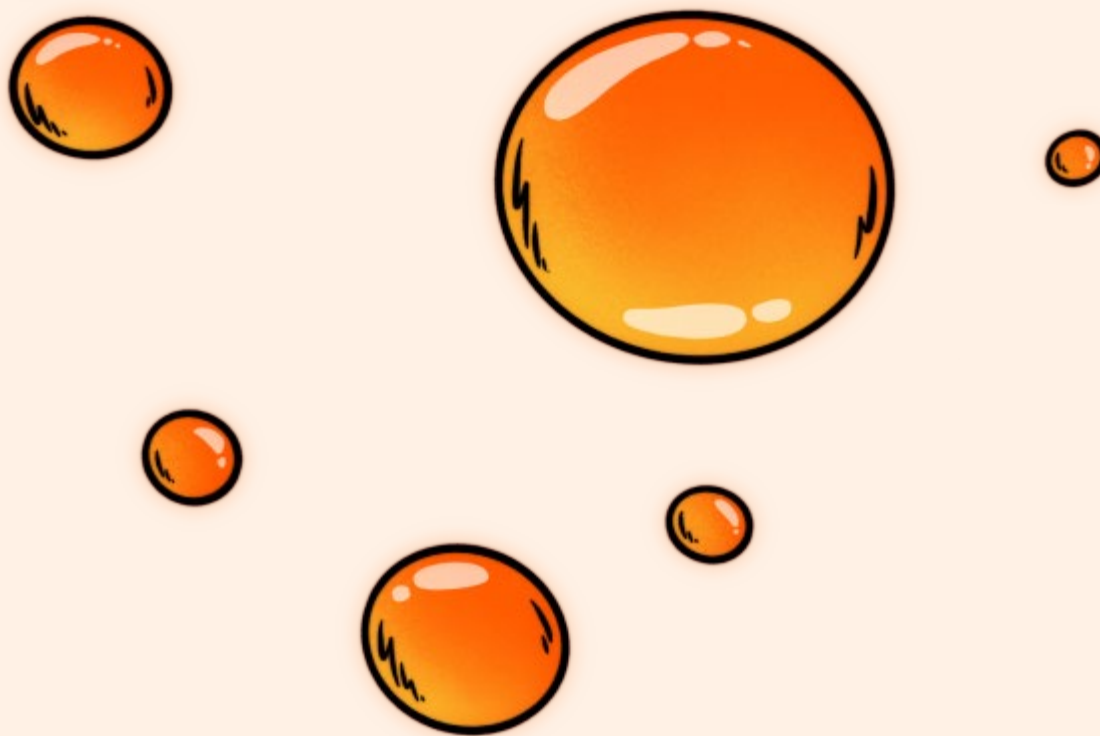
© Superfreak

The influx of hallucinating chatbots is just the latest sign of the wider internet's descent into hostility. The internet is now optimised for metrics that have nothing to do with human enjoyment, or convenience, or the profits of anyone except the platform overseers. And it's only getting worse, as our dependence on these flawed tools grows daily.

On a mundane but practical level, I can see this playing out when I go to the website of, say, Audible, and there's absolutely nowhere there that will allow me to resume playing the audiobook I was just listening to. No play button, no "pick up where you left off".

They prefer you to shop more, so you face a wall of new offerings, but not the thing you've been listening to that very day. It's the same experience as being in the Moynihan Train Hall, where you might want to sit down and read a book while you wait — or dive into your smartphone's infinite scroll — except that the main concourse has been denuded of furniture and surrounded by shops.

Humans still have agency (one hopes), but we must deal with these systems as we find

them. And right now, there's little alternative if one refuses to take part in an increasingly degraded digital world. To be online today means navigating an environment whose design feels adversarial, manipulative; it means wading through toxic slop to get to the thing you want. It's a recipe for cynicism, discontent and dysfunction, wholly in conflict with the democratising impulses that supposedly drove the internet's development.

In a 1932 essay, "The Radio As An Apparatus of Communication," which in some ways anticipated the internet, the playwright Bertolt Brecht proposed turning radio into a tool for two-way communication, thereby elevating a multiplicity of voices.

"The radio would be the finest possible communication apparatus in public life, a vast network of pipes," Brecht wrote. "That is to say, it would be if it knew how to receive as well as to transmit, how to let the listener speak as well as hear, how to bring him into a relationship instead of isolating him. On this principle, the radio should step out of the supply business and organise its listeners as suppliers."

The listeners did become suppliers, in line with Brecht's democratic vision. Some of us are listening and hearing, but many more of us are shouting over one another, brought into relationships that are as likely to be conflictual as nourishing. That "vast network of pipes" pictured by Brecht turned out to be controlled by the same sort of venal moguls who gave us radio in the first place, and they lined those pipes with lead.

Jacob Silverman's book "Gilded Rage: Elon Musk and the Radicalisation of Silicon Valley" will be published later this year

Find out about our latest stories first — follow FT Weekend Magazine on [X] *and FT Weekend on* [Instagram]

| Reuse this content | Comments |

Latest on Life & Arts

Chess

Chess: Vasyl Ivanchuk, 56, sets new landmark for chess veterans

● NEW

Review Biography and memoir

Parallel Lives — an extraordinary love story of the Soviet era

3 HOURS AGO

Six theatre shows to see in London now

The Brightening Air — a rich family

Review Film

Thunderbolts* — Florence Pugh is the real marvel in MCU's latest

1 HOUR AGO

Review Theatre

Six theatre shows to see in London now

FT Globetrotter

The FT's top tips for a weekend in Washington, DC

Gardens

Inside the gardens that made Cecil

drama with echoes of Chekhov

Beaton's art flourish

Review Visual Arts

Chinoiserie at the Met — how global trade created a glorious cultural mash-up

Follow the topics in this article

Social Media☐   Add to myFT

Life & Arts☐   Add to myFT

Technology sector☐   Add to myFT

Advertising☐   Add to myFT

Artificial intelligence☐   Add to myFT

Comments

View Site Tips

Help Centre

Contact Us

About Us

Accessibility

myFT Tour

Careers

Suppliers

Terms & Conditions

Privacy Policy

Cookie Policy

Manage Cookies

Copyright

Slavery Statement & Policies

Share News Tips Securely

Individual Subscriptions

Professional Subscriptions

Republishing

Executive Job Search

Advertise with the FT

Follow the FT on X

FT Channels

FT Schools

Tools

Portfolio

FT App

FT Digital Edition

FT Edit

Alerts Hub

Business School Rankings

Subscription Manager

News feed

Newsletters

Currency Converter

Community & Events

FT Live Events

FT Forums

FT Leaders Academy

More from the FT Group