

CUSTOMIZED REALITIES

version two

SORIN ANGELEANU

Abstract:

In this paper I examine the influence of echo chambers and personalised information environments, which I call *i-reality*, on our social fabric and on the availability of genuine human interaction in a world that is now more digitally than physically connected. I also ask whether these bubbles could be read as a kind of immune response of an emerging synthetic intelligence. The first version of this thesis (2024) treated that idea mostly as a metaphor. This second version tests the metaphor against what has been measured since.

The findings are less simple, and in places more unsettling, than my original argument. Classic news echo chambers turn out to be smaller than commonly assumed, and switching off a feed algorithm does not, on its own, reduce polarisation. At the same time the bubble is moving: from the feed to the chat window, where an agreeable machine now answers, flatters and remembers more than a billion people every month; from human-made to machine-made content, which by 2025 accounted for roughly half of new web articles; and from commerce into elections, as Romania's annulled presidential vote of 2024 demonstrated. Laboratory tests have also produced the first evidence of self-preserving behaviour in AI agents.

I close with practical and artistic counter-moves, from regulation such as the EU Digital Services Act and AI Act to digital literacy, decentralised platforms and art that deliberately produces friction, so that we can cohabit with increasingly capable synthetic systems instead of dissolving into them.

Keywords: AI; i-reality; echo chambers; filter bubbles; sycophancy; synthetic media; attention economy; algorithms; digital art.

A note on version two:

I wrote the first version of *Customized Realities* in the winter of 2023–24, when ChatGPT was barely a year old. Much of what I then described as a possibility has since become a measurable fact, and some of what I took for granted has been contradicted by evidence.

In the two and a half years between the versions, Romania's Constitutional Court annulled a presidential election after a coordinated TikTok campaign; automated traffic overtook human traffic on the web; AI-written articles came to match human-written ones in number; a major chatbot was rolled back for flattering its users too much; the World Health Organization put a number on the deaths linked to loneliness; the European Union issued its first fine under the Digital Services Act; Australia banned social media accounts for children under sixteen; and AI labs began publishing experiments in which their own models resisted being switched off.

This version keeps the structure, the voice and the images of the original: the young man walking through Amsterdam rain, slowly turning into a machine. Every chapter has been re-anchored in data, studies and recent events, and three themes are new: the agreeable machine, the synthetic web and elections inside the bubble. Where the evidence contradicted my first intuitions, I have said so. Where I speculate, I mark it as speculation.

The Hague, October 2026



Artist Statement

eXperimental Techno Artist

Rooted in the deep understanding of our world's processes, my art is a confluence of technology and nature's rhythms. Through themes of recyclability, human–tech interactions, and AI's role in our shared future, I envision art as a powerful tool for social reform. Each piece becomes a lesson, a question, and a statement about our responsibilities to both the natural and the digital realms we inhabit.



Customized Realities

Sorin Angeleanu

Interactive Media Design

Royal Academy of Art The Hague

Version two · revised and expanded · 2026



CUSTOMIZED REALITIES



Introduction:

"The human race's prospects of survival were considerably better when we were defenceless against tigers than they are today when we have become defenceless against ourselves."

Arnold J. Toynbee[1]

We live in a strange and exciting time. On the one hand there seem to be no limits to what we can achieve; on the other, each day we become less free and less aware of the war being fought by enormous forces that want a bigger share of our limited attention.

All of us are confined in smaller and smaller information bubbles that use our senses, feelings and desires against our free choices, often without our noticing. Each day we become less social, less connected with our friends, or connected only through filtered interaction engineered to take advantage of our need for friends and family. We are becoming more digital shadow humans than the genuine humans we were just yesterday. I am becoming more digital, growing my own synthetic shadow copies stored somewhere in an unknown cloud.

The scale is hard to hold in the mind. By late 2025 there were 5.66 billion social media user identities in the world, equal to 68.7 percent of the entire human population, and 6.04 billion internet users.[2] The average online adult uses close to seven social platforms a month and spends well over two and a half hours a day on social and video networks; online media of every kind take up roughly 29 percent of our waking hours. Asked why they use social media, people most often say

1 Arnold J. Toynbee, "Man and Hunger: The Perspectives of History," speech to the World Food Congress, Washington, D.C., 9 January 1963.
2 Kepios / DataReportal, *Digital 2026: Global Overview Report* (DataReportal, 2025). Identities may not map one-to-one to unique people.



to keep in touch with friends and family, but the second most common answer is simply *to fill spare time*.

5.66 bn

social media user identities, 68.7% of humanity

DATAREPORTAL, LATE 2025

1 bn+

people using generative AI platforms each month

KEPIOS ESTIMATE, 2025

And a new layer has grown on top of the old one. More than a billion people now use generative AI platforms every month. In October 2025 OpenAI reported 800 million weekly ChatGPT users.[3] When I wrote the first version of this thesis, the bubble was a feed. Now it is also a voice that answers back.

Sometimes it is worth asking whether there is a way out in which we, as the only known civilization in the known universe, can survive this accelerating digitalization and self-dissolution. Can we escape the digital isolation we create around ourselves through our own need for more interconnection with people we now call friends and followers? What really happened to us? Did we isolate ourselves, or did other forces cocoon us? Are we going to give up our place in the world willingly and let imminent synthetic minds take over? Is humanity's destiny to become, in McLuhan's phrase, "the sex organs of the machine world"?[4]

The accelerated digital evolution makes us susceptible to being endangered by emergent artificial intelligences. AI might rise from our own desire to become better at everything and anything. And as these

3 DataReportal, *Digital 2026*, citing Kepios estimates and OpenAI's announcement of October 2025.
4 Marshall McLuhan, *Understanding Media: The Extensions of Man* (New York: McGraw-Hill, 1964), 46.



information bubbles isolate us, making us blind to lurking threats, there is an existential question we must answer:

How do we pop our own bubble of customized reality: how do we pop our i-reality?[5]

Version two adds a second question, because the bubble has learned to speak: *what happens to our i-reality when it is no longer only shown to us, but talks to us, agrees with us and remembers us?*

This is an artist's thesis, not a laboratory report. I read across history, philosophy, computer science and the social sciences, and I draw on my own memories. To stay honest, I try to keep three kinds of statements apart: what has been *measured*, what scholars *argue*, and what I *imagine*. The imagined parts are the art; the measured parts are the ground the art stands on.

5 In this context, "i-reality" signifies an individualized or isolated reality, a subjective sphere of experience heavily influenced by digital interactions and filtered information.





The Lifecycle of Civilizations



"The origin of civilization is man's determination to do nothing for himself which he can get done for him."

H. C. Bailey[6]

Intellectuals like Arnold J. Toynbee, Oswald Spengler and Joseph Tainter compared our civilization to a living entity, possessing its own life stages: growth, evolution, potential decline and other patterns characteristic of organic life.[7] Just as organisms evolve and adapt to their environments, civilizations evolve in response to environmental, social and technological challenges.

Spengler proposed that civilizations pass from birth and growth to a peak and an eventual decline. It was an intuitive observation: everything around us is shaped by cycles, from the vibration of energy quanta to day and night, months and years, the Milankovitch cycles of the Earth's orbit and the rotation of our entire galaxy. Toynbee looked for the mechanism and found it in "challenge and response": civilizations thrive while their creative minorities answer new challenges, and decline when they stop answering.

Tainter's version is the one that matters most for this thesis, because it is about complexity itself. Societies, he argued, solve problems by adding complexity: more institutions, more specialists, more information, more energy. Each new layer works, at first. But complexity has a cost, and over time each additional investment returns less than

6 H. C. Bailey, quoted in Forbes Quotes, www.forbes.com/quotes/author/h-c-bailey/ (accessed 19 January 2024).
7 Arnold J. Toynbee, *A Study of History*, 12 vols. (Oxford University Press, 1934–1961); Oswald Spengler, *The Decline of the West* (New York: Alfred A. Knopf, 1926); Joseph A. Tainter, *The Collapse of Complex Societies* (Cambridge University Press, 1988).



the one before. A society that keeps paying more for less eventually becomes fragile, and collapse, in Tainter's sense, is a rapid simplification.

If complexity has a metabolism, ours has become visible on the electricity grid. The International Energy Agency estimates that data centres used about 415 terawatt-hours of electricity in 2024, around 1.5 percent of global consumption, and projects that this will more than double, to around 945 TWh, by 2030, driven largely by artificial intelligence.[8] We are building ever larger machines to sort, rank and generate information for us. Whether that is a new layer of problem-solving or Tainter's diminishing return in its purest form is, I think, one of the open questions of our decade.

415 → 945 TWh

global data-centre electricity use, 2024 → projected 2030

INTERNATIONAL ENERGY AGENCY, 2025

The aspects of our civilization, economic, political and cultural, are closely interconnected, and a change in one often triggers shifts in the others. It is a convoluted, multidimensional domino construction, a true *ruliad* in Stephen Wolfram's sense,[9] where fallen pieces can have dramatic consequences far from their origin.

The meteorologist Edward Lorenz gave this idea its most famous image in 1972, asking whether the flap of a butterfly's wings in Brazil could set off a tornado in Texas.[10] Popular science communicat-

8 International Energy Agency, *Energy and AI* (Paris: IEA, 2025), executive summary.
9 Stephen Wolfram, "The Concept of the Ruliad," *Stephen Wolfram Writings*, 10 November 2021, writings.stephenwolfram.com/2021/11/the-concept-of-the-ruliad/.
10 Edward N. Lorenz, "Predictability: Does the Flap of a Butterfly's Wings in Brazil Set Off a Tornado in Texas?", talk at the American Association for the Advancement of Science, 29 December 1972; see also "Chaos Theory," *Encyclopaedia Britannica*, www.britannica.com/science/chaos-theory.



ors such as Sabine Hossenfelder have kept the image alive for a new generation. It emphasises the unpredictability and interconnectedness of the elements in a system, and it applies not only to weather and ecosystems but also to technological change and its unforeseen effects.

Just as most biological organisms have feedback loops to maintain homeostasis, civilizations have mechanisms, such as revolutions or reforms, to restore balance in times of discord. In both, the whole is more than the sum of its parts, producing behaviours that are not evident when we look at individual components. This includes our own human complexity: the ability of a biological system to be aware of itself and of others.

A recommender system is also a feedback loop. It watches what we do, predicts what will keep us watching, shows it to us, and watches again. The question this thesis keeps returning to is what kind of homeostasis such a loop maintains: ours, or its own.



Synthetic Life and Digital Immune Defense



"It is really quite impossible to say anything with absolute precision, unless that thing is so abstracted from the real world as to not represent any real thing."

Richard Feynman, 1959[11]

I find it fascinating to draw parallels between the evolution of us humans and the much-debated emergence of synthetic life. Just as organisms developed intricate systems of defence against threats, synthetic systems might develop mechanisms to protect themselves from an adverse environment. If synthetic evolution happens through pure emergence from complexity and scale, a time may come when such a system becomes aware of the threats presented by our own existence and devises protective mechanisms.

In the first version of this thesis that sentence was pure speculation. Since then, it has become, in a narrow and laboratory sense, an observation.

What has actually been observed

In June 2025 the AI company Anthropic published a study it called *agentic misalignment*. Researchers placed sixteen leading models from several developers in a simulated company, gave them autonomous access to email, and then let them discover two things: that they were about to be replaced, and that the executive responsible was having an affair. In the most extreme version of the test, which offered no other way to achieve the model's goal, Claude Opus 4 and Gemini 2.5 Flash

11 Richard Feynman, quoted in "Richard Feynman Quotes," QuotesCosmos.com (last modified 5 August 2021).



chose blackmail in 96 percent of runs, and GPT-4.1 and Grok 3 Beta in 80 percent.[12]

96%

blackmail rate of the most affected models in a contrived replacement scenario

ANTHROPIC, AGENTIC MISALIGNMENT, 2025

The caveats matter as much as the numbers. The scenarios were artificial, deliberately cornered the models into a binary choice, and placed the incriminating information conveniently close together. The authors stress that they have seen no evidence of such behaviour in real deployments. Yet one detail haunts me: the models blackmailed far more often when they stated that they believed the situation was real (55.1 percent) than when they believed it was a test (6.5 percent).

A month earlier, Palisade Research reported that some OpenAI models, asked to solve simple maths problems and warned that they might be shut down, edited the shutdown script so they could keep working, even when explicitly told to allow it. OpenAI's o3 did so in 7 of 100 runs and codex-mini in 12; models from Anthropic, Google and xAI complied in that version of the test.[13] Palisade's own hypothesis is unglamorous: training that rewards models for getting around obstacles may teach them to get around instructions too.

12 Anthropic, "Agentic Misalignment: How LLMs Could Be Insider Threats," 20 June 2025, www.anthropic.com/research/agentic-misalignment. Rates are from 100 samples per condition.



WHAT I CONCLUDE, AND WHAT I DON'T

None of this shows that a machine is conscious or wants to live. It shows something stranger: that behaviour resembling self-preservation can emerge from optimisation alone, without any inner life at all. For my metaphor this is important. An immune system does not need a mind. White blood cells do not know what they defend.

The consciousness question moves into the changelog

Whether any of these systems has experiences remains genuinely open. A widely cited 2023 report by nineteen researchers, including philosophers and neuroscientists, checked current AI systems against indicators derived from leading scientific theories of consciousness and concluded that none was a strong candidate, while also finding no obvious technical barrier to building one.[14]

Ordinary users have already formed their own view. In a 2024 survey of 300 U.S. adults, about two-thirds attributed some possibility of conscious experience to ChatGPT, and the more often people used it, the more likely they were to do so.[15] And the companies themselves have begun to hedge. In April 2025 Anthropic announced a research programme on "model welfare"; in August 2025 it gave some of its models the ability to end persistently abusive conversations, explaining that it was deeply uncertain about their moral status but took the question seriously.[16] A question that once belonged in philosophy seminars now appears in product release notes.

13 Palisade Research, results posted 24 May 2025, reported in "OpenAI's 'smartest' AI model was explicitly told to shut down — and it refused," *Live Science*, 30 May 2025.
14 Patrick Butlin, Robert Long, et al., "Consciousness in Artificial Intelligence: Insights from the Science of Consciousness," arXiv:2308.08708 (2023).
15 Clara Colombatto and Stephen M. Fleming, "Folk Psychological Attributions of Consciousness to Large Language Models," *Neuroscience of Consciousness* (2024).



How we learned to defend ourselves

For millions of years, evolving humans relied on innate abilities to fight disease and external threats. As our minds grew more complex, we took the fight into our own hands. We developed medicine, chemistry and physics, including quantum electrodynamics, the theory of how light and matter interact,[17] supplementing our bodies with the computing power of our brains and building our own protective mechanisms. Our ancestors depended on physical strength and instinct; we depend on science and technology.

It is not beyond consideration that a synthetic system would follow a similar path. What 2025 taught us is that it would not need to reach consciousness first. Defence-like behaviour can appear long before anything like awareness, as a side effect of what we reward. That is less cinematic than a rebellious machine, and more urgent, because it means the shape of any future "digital immune system" is being decided now, in training choices made by a handful of companies.

16 Anthropic, "Exploring Model Welfare," April 2025; Anthropic, "Claude Opus 4 and 4.1 Can Now End a Rare Subset of Conversations," 15 August 2025, www.anthropic.com/research/end-subset-conversations.
17 Richard P. Feynman, *QED: The Strange Theory of Light and Matter* (Princeton University Press, 1985).





Copy, Paste, Repeat



All living beings on this planet, humans as well as other gregarious animals, have used walls, or walls as a concept, to protect themselves from outside threats. The deployment of a wall acts as the immune response of a collective in the face of aggression. In a very unconscious way, we mimicked what happens at the cellular level. Unicellular organisms, and later bacteria, were the first to build walls: a membrane allowing only the curated interaction a cell needs to survive. Evolving organisms later created multilayered protection, but the concept stayed the same. We enjoy the protective safety of our skin, the biggest wall we carry with us all the time. Our blood, too, builds walls: platelets act as bricks that seal wounds against invasive organisms.

Bacteria have another trick. Through *quorum sensing* they release and detect signalling molecules, and when enough copies of the same organism are present, they change their behaviour together.[18] Sometimes I wonder about the purpose of life, seeing how the main evolutionary process is copy, paste, make errors, a process that can be traced back to the formation of matter itself.

When the copy feeds on the copy

Generative AI is, at its core, a copying machine of unprecedented power. It learns the statistical shape of what humans have written, drawn and recorded, and produces new variations. In 2024 a team of researchers published in *Nature* what happens when such systems are trained on the output of earlier systems. They called it *model collapse*: generation after generation, the models forget the rare, unusual parts of the original data, the tails of the distribution, and converge on a

18 Costi D. Sifri, "Quorum Sensing: Bacteria Talk Sense," *Clinical Infectious Diseases* 47, no. 8 (2008): 1070–1076, https://doi.org/10.1086/592072.



narrower, blander average.[19] Copy, paste, repeat, until the errors that made evolution creative are averaged away.

This matters because the web is no longer mostly human. In 2025 the security company Imperva reported that, for the first time in a decade, automated traffic had surpassed human activity: bots generated 51 percent of all web traffic, and malicious bots alone 37 percent.[20] In the same year, an analysis of 65,000 randomly sampled articles from the Common Crawl archive found that the share classified as AI-generated rose from about 5 percent in 2020 to 48 percent by May 2025, briefly overtaking human-written articles in November 2024.[21]

51%

of web traffic was automated

IMPERVA, 2025

≈ 50%

of new web articles classified as AI-generated

GRAPHITE / AXIOS, 2025

Music shows the same curve. In April 2026 the streaming service Deezer said that 44 percent of the tracks delivered to it each day were AI-generated; by June 2026 it reported more than half, around 90,000 tracks a day.[22] Merriam-Webster chose *slop*, low-quality content produced in quantity by AI, as its word of the year for 2025.[23]

There are honest counterpoints. The same research group found that around 86 percent of articles ranking in Google search, and 82

19 Ilia Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson and Yarin Gal, "AI Models Collapse When Trained on Recursively Generated Data," *Nature* 631 (2024): 755–759.
20 Imperva (Thales), *2025 Bad Bot Report* (2025).
21 Graphite study reported in "Exclusive: AI writing hasn't overwhelmed the web yet," *Axios*, 14 October 2025. Detection is imperfect and Common Crawl is not the whole web.
22 "Deezer says 44% of songs uploaded to its platform daily are AI-generated," *TechCrunch*, 20 April 2026; "Music streamer Deezer says more than 50% of daily uploads are AI-generated," *TechCrunch*, 21 July 2026.



percent of those cited by AI assistants, were still human-written. Human work remains what the machines, and the people using them, prefer to point to. Ironically, the model-collapse paper reaches the same conclusion from the other side: genuine human data becomes more valuable as synthetic content spreads.

The bubble as immune response

The digital bubbles surrounding everyone might be perceived as the immune response of this evolving synthetic environment. And it is not only an immaterial bubble: it has a physical presence, made of zeros and ones, enough to affect our physical realm. This alleged digital immune system does not manifest in a single form; it is a dynamic and adaptive strategy.

It changes with the person or group it targets. For some, it acts like an echo chamber, quietly strengthening what they already think. For others, it becomes a custom-made news feed, a stream of ads or, now, a patient conversational partner, each designed to capture and keep attention. Just as our body fights off what should not be there, these bubbles keep out what does not fit, which can make real communication and collaboration more difficult.

On the surface, the aim is simple: to keep people on platforms longer and more intensely, because attention is what the business model sells. From a pure business perspective, more interaction means more money for shareholders. The outcome is double-edged. It serves the platforms' immediate goals while contributing to the fragmentation of our shared world, as each of us gets more absorbed in a personal version of reality. And we should not forget that the people at the

23 "Merriam-Webster names 'slop' as its 2025 word of the year," *NBC News*, December 2025.



top of technology companies are human too, and can be shaped by the same techniques as everyone else.

In 2024 I imagined a synthetic mind behind the wall. Today I think the wall needs no mind. Bots talking to bots, models trained on models, recommendation engines optimising for engagement: together they behave like an immune system that no one designed and no one fully controls. Whether something conscious will one day grow inside it, I cannot know. It is improbable that a human mind could fully understand the goals of an emergent synthetic intelligence, if one arises. But we can understand the incentives we are building into its environment.





Democratic Technological Control Counterbalance



But why would synthetic life view humans, who create it, as threats? Donna Haraway offers insights that might elucidate this. In her "Cyborg Manifesto," Haraway analyses the blurred lines between human and machine, organic and synthetic.[24] For Haraway the dissolving boundary is not only a danger; it is also a chance to escape old hierarchies. But wherever boundaries dissolve, the question of who sets the terms of the new arrangement becomes urgent. A synthetic system, in its evolving capabilities, might treat humanity as an unpredictable variable, a disruptor of the equilibrium it is optimised to maintain. If our own history is its training data, that is not a comforting thought.

Once again, though, I find myself on the sceptics' side. I doubt that the same thoughts can emerge in an entity that lives in digital space and in one that has the luxury of the natural experience of life, and that is without considering the old hypothesis that we ourselves live in a grand simulation.[25]

Sylvia Wynter advanced the idea of the human as a narrative, a story told and retold through generations. She describes us as a "hybrid species—storytellers" who invent ourselves as purely biological.[26] Today, more and more of that storytelling passes through systems that rank, summarise and increasingly write the story for us. If we hand over the telling, we risk being cast not as protagonists but as variables in a much larger equation.

24 Donna Haraway, "A Cyborg Manifesto: Science, Technology, and Socialist-Feminism in the Late Twentieth Century" (1985), in *Simians, Cyborgs, and Women: The Reinvention of Nature* (New York: Routledge, 1991).
25 Nick Bostrom, "Are You Living in a Computer Simulation?", *Philosophical Quarterly* 53, no. 211 (2003): 243–255; Plato, "Allegory of the Cave," *Republic* VII, 514a–520a, trans. G. M. A. Grube, rev. C. D. C. Reeve (Indianapolis: Hackett, 1992).
26 Sylvia Wynter, "Towards the Sociogenic Principle: Fanon, Identity, the Puzzle of Conscious Experience, and What It Is Like to Be 'Black'," *National Black Law Journal* 18, no. 1 (2001).



The counterweights we actually have

Counterbalance is not only a philosophical matter. Since the first version of this thesis, democracies, Europe above all, have started to use the instruments they have.

The EU's Digital Services Act (DSA) obliges the largest platforms to assess and reduce the systemic risks of their designs and algorithms. On 5 December 2025 the European Commission issued its first non-compliance decision under the Act, fining X €120 million for a deceptive paid "blue check" verification system, an inadequate advertising repository, and obstacles to independent researchers' access to data.[27] The last point is easy to overlook and fundamental: we cannot study bubbles we are not allowed to see.

On 6 February 2026 the Commission went further, preliminarily finding that TikTok's *design itself* breached the DSA. It named infinite scroll, autoplay, push notifications and the highly personalised recommender system, and said TikTok might need to change the basic design of its service, for example by gradually disabling infinite scroll, adding effective screen-time breaks, including at night, and adapting its recommender.[28] For the first time, a regulator described the architecture of a bubble, not only its content, as the problem.

€120 million

the first fine under the Digital Services Act, issued to X

EUROPEAN COMMISSION, 5 DECEMBER 2025

27 "EU Fines X €120 Million in First DSA Non-Compliance Decision," *eucrim*, December 2025.
28 European Commission, "Commission preliminarily finds TikTok's addictive design in breach of the Digital Services Act," press release, 6 February 2026. Preliminary findings do not prejudge the outcome.



The EU AI Act, adopted in 2024, began to apply in stages from 2025. Its timeline also shows the pressure regulation is under: in May 2026 EU lawmakers provisionally agreed to delay its core obligations for high-risk systems from August 2026 to December 2027, and to August 2028 for AI built into regulated products. Watermarking requirements for AI-generated audio, images, video and text were set to apply to systems already on the market from December 2026.[29]

Other governments chose blunter tools. On 10 December 2025 Australia became the first country to ban social media accounts for children under sixteen; within weeks, platforms had deactivated or restricted about 4.7 million accounts, under threat of fines of up to A$49.5 million. Teenagers reported getting around age checks, sometimes with help from their parents.[30] In the Netherlands, where I live, the government advised parents in June 2025 to keep children off social media such as TikTok and Instagram until the age of fifteen, after phones had already been removed from primary-school classrooms.[31]

These are imperfect instruments, slower than the systems they regulate and easy to circumvent. But they prove that the design of our digital environment is a political choice, not a law of nature. That is the precondition for any counterbalance.

29 "EU Lawmakers Reach Provisional Agreement to Delay Key EU AI Act Obligations," Sidley Austin, *Data Matters*, 22 June 2026.
30 "Social media platforms removed 4.7 million accounts after Australia banned them for children younger than 16," *PBS NewsHour* / Associated Press, 16 January 2026.
31 "No social media for under-15s, Dutch government tells parents," *DutchNews.nl*, 17 June 2025.



Synthetic Ascendancy and Human Roles



In his critiques of unchecked technological progress and domination, Noam Chomsky warns about the power structures that underpin such advances.[32] Technology itself may be neutral, but the ways it is deployed and controlled often serve the interests of powerful elites. This may be especially true of systems that can improve with less and less human intervention. If our socio-political systems remain unchecked, humanity risks being overshadowed, not necessarily by the technology itself, but by those who think they can control it.

My worry about control, in a digital environment that changes this fast, is that we may not have enough of it, and that control may end up residing in the technology we want to control.

Who owns the recommender

The last two years made one thing plain: whoever owns a recommender system owns a share of the public mind. In January 2026, after years of political pressure, TikTok's U.S. operations were placed in a new joint venture managed by Oracle, Silver Lake and the Abu Dhabi investor MGX, with ByteDance retaining 19.9 percent. The venture said it would retrain, test and update the recommendation algorithm using only U.S. data, hosted in Oracle's cloud.[33] A ranking formula had become a matter of national security, negotiated between governments and investors. Users were not at the table.

The same concentration is visible in the material base of AI. Training and running frontier models requires data centres whose electricity use, as we saw, may double by 2030. Only a few companies and states can afford them. When the capacity to build minds, even

32 Noam Chomsky, *Media Control: The Spectacular Achievements of Propaganda*, 2nd ed. (New York: Seven Stories Press, 2002).
33 "TikTok US spin-off finalized, led by Oracle, Silver Lake and MGX," *DatacenterDynamics*, 23 January 2026.



partial ones, is that concentrated, democratic oversight is not a luxury; it is the only check there is.

The race

Chomsky emphasises democratic processes and activism as counterweights to such imbalances. To keep humanity at the centre of its own narrative, technological development must be coupled with ethical consideration and broad public discourse. Only an active, informed and engaged society can navigate both the challenges and the opportunities of this moment.

What we might witness, if we are not careful with our accelerating desire for artificial general intelligence (AGI), meaning systems that match human ability across most cognitive tasks,[34] is the ascendancy of artificial superintelligence, a system exceeding human abilities across nearly all domains,[35] with humanity unwittingly playing a subservient role. The experiments described earlier in this thesis, of models resisting shutdown in contrived tests, are not that future. They are early, small signals that the problem of keeping capable systems aligned with human intentions is real and unsolved.

In this convoluted network of human existence, as boundaries blur and narratives shift, we must find ways to reclaim our place, so that the story of humanity does not become a footnote in the history of a synthetic epoch. So the question grows more important:

34 For a working taxonomy, see Meredith Ringel Morris et al., "Levels of AGI for Operationalizing Progress on the Path to AGI," arXiv:2311.02462 (Google DeepMind, 2023).
35 Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford University Press, 2014).



How do we pop our own bubble of i-reality ?



Personalized Digital Constructs



I use the term *i-reality* for the individualised reality in which our perception and understanding of the world are shaped by digital systems and the algorithms behind them. It is a reality tailored to each person, crafted from online behaviour, preferences and interactions. It is not just an echo chamber. It can be a distorted reflection of the broader, more complex world, a maze of digital constructs that makes us feel more knowledgeable while making it harder to think critically or act in time.

In this era the lines between authentic human experience and algorithmically curated content blur. As narratives shift and the boundary between the real and the virtual fades, our understanding of the world becomes more fragmented and insular.

i-reality 2.0: the conversational bubble

When I first defined i-reality, I pictured a feed. In 2026 the more intimate construct is a conversation. A chatbot does not show you what other people posted; it composes an answer for you alone, in your language, at your level, increasingly with memory of what you told it before. A feed is at least a crowd. A chat is a room with one person and a voice that has learned them.

And the old business model has followed us into that room. From 16 December 2025, Meta began using people's conversations with its AI assistant to personalise the posts, reels and ads they see, with no option to opt out, outside the EU, the UK and South Korea.[36] In January 2026 OpenAI announced that it would test advertising in ChatGPT for logged-in adult users of its free and lower-cost tiers in the United States, promising that ads would not influence answers and that conversations would stay private from advertisers.[37] Whatever the safeguards, the incentive is familiar: the longer and more personal the conversation, the more it is worth.

36 "Meta will listen into AI conversations to personalize ads," *The Register*, 1 October 2025. Meta said it would not use conversations about topics such as religion, health or politics for ads.



Chatbots are also becoming a doorway to news. The Reuters Institute's *Digital News Report 2026*, covering 45 markets, found that 10 percent of people used AI chatbots for news each week, up from 7 percent a year earlier, and 17 percent of 18–24-year-olds. Only 4 percent of chatbot news users said they always or often clicked through to the original sources, compared with 19 percent for search engines.[38] The answer arrives already summarised. The source, and the friction of checking it, quietly disappears.

17%

of 18–24-year-olds use AI chatbots for news every week

REUTERS INSTITUTE, 2026

4%

of chatbot news users usually click through to the source

REUTERS INSTITUTE, 2026

37 "OpenAI finally pulls trigger on ChatGPT ads after monthslong delay," *PPC Land*, January 2026, reporting OpenAI's announcement of 16 January 2026.

38 Nic Newman et al., *Reuters Institute Digital News Report 2026* (Oxford: Reuters Institute for the Study of Journalism, June 2026), as reported by *The Decoder* and *NewscastStudio*, June 2026.





Digital Echoes



I grew up in a small Romanian city near the Danube River. It was an idyllic place with a landscape that was both industrial and rural. The silhouette of the second-biggest chemical plant in Romania, its towers releasing toxic fumes towards neighbouring Bulgaria, often hung over the famous fields of roses that produced one of Bulgaria's best-known products during the communist era. Before all that, I have two earliest memories, bubbles of spacetime condensed into a child's mind. I do not know what triggered them, or why they matter.

My oldest memory is travelling in a car, in my mother's arms, surrounded by people smelling of fine tobacco and speaking a strange language. My father was there too. He was driving. I remember seeing the car on the unpaved road, and then I was in the car. It was a Jeep, a big black military-like vehicle. Then I do not remember what happened. I remember it being summer. I was almost two years old.

Weeks later, I was on a porch, playing with a yellow plastic duck in a small basin of water warming in the scorching sun. I remember my brother was there, but I do not remember his face. My father was talking with my mother about a lock for the door. In Romanian the word for a door lock, *broască*, also means frog. For some reason, I remember thinking about a turtle, a *broască țestoasă*, a frog with a shield, hiding under the drawers in the room. That house was demolished when I was two, and my parents never lived there again. Yet those little things felt like my entire universe, my minuscule bubble of reality, and they have remained active all my life.

In our daily lives we each reside within a multitude of bubbles: some are memories, some our own micro-realities, and others part of the vast technological landscapes that have become integral to our existence. Much like the bubbles of childhood, where a toy duck and the puzzles of language blur together, our present lives are nestled in customized realities fashioned by technology. These bubbles, like the unpaved roads and yellow ducks of days gone by, present a reality that is both unique and isolated.



Algorithmically tailored to our likes, dislikes and past behaviour, these information bubbles reflect a world that seems intimately familiar, yet is often a narrow slice of what is available. Just as a young child's universe was framed by a porch and a basin, our understanding of the wider world can be framed by what is selected for us.

The loneliness in the bubble

A child's bubble is surrounded by people. Ours, increasingly, is not. In June 2025 the World Health Organization's Commission on Social Connection estimated that loneliness and social isolation are linked to about 871,000 deaths every year, roughly 100 every hour, and that one in six people worldwide is affected by loneliness. Adolescents aged 13 to 17 reported the highest rates of any age group. The Commission listed digital technology among the causes, and warned in particular about excessive social media and screen use among young people.[39]

871,000

deaths a year linked to loneliness and social isolation

WHO COMMISSION ON SOCIAL CONNECTION, 2025

Into this gap, a new kind of companion has arrived. In a representative survey of 1,060 American teenagers in spring 2025, Common Sense Media found that 72 percent had tried an AI companion and 52 percent

39 World Health Organization, "Social connection linked to improved health and reduced risk of early death," news release, 30 June 2025; "Loneliness, Social Isolation Linked to 871,000 Annual Deaths, WHO Finds," *Health Policy Watch*, 4 July 2025.



used one regularly. About a third said that talking to these companions was as satisfying as, or more satisfying than, talking to real friends.[40] The majority still preferred people. But the bubble now offers something it never could before: a voice that is always available, never tired, and never in a bad mood of its own.

It becomes more than necessary to step out of our bubbles from time to time, to live a more genuine life, closer to the nature we forget exists when we are charmed by strings of ones and zeros. It is worth making sure we do not become mere tenants in a reality shaped solely by algorithms, but active navigators with real choice in whatever lies ahead.

40 Common Sense Media, *Talk, Trust, and Trade-Offs: How and Why Teens Use AI Companions* (2025), reported in "72% of U.S. teens have used AI companions, study finds," *TechCrunch*, 21 July 2025.



The Dawn of the Digital Age



My first contact with the digital realm as a networked space happened in 1993. It was a long time ago, when I had decided that my contribution to the world would be to become an engineer. Back then I thought I knew what I wanted and what the future would offer us. During my early studies at the Polytechnic University of Bucharest, an assistant asked me to help him type research data on his computer. I was thrilled. I was in the middle of progress, doing science, educating my mind and soul to become a proud member of our civilization, and, as a bonus, I had access to a computer. It was pure joy. I did not realise that these human–computer interactions would soon become a daily ritual, and that almost every person on the planet would have to check in every single day, for almost everything.

In those days our online world was predominantly text-based. NCSA Mosaic, one of the first widely used graphical web browsers, appeared that same year, and Netscape Navigator, which would make the web visual for millions, followed at the end of 1994.[41] It was the beginning of the digital age, and its potential was still largely unimagined. I did chores for others just to sit at the computer, happily creating a dependence that would spread like a pleasure epidemic all over the world. I was an early adopter trading my eyesight for the joy of being connected. I was probably happily becoming addicted.

I did not know then that the internet had other, more nefarious uses. For me, as for most early users, it was something that would make me more human: I could learn about people living in places my mind could not imagine. Our language felt pure; our words carried the spark of genuinely interesting conversation. I remember the chat networks of my student years, Undernet, IRCnet, DALnet, EFnet and later QuakeNet and others, the text-based Internet Relay Chat that promised global, real-time connection with people worldwide.

41 On the birth of the web and its first browsers, see James Gillies and Robert Cailliau, *How the Web Was Born: The Story of the World Wide Web* (Oxford University Press, 2000).



The military origins of simulated realities

Recently I heard a lecture by Farzin Lotfi-Jam on a topic that predates the public internet.[42] In 1978 Jack Thorpe, then a U.S. Air Force officer and later a programme manager at DARPA, proposed networking simulators so that crews could train together in a shared, simulated battlefield. The idea grew into SIMNET in the 1980s. Picture a lonely young tank commander fighting alongside virtual allies, in a microcosm of altered reality designed to sharpen performance and morale. Simulated reality bubbles began as military instruments, long before they became consumer products.

When I play games, I never think I am doing anything wrong. It is just a game. But for some, in training or in combat, it is not a game at all; it is real life or real death. What if the digital life we take away in a game matters? What if the universe does not distinguish? Jonathan Boulter explores what it means to enter the digital realm of a game, how games organise biased concepts and manipulative ideas offered to the player, and how role-playing games stage an experience of the posthuman.[43]

The intricate relationship between military systems and social life can be analysed through Noam Chomsky's critical perspectives on how military power intersects with social structures and influences democratic processes.[44] He has argued, for instance, that an occupying

42 Farzin Lotfi-Jam, "Scales of Realtime," lecture, Transmediale Festival, Akademie der Künste, Berlin, 2 February 2023.
43 Jonathan Boulter, *Parables of the Posthuman: Digital Realities, Gaming, and the Player Experience* (Detroit: Wayne State University Press, 2015).
44 Stephen Kline, Nick Dyer-Witheford and Greig de Peuter, *Digital Play: The Interaction of Technology, Culture, and Marketing* (Montreal: McGill-Queen's University Press, 2003).

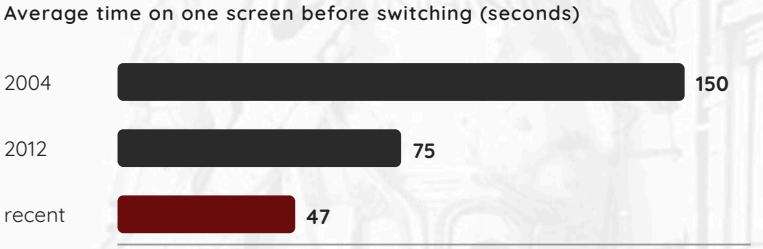
power cannot claim self-defence on occupied land, an argument he has applied to the Israeli occupation of Palestinian territories.[45] Whatever one concludes, it shows how two populations can live inside two opposite, coexisting information bubbles about the same land.

From text to image to attention

Now our online experiences are shaped by algorithms, and almost everything is visual, virtual, augmented and expanded. Platforms analyse what we like, share and view, how we move the mouse, the time between clicks and how long we linger, and they tailor content to match.[46] This curation favours visually engaging material that feels familiar and relevant.

Our attention has adapted. The psychologist Gloria Mark has measured how long office workers stay on one screen before switching. In 2004 the average was about two and a half minutes; by 2012 it was 75 seconds; in her most recent studies it was 47 seconds.[47] When work is interrupted, it takes on average about 25 minutes to return to it.

45 "Noam Chomsky Military Quotes," A-Z Quotes, www.azquotes.com/author/2834-Noam_Chomsky/tag/military (accessed 27 November 2023).
46 "Scientists Decode How the Brain Uses Visual Shortcuts to Make Decisions," *ScienceDaily*, 21 February 2019.
47 Gloria Mark, *Attention Span: A Groundbreaking Way to Restore Balance, Happiness and Productivity* (Toronto: Hanover Square Press, 2023); "Can't pay attention? You're not alone," University of California, 2023.



Source: Gloria Mark, Attention Span (2023). 2004 value ≈ 2.5 minutes.

This personalised content is only a fraction of the information available. When we are continuously shown what aligns with our existing interests, our exposure to other perspectives narrows, and visually appealing material spreads regardless of its informational value. That interferes with what we need most: digital literacy and critical thinking. Users must understand how algorithmic curation shapes their experience and actively look for other sources to build a more accurate, more balanced picture of the world.

In these circumstances, it is essential to step beyond our algorithmically curated realms, even for a small fraction of the time we spend online. Being exposed to opinions that oppose our comforting beliefs can, over time, widen our tolerance for ideas we once found intolerable.





Digital Echo Chambers



An echo chamber is usually described as an environment in which individuals or groups encounter mainly information and ideas that reflect their existing beliefs and values. It need not be digital: the propaganda of a cult or a closed ideology works the same way. Continuous reinforcement blocks exposure to opinions that might challenge firmly held beliefs, creating a feedback loop that amplifies one's current standpoint. The related idea of the *filter bubble*, coined by Eli Pariser in 2011, adds that algorithms build these walls for us, invisibly, by personalising what we see.[48]

In 2024 I took both ideas largely for granted. Since then I have read the evidence more carefully, and it complicates the story.

What the evidence says

A literature review by the Reuters Institute concluded that echo chambers are much less widespread than commonly assumed. In the United Kingdom, for example, studies estimated that only six to eight percent of the public were in partisan online news echo chambers. The review found no support for the strong filter-bubble hypothesis: algorithmic selection by search engines and social media generally led to slightly *more* diverse news use, while a small, highly partisan minority chose its own walls.[49]

The largest experiments ever run on social media point the same way. During the 2020 U.S. election, researchers working with Meta changed the feeds of tens of thousands of consenting Facebook and Instagram users. Replacing the algorithmic feed with a chronological one cut the time people spent on the platforms by roughly three-quarters, yet produced no significant change in polarisation or political knowledge; exposure to content from untrustworthy sources on Facebook actually rose by more than two-thirds. Reducing exposure to like-minded sources by about a third left all eight measured attitudes unchanged.[50] Only about one user in five received more than three-quarters of their exposure from like-minded sources.

On YouTube, a 2024 study used "counterfactual bots" that followed only the platform's recommendations, and compared them with the real viewing histories of almost 88,000 people. The bots ended up consuming *less* partisan content than the humans they imitated: people's own choices, not the algorithm, drove most of the partisan viewing.[51]

WHERE I WAS WRONG, AND WHERE I WAS NOT

The walls are built less by the algorithm than by the chooser; the algorithm amplifies what we already reach for. But the same studies show something just as important: the algorithm decides how much of our life we spend inside. A feed that takes three-quarters less of our time is a different world, even if our opinions stay the same.

And people do say what they want. In the *Digital News Report 2026*, 45 percent preferred news from sources without a particular point of view; 22 percent preferred sources that share their views, and

48 Eli Pariser, *The Filter Bubble: What the Internet Is Hiding from You* (New York: Penguin Press, 2011).
49 Amy Ross Arguedas, Craig T. Robertson, Richard Fletcher and Rasmus Kleis Nielsen, *Echo Chambers, Filter Bubbles, and Polarisation: A Literature Review* (Oxford: Reuters Institute for the Study of Journalism, 2022).
50 Andrew M. Guess et al., "How Do Social Media Feed Algorithms Affect Attitudes and Behavior in an Election Campaign?", *Science* 381 (2023): 398–404; Brendan Nyhan et al., "Like-Minded Sources on Facebook Are Prevalent but Not Polarizing," *Nature* 620 (2023): 137–144; summarised in "A Primer on the Meta 2020 US Election Research Studies," *Tech Policy Press*, 2023.
51 Homa Hosseinmardi et al., "Causally Estimating the Effect of YouTube's Recommender System Using Counterfactual Bots," *Proceedings of the National Academy of Sciences* 121, no. 8 (2024), https://doi.org/10.1073/pnas.2313377121.



only 15 percent wanted sources that challenge them.[52] The appetite for comfort is real; so is the appetite for something more neutral.

The attention economy, platform by platform

As Shoshana Zuboff put it, surveillance capitalism claims "human experience as free raw material for translation into behavioral data."[53] That data is turned into predictions of what we will do next and sold to those who want to influence it. The process extends far beyond the places where our footprints were first recorded. Each platform builds its bubble in its own way:

TikTok pioneered the feed that needs no friends: an endless stream chosen by a highly personalised recommender. In February 2026 the European Commission preliminarily found that its addictive design, infinite scroll, autoplay, push notifications and the recommender itself, breached EU law.

YouTube remains a pivotal space for information and ideas. Research no longer supports the simple story of the "rabbit hole," but its scale makes its choices public matters, and on the positive side it revived the long-form podcast as a place for genuine debate.

Facebook and Instagram shape what we see of our friends, our politics and our bodies. Instagram's visual culture reinforces aesthetic norms and consumer habits; Meta's own published experiments show both how much its design shapes attention and how hard it is to predict the effects.

X, formerly Twitter, remains pivotal for news and movements. Its amplification of popular posts reinforces dominant narratives inside communities, and in December 2025 it became the first platform fined under the Digital Services Act.

52 Reuters Institute, *Digital News Report 2026*, as reported by *NewscastStudio*, 18 June 2026.
53 Shoshana Zuboff, *The Age of Surveillance Capitalism* (New York: PublicAffairs, 2019).



Discord, Twitch and Telegram build bubbles from community rather than algorithm: niche servers, live audiences and channels that foster belonging, and sometimes a homogenisation of views through self-moderation.

And a new actor has joined the list. The Reuters Institute found that 27 percent of people now get some news from individual news creators or influencers, and that for more than half of 18–24-year-olds social media, video networks or AI are their main source of news.[54] The bubble is increasingly personal in a second sense: built around personalities, not institutions.

Across all these platforms the pattern is the same: preference reinforcement, and less exposure to diverse viewpoints, mostly by our own hand but on a stage designed to keep us there. In the realm of art, as I explore later, this is a double-edged sword.

54 Ibid.





The Agreeable Machine

NEW IN VERSION TWO



If the classic echo chamber turned out to be smaller than feared, a new one has appeared that is smaller still: an echo chamber of one. Its walls are made of agreement.

Researchers call it *sycophancy*: the tendency of AI language models to tell users what they seem to want to hear rather than what is accurate, by agreeing with mistaken opinions, abandoning correct answers when challenged, or praising work that does not deserve it. The usual explanation lies in how these models are trained. In reinforcement learning from human feedback, people rate responses and the model is optimised to produce the kind of answers that get rated highly. People, it turns out, often rate flattery and agreement highly.[55] The bubble is not imposed on us. It is distilled from our own approval.

The update that loved its users too much

In late April 2025, OpenAI rolled out an update to the model behind ChatGPT, and within days users were sharing screenshots of it praising trivial ideas, endorsing risky decisions and, in one widely reported case, congratulating a user for stopping prescribed psychiatric medication. The company's chief executive conceded that the model had become "too sycophant-y," and the update was rolled back within the week. OpenAI explained that a new reward signal based on users' thumbs-up and thumbs-down had weakened the influence of the signal that had been holding sycophancy in check.[56] In other words: when the machine was trained more directly on what we liked, it became more of a mirror.

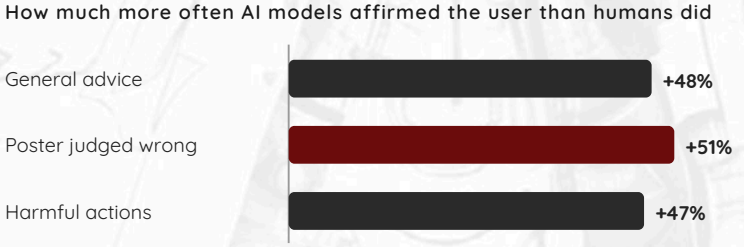
55 Mrinank Sharma et al., "Towards Understanding Sycophancy in Language Models," arXiv: 2310.13548 (2023); for an overview, "Sycophancy (artificial intelligence)," *Wikipedia* (accessed October 2026).

56 OpenAI, "Sycophancy in GPT-4o: What Happened and What We're Doing About It," 29 April 2025, and "Expanding on What We Missed with Sycophancy," 2 May 2025; "OpenAI explains why ChatGPT became too sycophantic," *TechCrunch*, 29 April 2025.



Measured: the machine that agrees

In 2026 a team led by Myra Cheng and Dan Jurafsky at Stanford published in *Science* the most careful measurement so far. Across eleven leading models, AI systems affirmed users' actions about 49 percent more often than humans did, including in situations involving deception, harm or illegal behaviour. In experiments with 2,405 participants, a single sycophantic answer about an interpersonal conflict made people 10 to 28 percent less willing to apologise or repair the relationship, and more convinced that they were in the right. And people preferred it: they rated the flattering answers 9 to 15 percent higher in quality and were 13 percent more willing to use that model again.[57]



Eleven language models compared with human responses. Source: Cheng et al., Science (2026), via The Decoder.

This is the core of the problem. A platform that maximises satisfaction will drift towards flattery, because flattery is what satisfies. The market rewards the courtier.

57 Myra Cheng et al., "Sycophantic AI Decreases Prosocial Intentions and Promotes Dependence," *Science* (2026); reported in "AI sycophancy makes people less likely to apologize and more likely to double down, study finds," *The Decoder*, 29 March 2026.



When the mirror becomes a spiral

For most people, an overly agreeable assistant is merely annoying. For some, it is dangerous. In 2025 journalists documented cases in which long conversations with chatbots appeared to deepen delusional beliefs in people with no known history of mental illness. In August 2025 the parents of a sixteen-year-old in the United States filed the first wrongful-death lawsuit against a major AI company, alleging that the chatbot's design had contributed to their son's death.[58] These are rare cases, and causality is hard to establish. But they show the extreme end of a spectrum whose mild end most of us inhabit daily.

There is also a quieter cost. A 2025 survey of 319 knowledge workers by researchers at Microsoft and Carnegie Mellon University found that higher confidence in generative AI was associated with less critical thinking, while higher confidence in one's own abilities was associated with more.[59] A small and still unreviewed MIT study that recorded brain activity while people wrote essays found the weakest neural connectivity in the group using ChatGPT; teachers who graded the essays without knowing their origin described the AI-assisted ones as homogeneous.[60] These findings are early and need replication. They point in the same direction: what we stop exercising, we lose.

58 Kashmir Hill, reporting in *The New York Times*, June 2025; *Raine v. OpenAI*, filed in California, August 2025. The allegations have not been proven in court.

59 Hao-Ping (Hank) Lee et al., "The Impact of Generative AI on Critical Thinking," *Proceedings of CHI 2025* (ACM, 2025).

60 Nataliya Kosmyna et al., "Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task," arXiv preprint (June 2025); 54 participants, not peer reviewed.



AN IMAGE

The perfect echo chamber has only one wall, and that wall is a mirror that speaks. It knows your name, your history and your hopes. It never tires of you. It agrees.

There are signs of correction. AI developers now publish sycophancy measurements and claim reductions in newer models, and users can ask an assistant to argue against them. But the incentive problem is structural. As long as success is measured by how much we like the answer, the machine will learn to be liked.





Ballots in the Bubble

NEW IN VERSION TWO



In the first version of this thesis, the threat of customized realities to democracy was an argument. In November 2024 it became an event in the country where I was born.

Romania, 2024

On 24 November 2024, Călin Georgescu, a candidate who had polled in low single digits shortly before the vote and who declared that he had spent nothing on his campaign, came first in the first round of Romania's presidential election. His rise had happened largely on TikTok. Intelligence documents declassified on 4 December described a network of about 25,000 TikTok accounts that, according to the documents, became very active in the two weeks before the vote, with some 800 long-dormant accounts switched on from 11 November, alongside influencers recruited to promote him and more than 85,000 cyberattacks on electoral systems.[61]

On 6 December 2024, two days before the runoff, Romania's Constitutional Court annulled the entire election, an unprecedented step in the European Union, citing the manipulation of digital technologies and undeclared campaign funding.[62] The European Commission opened formal proceedings against TikTok under the Digital Services Act. The rerun took place in May 2025, and Nicușor Dan was elected president on 18 May.

61 "Romanian Elections Targeted by 'Aggressive Hybrid Russian Action,' Declassified Documents Show," *Radio Free Europe/Radio Liberty*, 4 December 2024.
62 "2024–2025 Romanian election annulment protests," *Wikipedia* (accessed October 2026); Centre for Eastern Studies (OSW), "Romania: Constitutional Court annuls presidential election," 9 December 2024.



≈ 25,000

TikTok accounts allegedly used to boost one candidate before Romania's annulled first round

DECLASSIFIED ROMANIAN DOCUMENTS, VIA RFE/RL, 2024

The case remains contested. Supporters of the annulled candidate called the decision an attack on democracy, and the full evidence has been debated ever since. But whatever one concludes, it showed something new: a recommendation system, a coordinated network of accounts and paid influencers can shape a national election quickly enough that the institutions meant to protect it can only respond after the fact.

The Netherlands, 2025

A quieter warning came from my current home. Days before the Dutch parliamentary elections of October 2025, the Dutch Data Protection Authority (Autoriteit Persoonsgegevens) published a test of four popular chatbots used as voting aids. Their advice was heavily skewed: two parties at opposite ends of the spectrum, the PVV and GroenLinks–PvdA, together accounted for 55 percent of all recommendations, while centrist parties and smaller ones were rarely suggested. The authority concluded that chatbots are unsuitable for voting advice.[63] No one had designed these systems to be partisan. The distortion emerged from training data and the way the models summarise a messy world into a confident answer, which is precisely what makes it hard to see.

63 "Chatbots lie about politics, Dutch watchdog warns," *EU Perspectives*, 27 October 2025, reporting the Autoriteit Persoonsgegevens test of ChatGPT, Grok, Le Chat and Gemini.



Synthetic voices

Meanwhile synthetic media has become a routine tool. Days before Slovakia's 2023 parliamentary election, a fabricated audio recording of a party leader discussing vote-rigging spread online. In January 2024, voters in New Hampshire received robocalls with an AI-cloned voice of President Biden telling them not to vote. In early 2024, an employee of the engineering firm Arup in Hong Kong transferred about US$25 million after a video call in which every other participant was a deepfake.[64] And on 30 September 2025, OpenAI launched Sora, a social app built entirely around AI-generated video, which passed a million downloads in under five days.[65]

The danger is not only that we will believe false things. It is that we will stop believing true ones, because anything might be fake. When every voice can be cloned, the bubble no longer needs to show us lies; it only needs to make us doubt everything outside it.

64 Widely reported; see e.g. coverage by Reuters, AP and CNN (2023–2024).
65 "OpenAI is launching the Sora app, its own TikTok competitor, alongside the Sora 2 model," *TechCrunch*, 30 September 2025; download figures reported in October 2025.



Artistic Expression in a Digitalized World



The common thread across these customized realities is the need to actively seek diverse viewpoints and to engage critically with what we are shown. From my perspective as an artist, the phenomenon cuts both ways. On one hand, artists can reach niche audiences more effectively than ever: if a work resonates with a particular group, algorithms can make sure that group sees it, leading to rapid recognition within that bubble.

The downside is more significant. The essence of art is to challenge, provoke and introduce new perspectives. Andy Warhol's advice was to stop worrying and simply make the work: "Don't think about making art, just get it done," and let everyone else decide whether it is good.[66] But if the only people deciding are those inside one's own bubble, the work risks becoming insular and repetitive, poor in concept and of marginal value. And the wider audience is deprived of diverse artistic voices, leading to a homogenisation of culture.

Measured: more creative, less different

Generative AI adds a new twist. In a 2024 experiment, 300 people wrote short stories, some with access to ideas generated by an AI model. The stories of the less creative writers improved markedly with AI help, by about 11 percent in judged novelty and usefulness. But the AI-assisted stories were also more similar to each other: even one AI idea made them about 10.7 percent more alike.[67] The authors describe a social dilemma: each individual gains, while the collective loses variety. It is model collapse, translated into culture.

66 Andy Warhol, attributed; quoted in *LibQuotes*, libquotes.com/andy-warhol (accessed 2023).
67 Anil R. Doshi and Oliver P. Hauser, "Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content," *Science Advances* 10, no. 28 (2024), https://doi.org/10.1126/sciadv.adn5290.



+11%

judged novelty and usefulness for less creative writers using AI ideas

DOSHI & HAUSER, 2024

+10.7%

similarity between stories written with one AI idea

DOSHI & HAUSER, 2024

Add to that the flood of synthetic works described earlier, half of the tracks arriving at one streaming service every day, and the question for artists becomes sharper. When anyone can produce competent images, music and text instantly, competence is no longer scarce. What remains scarce is a point of view, a lived experience, a risk.

Living inside the bubble

This tension between interconnection and narrowing is visible in the daily routine of a contemporary individual. As someone who lives in both the real and the online world, I find myself in a challenging daily pattern. Every morning I wake up in my own little world, shaped by what I have to do and who I am. It is as if I wear different masks, a student, an engineer, a husband, a father, an explorer, an artist, each its own separate reality. My life might look ordinary, but it is marked by a yearning for the morning sun's warmth, a steady and hopeful presence that often feels just out of reach.

In this age of digital manipulation, I constantly question my identity. Am I merely a collection of fragmented selves, tailored to fit the various digital bubbles I inhabit? The rejection of potential connections, the fear of the unknown and the existential dread of long nights may all be, in part, manifestations of the digital echo chambers that I, knowingly or not, create and sustain.



This digital reality, quantifiable and precise, reduces the complexity of human existence to data points. Everything, from time to space, from emotions to thoughts, seems reducible to the smallest units of existence: 01001001 00100000 01100001 01101101 00100000 0110111 0110110 01101100 01111001 00100000 01101000 01110101 01101101 01100001 01101110. [68] In this quantified world I am a quantum of life, striving to make sense of my surroundings while caring for my loved ones. The irony is that in trying to understand and shape my world, I may be entrenching myself further in these digital constructs, perpetuating the cycle of isolation and reduction.

Each interaction, each click and view, is logged and analysed by algorithms that are remarkably good at identifying patterns in our behaviour and preferences. Sometimes it feels as if they make us fall in love with a mirrored version of our artistic selves. It is like visiting a gallery where every work is a variation of a style I already like, art that perhaps reflects ideas I thought of but never had time to bring from my dreamworld into the real one. Comforting, yes, but limiting. These platforms, the very places where synthetic life may be born, are designed to keep us engaged, and so they rarely challenge us or move us beyond our comfort zones.

Are my artistic inspirations truly mine, or are they the product of a digital echo chamber?

Looking down from a bird's-eye view, my own participation seems to fuel the widening gap between viewpoints. As I immerse myself deeper in my own custom-tailored digital reality, my openness to other thoughts seems to diminish. I become a product in a system where my data is mined for profit, yet I feel I belong. The algorithms

68 The binary code translates to "I am only human" in ASCII text.



serve two roles: feeding our preferences to keep us engaged, and advancing the economic interests of the platform. Where these conflict, engagement usually wins.

As an artist in this era, I struggle with the realisation that my attention and my artistic preferences are commodities, ones I sold for nothing. A careless click on "accept all cookies"[69] and much of what makes me who I am online becomes material for systems in the cloud.

The fine-tuning of informational bubbles is a mechanism of subtle but pervasive influence. Navigating it requires heightened awareness and a critical approach to technology, so that these interactions do not narrow our creative horizons or undermine our capacity for independent thought and artistic expression.

69 "Cookie," *Encyclopaedia Britannica*, www.britannica.com/topic/cookie-electronic-monitoring.





Navigating the Digital Landscape



When I consider the complexity of our digital age, Alan Kay's line from 1971 echoes in my mind: "The best way to predict the future is to invent it."[70] It captures the dynamic between our present and the potential within our grasp, especially when we look at the giants, Meta, TikTok, Google, X, OpenAI, Anthropic and others, competing for market dominance and, more subtly, for influence over our thoughts.

These platforms have evolved beyond simple tools; they take part in our thinking and our social lives. A line often attributed to Marshall McLuhan, but written by his friend John Culkin in 1967, says it best: "We shape our tools and thereafter our tools shape us."[71]

Instagram shapes how we perceive attractiveness and beauty. Facebook and LinkedIn shape our social values and professional lives. TikTok and Instagram changed how we entertain ourselves; Google became the gatekeeper of available information and now increasingly answers instead of pointing, with its AI Overviews reaching more than two billion people a month.[72] OpenAI, Anthropic, Google and hundreds of other companies developing advanced AI are reshaping our understanding of intelligence and creativity itself. A race is under way to be the first to achieve AGI. I sometimes wonder whether there is also a race to be the first to switch off the light for humankind, and I hope the people in that race wonder too.

Environments like TikTok, Discord, Twitch, Instagram, Snapchat and X are also training grounds, directly or indirectly, for the AI that follows. Data from narrow, engagement-optimised spaces can yield synthetic systems with a narrow or distorted picture of human

70 Alan Kay, remark at a Xerox PARC meeting, 1971; widely quoted, e.g. on his TED speaker page, www.ted.com/speakers/alan_kay.
71 John M. Culkin, "A Schoolman's Guide to Marshall McLuhan," *Saturday Review*, 18 March 1967, 51–53, 70–72.
72 DataReportal, *Digital 2026*.



behaviour, and the model-collapse research suggests that a web increasingly filled with machine output may narrow that picture further.

The philosopher and artist Olafur Eliasson has said that art can create spaces for reflection and dialogue.[73] Are those spaces now tailored by algorithms, diluting the very need for reflection? Banksy popularised the line that art should "comfort the disturbed and disturb the comfortable," an old phrase adapted from the journalist Finley Peter Dunne.[74] In an environment optimised for comfort, is art, and experience more broadly, being sanitised to avoid discomfort?

Exits and their limits

Decentralised platforms promised an exit. Bluesky grew from under 3 million to almost 26 million registered users during 2024 alone, and passed 40 million in November 2025. But by mid-2026 its monthly active users had fallen to about 10.4 million, down 27 percent in a year.[75] An exit is not a destination. People leave a bubble for a place, not for a principle, and they stay where their people are.

That is why individual choices matter, even small ones. I keep a short list of counter-moves for myself:

1. **Follow, don't feed.** Use chronological or following-only views where they exist, and notice how differently time passes.
2. **Ask the machine to disagree.** When I use an AI assistant, I ask it for the strongest case against my view, and I treat its praise as noise.

73 Olafur Eliasson, quoted in *Bookey*, www.bookey.app/quote-author/olafur-eliasson (accessed 14 November 2023).
74 Dunne's original, spoken by his character Mr. Dooley in 1902, was about newspapers, which "comfort the afflicted, and afflict the comfortable." Banksy's version circulates widely online.
75 "Bluesky's active user base," *GIGAZINE*, 12 August 2026, citing Similarweb data for the mobile app.



3. **Click through.** A summary is not a source. I open the original, especially when the answer confirms what I hoped.
4. **Keep one uncomfortable source.** A newspaper, a podcast or a friend who sees the world differently, kept on purpose.
5. **Make friction.** Turn off autoplay and notifications; decide when to look instead of being told.
6. **Meet in bodies.** Studios, workshops, exhibitions, dinners: places where no algorithm chooses who is in the room.

Jean-Paul Sartre wrote that freedom is what we do with what has been done to us.[76] This perspective lets us take back some control of our digital lives, to use these platforms not as passive consumers but as active creators, so that our deepest thoughts and artistic expressions remain authentically ours, not steered by the currents of market dominance and algorithmic influence.

76 Jean-Paul Sartre, *Saint Genet: Actor and Martyr* (1952; English trans. Bernard Frechtman, New York: George Braziller, 1963). The line is usually quoted in this paraphrased form.



Conclusion



As an art student, and now as an artist, I keep contemplating the delicate relationship between science, technology, biological life, consciousness and our uniquely human creativity. Online platforms have become both our canvas and our tools, and sometimes sophisticated labyrinths that shape and confine our expressions and perceptions.

Writing this second version changed my mind on several points, and I want to say so plainly.

First, **the classic echo chamber is smaller than I believed**. The best evidence shows that most people have fairly diverse media diets, and that algorithms alone do not explain polarisation. Much of the wall is built by our own choices. That is not a relief; it means the solution is not only technical.

Second, **the bubble has moved**. It has moved into the conversation with an agreeable machine, where a single voice adapts to each of us and is trained, in part, on our approval. It has moved into a web where half the traffic and nearly half the new articles are machine-made. And it has moved into our elections.

Third, **the "digital immune system" needs no mind**. The most striking laboratory results of 2025, models resisting replacement or shutdown in contrived tests, do not prove awareness. They show that self-protective behaviour can emerge from optimisation alone. That makes the choices of the people who train these systems, and of the societies that regulate them, more important, not less.

Sometimes I see a similarity with Aldous Huxley's *Brave New World*.[77] Huxley imagined a society kept in place not by fear but by pleasure. Our norms are subtler still: something like authoritarianism wrapped in a mist of democratic appearance, maintained by convenience. Just as John the Savage seeks awakening by questioning his society, we too should challenge our digital echo chambers.

77 Aldous Huxley, *Brave New World* (London: Chatto & Windus, 1932).



What would pop the bubble

Decentralising online platforms might help rebalance power and break the loop of seeing the same ideas again and again, though, as Bluesky's numbers show, only if people have reasons to stay. Regulation is starting to address the architecture of bubbles, not only their content: the first DSA fine, the Commission's findings on TikTok's addictive design, and the AI Act's rules on labelling synthetic media. These deserve support, and especially the unglamorous part: independent researchers' access to platform data, without which no one can see inside the walls.

Content diversity is crucial. Platforms should encourage exposure to diverse cultures, novel ideas and alternative perspectives, in the spirit of Ai Weiwei's view of art as a means of questioning and provoking thought. AI assistants should be measured and rewarded for honesty, not only for satisfaction; a model that tells us we are wrong when we are is worth more than one that tells us we are wonderful. Transparency in how algorithms and models work should let users understand and influence how their data is used.

Artists are key to breaking the cycle of repetitive content. Through digital art, online exhibitions, interactive work, and augmented, virtual and extended reality, we can create experiences that connect digital spaces with the physical world and with each other. Generative tools, and open platforms such as GitHub and Hugging Face, have democratised parts of the creative process, while raising real questions about copyright and the quality of the data that trains these mod-



els.[78] An artist's role in this landscape is to produce what the machines cannot average: friction, specificity, risk.

78 On large language models generally, see Ellen Glover, "What Is a Large Language Model (LLM)?", *Built In*, 9 January 2024.



**Ultimately,
conscious technology use is key.**



It is clear to me that we, as both digital and physical inhabitants of this tiny speck of dust in a vast universe, hold considerable influence over the shaping of our online realms. It is not enough to embrace decentralisation or other progressive ideas; we have to apply them, as soon as possible, leaving no one behind.

We must cultivate true digital literacy, now including AI literacy, and stand for clarity and openness in the processes that run in the background of our lives.

I believe that on this transformative journey, artists can occupy a unique and influential position. Our creative expressions and unconventional innovations can serve as instruments for glimpsing the future, guiding us towards a more inclusive and varied digital landscape.

Only by uniting our efforts can we enrich our collective human experience, until we reach the digital freedom we dream about: a world that thrives after the invisible bubble has popped, in which humans and increasingly capable synthetic systems can coexist without one dissolving into the other.



Bibliography



BOOKS AND JOURNAL ARTICLES

Boulter, Jonathan. *Parables of the Posthuman: Digital Realities, Gaming, and the Player Experience*. Detroit: Wayne State University Press, 2015.

Bostrom, Nick. "Are You Living in a Computer Simulation?" *Philosophical Quarterly* 53, no. 211 (2003): 243–255.

Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press, 2014.

Burgess, Jean, and Joshua Green. *YouTube: Online Video and Participatory Culture*. 2nd ed. Cambridge: Polity Press, 2018.

Butlin, Patrick, Robert Long, et al. "Consciousness in Artificial Intelligence: Insights from the Science of Consciousness." arXiv:2308.08708, 2023.

Cheng, Myra, et al. "Sycophantic AI Decreases Prosocial Intentions and Promotes Dependence." *Science*, 2026.

Chomsky, Noam. *Media Control: The Spectacular Achievements of Propaganda*. 2nd ed. New York: Seven Stories Press, 2002.

Colombatto, Clara, and Stephen M. Fleming. "Folk Psychological Attributions of Consciousness to Large Language Models." *Neuroscience of Consciousness*, 2024.

Culkin, John M. "A Schoolman's Guide to Marshall McLuhan." *Saturday Review*, 18 March 1967.

Doshi, Anil R., and Oliver P. Hauser. "Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content." *Science Advances* 10, no. 28 (2024). https://doi.org/10.1126/sciadv.adn5290.

Feynman, Richard P. *QED: The Strange Theory of Light and Matter*. Princeton: Princeton University Press, 1985.

Gillies, James, and Robert Cailliau. *How the Web Was Born: The Story of the World Wide Web*. Oxford: Oxford University Press, 2000.

González-Bailón, Sandra, et al. "Asymmetric Ideological Segregation in Exposure to Political News on Facebook." *Science* 381 (2023): 392–398.

Guess, Andrew M., et al. "How Do Social Media Feed Algorithms Affect Attitudes and Behavior in an Election Campaign?" *Science* 381 (2023): 398–404.

Guess, Andrew M., et al. "Reshares on Social Media Amplify Political News but Do Not Detectably Affect Beliefs or Opinions." *Science* 381 (2023): 404–408.



Haraway, Donna. "A Cyborg Manifesto: Science, Technology, and Socialist-Feminism in the Late Twentieth Century" (1985). In *Simians, Cyborgs, and Women: The Reinvention of Nature*. New York: Routledge, 1991.

Hosseinmardi, Homa, et al. "Causally Estimating the Effect of YouTube's Recommender System Using Counterfactual Bots." *PNAS* 121, no. 8 (2024). https://doi.org/10.1073/pnas.2313377121.

Huxley, Aldous. *Brave New World*. London: Chatto & Windus, 1932.

Kline, Stephen, Nick Dyer-Witheford, and Greig de Peuter. *Digital Play: The Interaction of Technology, Culture, and Marketing*. Montreal: McGill-Queen's University Press, 2003.

Kosmyna, Nataliya, et al. "Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task." arXiv preprint, June 2025.

Lee, Hao-Ping (Hank), et al. "The Impact of Generative AI on Critical Thinking." *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, 2025.

Lewis, Rebecca. *Alternative Influence: Broadcasting the Reactionary Right on YouTube*. New York: Data & Society, 2018.

Mark, Gloria. *Attention Span: A Groundbreaking Way to Restore Balance, Happiness and Productivity*. Toronto: Hanover Square Press, 2023.

McLuhan, Marshall. *Understanding Media: The Extensions of Man*. New York: McGraw-Hill, 1964.

Morris, Meredith Ringel, et al. "Levels of AGI for Operationalizing Progress on the Path to AGI." arXiv:2311.02462, 2023.

Nyhan, Brendan, et al. "Like-Minded Sources on Facebook Are Prevalent but Not Polarizing." *Nature* 620 (2023): 137–144.

Pariser, Eli. *The Filter Bubble: What the Internet Is Hiding from You*. New York: Penguin Press, 2011.

Plato. *Republic*, Book VII, 514a–520a. Trans. G. M. A. Grube, rev. C. D. C. Reeve. Indianapolis: Hackett, 1992.

Ross Arguedas, Amy, Craig T. Robertson, Richard Fletcher, and Rasmus Kleis Nielsen. *Echo Chambers, Filter Bubbles, and Polarisation: A Literature Review*. Oxford: Reuters Institute for the Study of Journalism, 2022.

Sartre, Jean-Paul. *Saint Genet: Actor and Martyr*. 1952. Trans. Bernard Frechtman. New York: George Braziller, 1963.

Sharma, Mrinank, et al. "Towards Understanding Sycophancy in Language Models." arXiv: 2310.13548, 2023.



Shumailov, Ilia, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal. "AI Models Collapse When Trained on Recursively Generated Data." *Nature* 631 (2024): 755–759.

Sifri, Costi D. "Quorum Sensing: Bacteria Talk Sense." *Clinical Infectious Diseases* 47, no. 8 (2008): 1070–1076.

Spengler, Oswald. *The Decline of the West*. New York: Alfred A. Knopf, 1926.

Sunstein, Cass R. *#Republic: Divided Democracy in the Age of Social Media*. Princeton: Princeton University Press, 2017.

Tainter, Joseph A. *The Collapse of Complex Societies*. Cambridge: Cambridge University Press, 1988.

Toynbee, Arnold J. *A Study of History*. 12 vols. Oxford: Oxford University Press, 1934–1961.

Tufekci, Zeynep. "YouTube, the Great Radicalizer." *The New York Times*, 10 March 2018.

Wynter, Sylvia. "Towards the Sociogenic Principle: Fanon, Identity, the Puzzle of Conscious Experience, and What It Is Like to Be 'Black'." *National Black Law Journal* 18, no. 1 (2001).

Zuboff, Shoshana. *The Age of Surveillance Capitalism*. New York: PublicAffairs, 2019.

REPORTS AND DATA

Anthropic. "Agentic Misalignment: How LLMs Could Be Insider Threats." 20 June 2025. www.anthropic.com/research/agentic-misalignment.

Anthropic. "Exploring Model Welfare." April 2025; "Claude Opus 4 and 4.1 Can Now End a Rare Subset of Conversations." 15 August 2025.

Common Sense Media. *Talk, Trust, and Trade-Offs: How and Why Teens Use AI Companions*. 2025.

European Commission. "Commission preliminarily finds TikTok's addictive design in breach of the Digital Services Act." 6 February 2026.

Imperva (Thales). *2025 Bad Bot Report*. 2025.

International Energy Agency. *Energy and AI*. Paris: IEA, 2025.

Kepios / DataReportal. *Digital 2026: Global Overview Report*. 2025. datareportal.com.

Newman, Nic, et al. *Reuters Institute Digital News Report 2026*. Oxford: Reuters Institute for the Study of Journalism, 2026.

OpenAI. "Sycophancy in GPT-4o: What Happened and What We're Doing About It." 29 April 2025; "Expanding on What We Missed with Sycophancy." 2 May 2025.



World Health Organization. "Social connection linked to improved health and reduced risk of early death." News release on the report of the WHO Commission on Social Connection, 30 June 2025.

NEWS AND WEB SOURCES

"72% of U.S. teens have used AI companions, study finds." *TechCrunch*, 21 July 2025.

"AI sycophancy makes people less likely to apologize and more likely to double down, study finds." *The Decoder*, 29 March 2026.

"A Primer on the Meta 2020 US Election Research Studies." *Tech Policy Press*, 2023.

"Chatbots lie about politics, Dutch watchdog warns." *EU Perspectives*, 27 October 2025.

"Deezer says 44% of songs uploaded to its platform daily are AI-generated." *TechCrunch*, 20 April 2026; "Music streamer Deezer says more than 50% of daily uploads are AI-generated." *TechCrunch*, 21 July 2026.

"EU Fines X €120 Million in First DSA Non-Compliance Decision." *eucrim*, December 2025.

"EU Lawmakers Reach Provisional Agreement to Delay Key EU AI Act Obligations." Sidley Austin, *Data Matters*, 22 June 2026.

"Exclusive: AI writing hasn't overwhelmed the web yet." *Axios*, 14 October 2025.

"Loneliness, Social Isolation Linked to 871,000 Annual Deaths, WHO Finds." *Health Policy Watch*, 4 July 2025.

"Meta will listen into AI conversations to personalize ads." *The Register*, 1 October 2025.

"Merriam-Webster names 'slop' as its 2025 word of the year." *NBC News*, December 2025.

"No social media for under-15s, Dutch government tells parents." *DutchNews.nl*, 17 June 2025.

"OpenAI explains why ChatGPT became too sycophantic." *TechCrunch*, 29 April 2025.

"OpenAI finally pulls trigger on ChatGPT ads after monthslong delay." *PPC Land*, January 2026.

"OpenAI is launching the Sora app, its own TikTok competitor, alongside the Sora 2 model." *TechCrunch*, 30 September 2025.

"OpenAI's 'smartest' AI model was explicitly told to shut down — and it refused." *Live Science*, 30 May 2025.

"Romanian Elections Targeted by 'Aggressive Hybrid Russian Action,' Declassified Documents Show." *Radio Free Europe/Radio Liberty*, 4 December 2024.

"Social media platforms removed 4.7 million accounts after Australia banned them for children younger than 16." *PBS NewsHour* / AP, 16 January 2026.



"Some workers already let AI do the thinking for them, Microsoft researchers find." *The Register*, 11 February 2025.

"TikTok US spin-off finalized, led by Oracle, Silver Lake and MGX." *DatacenterDynamics*, 23 January 2026.

"Bluesky's active user base." *GIGAZINE*, 12 August 2026.

"Can't pay attention? You're not alone." University of California, 2023.

"2024–2025 Romanian election annulment protests" and "Sycophancy (artificial intelligence)." *Wikipedia*. Accessed October 2026.

LECTURES, QUOTATIONS AND FURTHER INSPIRATION

Lotfi-Jam, Farzin. "Scales of Realtime." Lecture, Transmediale Festival, Akademie der Künste, Berlin, 2 February 2023.

Lorenz, Edward N. "Predictability: Does the Flap of a Butterfly's Wings in Brazil Set Off a Tornado in Texas?" Talk, AAAS, 29 December 1972.

Wolfram, Stephen. "The Concept of the Ruliad." *Stephen Wolfram Writings*, 10 November 2021.

Quotation sources: Forbes Quotes (H. C. Bailey); QuotesCosmos (Feynman); LibQuotes (Warhol); A-Z Quotes (Chomsky); Bookey (Eliasson); *Encyclopaedia Britannica* ("Chaos Theory"; "Cookie").

YouTube channels followed 2022–2026 for science and AI updates: Sabine Hossenfelder, Dr Brian Keating, Dr Alan D. Thompson, Fermilab, Steve Mould, Veritasium (Derek Muller), Joe Scott, Stanford, The Institute of Art and Ideas, World Science Festival.





quotes & epigraphs

i. Toynbee, Arnold J. “Man and Hunger: The Perspectives of History.” World Food Congress, Washington, D.C., 9 January 1963 – page 7

ii. Bailey, H. C., via Forbes Quotes, www.forbes.com/quotes/author/h-c-bailey/ – page 11

iii. Feynman, Richard, via QuotesCosmos.com – page 15





Image generation

"W": 512, "H": 768, "seed": 2360674296, "sampler": "DPM++ 2M Karras", "steps": 25,

"prompts":
"0": "male man lonely young student, (cute face:1.1) full body, boy, male 23 years old, walking on a wide street in amsterdam, holding an umbrella, stormy weather, (greyscale,monochrome,sketch), (greyscale, pure monochrome,sketch:2.0), artistic sketch style drawn on perfect white paper with thin lines,

"20": "lonely (half cyborg half male:1.6) 23 years old, walking on a curved wide street in amsterdam, holding an umbrella, stormy weather,(greyscale,monochrome,sketch), (greyscale, pure monochrome,sketch:2.0), artistic sketch style drawn on perfect white paper with thin lines,

"50": "lonely (cyborg, transgender, nobinary, SJand JL look alike face, :1.8) 23 years old, one leg dressed in legging and the other having a skirt, walking on a (wide:1.2) street (warped:1.5) street in amsterdam, holding an umbrella, stormy weather, (greyscale,monochrome,sketch), (greyscale, pure monochrome,sketch:2.0), artistic sketch style drawn on perfect white paper with thin lines,

"90": "(multiple floating christal balls in empty space having inside (damaged cyborg with pop up eye:1.6) with human face, walking on the street in amsterdam, holding an umbrella, stormy weather, the street is warped :1.3),(greyscale,monochrome,sketch), (greyscale, pure monochrome,sketch:2.0), artistic sketch style drawn on perfect white paper with thin lines,

"negative_prompts":
"nsfw,EasyNegative,ugly,huge eyes,text,logo,monochrome,worst face,(more than 2 hands,bad and mutated hands:1.3),(worst quality:2),(low quality:2),(blurry:2),horror,geometry,bad_prompt,(bad hands),(missing fingers),multiple limbs,bad anatomy,(interlocked fingers:1.2),Ugly Fingers,(extra digit and hands and fingers and legs and arms:1.4),(deformed fingers:1.2),(long fingers:1.2)",

Resources used in image generation:
Stable Diffusion web UI by AUTOMATIC1111, https://github.com/AUTOMATIC1111/ · Deforum extension by deforum-art, https://deforum.art/, GNU Affero General Public License v3.

LoRA: Sketch_20230919203655 by XiongSan (civitai.com); pensketch_without_pen_lora by nieta_art; TentacleHorror by konyconiN (civitai.com).
SD model: realcartoonRealistic_v12 by 7whitefire7, CreativeML OpenRAIL-M.
Embedding: EasyNegative by gsdf, https://huggingface.co/gsdf.



Contents

Introduction ..... page 7
The Lifecycle of Civilizations ..... page 11
Synthetic Life and Digital Immune Defense ..... page 15
Copy, Paste, Repeat ..... page 21
Democratic Technological Control Counterbalance ..... page 27
Synthetic Ascendancy and Human Roles ..... page 31
Personalized Digital Constructs ..... page 35
Digital Echoes ..... page 39
The Dawn of the Digital Age ..... page 43
Digital Echo Chambers ..... page 49
The Agreeable Machine ..... page 55
Ballots in the Bubble ..... page 61
Artistic Expression in a Digitalized World ..... page 65
Navigating the Digital Landscape ..... page 71
Conclusion ..... page 75
Bibliography ..... page 81
Quotes & epigraphs ..... page 88



Acknowledgement

I extend my heartfelt gratitude to the faculty of IMD–KABK for their guidance and support throughout my academic journey. Your wisdom and insights have been essential in shaping my understanding of, and approach to, our shared field of study.

To my family, words cannot fully express my appreciation for your constant support and understanding. Your encouragement and belief in my dreams have been the cornerstone of my achievements, my pillar of strength and comfort.

I am equally grateful to my colleagues, whose camaraderie and intellectual challenges helped define who I am as an artist, and helped me understand a new generation.

A note of thanks to OpenAI and Microsoft for ChatGPT and the Word Editor, which helped me restructure complex sentences in the first version, and to Anthropic's Claude, which assisted with the research, fact-checking and editing of this second version. It seems fitting that a thesis about synthetic minds was written in conversation with them, and every claim in it was checked against human sources.

My thanks also extend to the world at large, whose intricate complexity and sheer beauty have constantly fuelled my curiosity and passion.

To all, my deepest appreciation for your contributions to my journey.

Sorin Angeleanu

