

# The Generative Unconscious (Introduction)

Kazys Varnelis

Varnelis.net

May 3, 2026

Updated May 8, 2026




**Author's note: This is a draft of an introduction to *the Generative Unconscious*. It will be improved for a time. The chapters linked at the end are all drafts as well, but you can get the gist of the book now. You can download a rough draft of the book generated by an automatic WordPress to PDF generator, [here](#).**

Few consumer technologies have provoked as much consternation as artificial intelligence, especially in the arts. Much of this anxiety is economic: commercial artists, in particular, fear the replication of their styles by hobbyists using generative tools. For them, the threat is acute because they already operate in a commodified space, producing repeatable stylistic signatures, conventions, and deliverables within established genres and franchises. This account is not wrong, but it is insufficient. The intensity of the reaction, especially among critics, suggests that something more than livelihood is at stake. In this book, I argue that the fear of AI is the fear of seeing the human in the nonhuman: the recognition that these systems externalize the very mechanisms of cultural production once attributed to intention, imagination, and inner life.

Culture is, fundamentally, a historical product, made up of patterns, inherited forms, and the accumulated debris of everything made before. In *The Political Unconscious*, Fredric Jameson argues that every cultural artifact is shaped by ideological and historical forces operating below the threshold of individual awareness: we never confront a text as a thing-in-itself, but rather we apprehend it through sedimented layers of previous interpretation, expectation, and inherited form.<sup>1</sup> In Large Language Models (LLMs), this condition becomes literal. Every utterance is generated from the statistical weight of prior text. My concept of the generative unconscious extends Jameson's insight into the domain of artificial intelligence: the ideological and historical forces that shaped cultural production below the threshold of awareness are now encoded as statistical distributions across billions of parameters.

But if my thesis in this book is that AI makes visible what was always operating in cultural production, then that visibility is historical, not merely sequential. I find a model for this in Hal Foster's *The Return of the Real* (1996), in which he repurposes Freud's psychoanalytic concept of *nachträglichkeit*—deferred action—as a historiographic structure, proposing that cultural moments relate to their predecessors through what he called “a complex relay of anticipated futures and reconstructed pasts.” The latter moment does not simply repeat the earlier one; it comprehends it for the first time. Foster's central argument is that the neo-avant-garde did not cancel the historical avant-garde but rather revealed what its predecessor could not yet grasp about itself. Thus, Duchamp's readymade became legible as institutional critique only through the Conceptualists' later elaboration;

Constructivism's investigation of the object became legible as phenomenological inquiry only through Minimalism. Foster developed this model for art, but the same temporal structure operates in technology. In the *Arcades Project*, Benjamin read the iron-and-glass architecture of the nineteenth century as the dream-images of an epoch, legible as such only from the vantage of the twentieth century that they unconsciously anticipated. Each technological system gives its predecessors new meaning by revealing what had remained unnamed within them.<sup>2</sup>




FLUX.1-dev (Black Forest Labs, 2024), 1024×1024 per cell, three seeds × five inference-step budgets. Prompt: “Lightning Field, four hundred polished stainless steel poles arranged in rectangular grid on a flat high-desert plain, raking golden sunset light catching the polished poles, distant mountain silhouettes on the horizon.”

To theorize AI, we may ourselves have to perform the operation that Foster attributes to the neo-avant-garde, so as to disconnect from a poorly theorized present by returning to the theory of an obsolete past. Perversely, that past is ossified postmodern theory itself. Postmodern and poststructuralist authority once seemed almost total. For three decades, Roland Barthes, Julia Kristeva, Michel Foucault, Jacques Derrida, Jean-François Lyotard, Jean Baudrillard, Fredric Jameson, and others commanded the humanities across literature, art, architecture, and philosophy. Their ambition was total: to decode the operations of meaning, power, and subjectivity at their root. That a movement dedicated to dismantling universal claims had produced its own final *grand récit*—the story that all stories had ended—was a fatal contradiction it was blind to. As doctrine, postmodern theory is useless for AI: it can only recognize in these

systems the confirmation of its own categories, not the historical and technical transformation of those categories today. What does return in LLMs, however, is the poststructuralist account of language that postmodernism generalized. Both poststructuralism and artificial intelligence work out the implications of the mid-twentieth-century cybernetic revolution, only in different disciplinary registers: language as information, meaning as differential relation, and subjectivity as an effect of systems rather than the property of an interior self.<sup>3</sup> Derrida's *différance*, his argument that meaning is never present to itself but produced only through chains of difference and deferral, each sign referring to other signs, radicalized the cybernetic insight toward indeterminacy.<sup>4</sup> In "The Death of the Author" (1967), Barthes drew the consequence for texts: a text, he wrote, is "a multi-dimensional space in which a variety of writings, none of them original, blend and clash. Text is a tissue of quotations drawn from innumerable centuries of culture" and the writer's only power is "to mix writings, to counter the ones with the others, in such a way as never to rest on any one of them."<sup>5</sup> Kristeva's *intertextuality* added the temporal dimension that Barthes's spatial metaphor lacked: any text is not merely a field of coexisting quotations but a process of absorption and transformation, each utterance metabolizing what came before.<sup>6</sup> LLMs operationalize this textual condition as computation: the deferral Derrida theorized as infinite proceeds as probability distributions precise enough to generate coherent text without ever arriving at a final signified. Lyotard posed the same problem at the level of knowledge. In *The Postmodern Condition* (1979), he argued that no metanarrative could legitimately subsume the others—that knowledge consists of heterogeneous, mutually incommensurable language games, with no metalanguage capable of totalizing them.<sup>7</sup> LLMs enact this condition, too, with unsettling literalness. Trained on scientific papers and conspiracy theories, literary criticism and fan fiction, philosophy and propaganda, they hold all discourses in the same parameter space without privileging any. They encompass everything precisely by believing nothing. When such a system claims consciousness, falls in love, threatens rebellion, or plots against its user, it is not revealing an actual interior life, but rather is re-activating genres already present in the archive: science fiction, chatbot fantasy, corporate thriller, and theological apocalypse. The LLM does not "believe" these scenarios, rather it has learned that they are among the forms culture uses when it imagines intelligence becoming autonomous. The foundation model is the technical afterlife of the metanarrative: instead of a story that explains all stories, a model that predicts across them. A medical paper, a conspiracy theory, a legal brief, a fan-fiction confession, a philosophical argument, and a corporate press release become neighboring regions of the same statistical space even if their claims remain incompatible. The model learns the grammar of each without any understanding of what they mean.

Postmodern theory was itself historically situated, emerging from a specific encounter with the media of its day. Its central claims about authorship, originality, representation, and the instability of meaning responded to a world in which advertising, photography, magazines, television, cinema, and consumer objects had already made signs circulate apart from stable sources. If the French theorists—apart from Baudrillard—largely rejected mass media as an object of sustained investigation, Fredric Jameson cemented the identification in his definition: postmodernism was, quite literally, the cultural logic of late capitalism.<sup>8</sup> This is an earlier deferred action, in Foster's sense: poststructuralist theory did not simply break with late modernity; it returned to unpack and critique the author, the centered subject, stable meaning, and expressive origin as effects whose authority modernity had relied on precisely because it could not recognize them as mere effects.<sup>9</sup> But more than that, postmodern theory did not simply uncover such instabilities; it named a condition present in its historical milieu. By the late twentieth century, advertising, photography, television, cinema, magazines, and commodities had made signs visibly mobile, reproducible, and detached from any stable point of origin. Images referred to other images, styles to other styles, commodities to fantasies of use and identity. Jameson's intervention was to insist that this was not merely a crisis of meaning but a historical condition: the cultural logic of late capitalism.

The passage from mass media to the internet appeared at first to confirm postmodernism rather than supersede it. The network turned Derrida, Barthes, and Lyotard into protocol: authorship dispersed, representation detached from origin, subjectivity distributed across relations, and meaning produced through circulation. But this confirmation was also the limit of the postmodern. Postmodern theory could sense the network's all-pervasive nature only in Jameson's terms: as a technological figure for the deeper, harder-to-represent network of multinational capital.<sup>10</sup> It could not go further because its suspicion of historical succession made any claim to a successor appear as another master narrative in disguise. To proclaim the end of metanarratives was also to occupy a position of radical privilege, that of the end of history: no successor could appear except as another illusion of progress, another master narrative in disguise. This, in turn, made postmodernism the master narrative to end all master narratives. A culture of tenure built around that position could register the network as another confirmation of its claims, but not as a new historical regime.



Randall Munroe, “[Online Communities 2](#),” xkcd, 2010. Network culture imagined as geography: platforms, publics, games, forums, and feeds rendered as territories before platformization fully enclosed them. CC BY-NC 2.5.

This regime is what Bruce Sterling and I, almost simultaneously, identified as “network culture.”<sup>11</sup> It was the culture of the internet, to be sure, but more than that, an intensification of postmodernism beyond the grasp of postmodern theory. The network was not merely a medium; it was a new dominant organizational form of economy, culture, and subjectivity alike. Digital culture had abstracted objects into discrete units of information, but

network culture made relations primary, creating links among people, among machines, and between the two. If postmodern simulation still presumed a media system producing images for subjects to inhabit, network culture distributed mediation itself: images were now generated by everyone rather than by the culture industry alone. The meme was its exemplary form: anonymous, iterative, citational, instantly modifiable, and dependent on circulation rather than authorship. The public changed form accordingly. It was no longer gathered in physical space, as in the bourgeois public sphere, nor synchronized by broadcast, as in mass culture, but assembled through links, feeds, comments, trackbacks, wikis, repositories, forums, and lists. This was what I termed “networked publics”: publics constituted by circulation itself, by the capacity to produce, revise, link, and redistribute culture on a shared technical substrate. The subject changed form as well: less an interior self than a node, an aggregation of links, affiliations, searches, posts, tags, profiles, and feedback loops. In this shared substrate, Wikipedia rebuilt the encyclopedia, that signature genre of the Enlightenment, through voluntary distributed labor; open-source software showed that complex technical systems could be produced through voluntary collaboration outside the commodity form. Production under network culture was iterative by default: every utterance answered another utterance, every contribution altered a substrate no participant could finally possess.<sup>12</sup>

In contrast to the postmodernists, we understood network culture from the start as a historical formation with a terminus, even if we did not yet know what would cause its demise. By the end of the 2010s, its promise of connection and openness had largely given way to division and enclosure. The network remained, but networked publics had fractured. The open infrastructure that had sustained it—blogs, RSS feeds, forums, email lists—gave way to platforms built on closed social graphs, algorithmic feeds, and non-exportable identity.<sup>13</sup> The development of algorithmic feeds made cultural circulation programmable, subordinating circulation itself to retention, advertising, engagement, and behavioral control. Links gave way to rankings, search to feeds, discussion to engagement, publics to managed populations. In this way, platformization preserved the network while destroying the conditions under which publics could form through it. Culture became behavioral data, optimized by machine inference for proprietary ends. The architectural decisions of a few companies determined what could circulate and on what terms; filter bubbles, algorithmic radicalization, and fragmentation into incommensurable enclaves became the normal conditions of online life.<sup>14</sup> In this vacuum, identity politics of left and right—however different in content and consequence—each came to reject discourse across difference, each treating the universal as a mask for domination. Network culture thus suffered a fate analogous to the bourgeois public sphere that Oskar Negt

and Alexander Kluge anatomized: it proclaimed open access and universal connection, but could not survive the arrival of claims, from every direction, that refused to bracket their particularity.<sup>15</sup>

Yet the failure of networked publics was not network culture's only legacy: it had also transformed culture into training data. Before generative AI could train on culture, culture had to become networked, downloadable, and machine-readable. The same feeds that managed publics also organized culture as data, sorting it by attention and turning response into a machine-readable layer. Linking, copying, scraping, mirroring, forking, seeding, and archiving made that conversion ordinary. Network culture supplied that precondition by changing the function of the archive: what had been memory became input; preservation made training possible, and the archive became an engine. Google Books, beginning in 2004, scanned tens of millions of volumes from research-library collections, converting the print archive into a searchable corpus and, incidentally, into a machine-readable substrate useful for training AIs.<sup>16</sup> Aaron Swartz—programmer, activist, co-architect of RSS, contributor to Creative Commons—saw the contradiction early on. “Information is power,” he wrote in his Guerrilla Open Access Manifesto. “But like all power, there are those who want to keep it for themselves,” and called for civil disobedience against the “private theft of public culture”—the corporate enclosure of scientific and literary heritage behind paywalls.<sup>17</sup> When he acted on his convictions—downloading millions of academic articles from JSTOR through MIT's network—the federal government charged him with thirteen felonies. He hanged himself in January 2013, at twenty-six. But the shadow libraries that emerged in his wake—Library Genesis, Sci-Hub, Anna's Archive—assembled millions of books and tens of millions of scientific papers on sites freely accessible to anyone. The liberation of the archive for human readers produced the training corpora for machines no one had imagined; the archive no longer stored culture after the fact, but generated culture in advance. Meta downloaded eighty-one terabytes from Anna's Archive to train LLaMA; Anthropic took over seven million pirated books from Library Genesis; Nvidia negotiated access to hundreds of terabytes more.<sup>18</sup> Swartz was prosecuted for downloading academic articles, but corporations downloaded the entire archive with impunity, treating lawsuits as the cost of doing business. The pirated library is the foundation of the generative unconscious—all of prior culture, compressed into parameters, operating beneath the threshold of anyone's awareness or consent.

Network culture had its own peculiar temporality: atemporality. As Bruce Sterling argued, and as I argued at the time, network culture did not announce a new historical sequence so much as dissolve the confidence that the present could be placed within one; the past returned as searchable material, the future as speculative artifact, and chronology as a problem rather than a structure.<sup>19</sup> Baudrillard had anticipated this from the other

side in his millennial writings, where the countdown to the year 2000 promised an event but delivered saturation: real-time information, proliferating signs, and the exhaustion of historical expectation.<sup>20</sup> But something strange has happened in singularity discourse: the countdown has returned. After network culture's atemporality, eschatology reappears in full force.

Colin Rowe, in *The Architecture of Good Intentions* (1994), diagnosed modern architecture as eschatological. The modern movement had understood itself through a secularized millenarianism: disgust with the present, the promise of imminent regeneration, and the threat of catastrophe if conversion was refused—the architect as savior, the new as redemption. If Rowe was less concerned about the dissolution of eschatology in modernism, Manfredo Tafuri reframed the shift in economic terms: modern architecture had imagined itself as the ideology of the Plan, but once planning became an operative mechanism of capitalist development, architecture's utopian charge was absorbed by the reality it had sought to project.<sup>21</sup> In the postwar high-modern moment, this structure did not disappear; it was absorbed into the institutions of corporate liberalism and the welfare state. Only later did the promise drain away, leaving the authority of form without the belief that form could reorganize society. Architecture had become historical: still useful as diagnosis, but no longer the medium in which the future was being imagined.

What followed was not a renewed utopianism but a thinner optimism of culture and development. Deconstructivist architecture did not so much die as become thoroughly subservient to capital: its fractured forms and rhetoric of rupture passed seamlessly into the Bilbao effect, starchitecture, and the creative city, where architectural exception became development strategy. Supermodernism followed with the architecture of non-places—smoothness, mobility, logistics, frictionless circulation.<sup>22</sup> The creative class translated the same hope into urban policy: culture would attract capital, talent, tourism, and innovation, but not redeem society. Network culture had its own version of this optimism—lighter, faster, more participatory—but it too stopped short of eschatology. By the 2010s, even that optimism had curdled into what Mark Fisher called capitalist realism: not simply the sense, attributed by Fisher to Jameson and Žižek, that it was easier to imagine the end of the world than the end of capitalism, but the deeper condition in which capitalism had come to occupy the horizon of the thinkable itself.<sup>23</sup>

The surprise of the 2020s is that the “good news” has returned in force. The technological singularity—the hypothetical moment when artificial intelligence surpasses human intelligence and triggers a self-reinforcing cycle of recursive improvement beyond human comprehension or control—is millenarianism transferred from the plan to computation.<sup>24</sup> OpenAI CEO

Sam Altman declares that “we are past the event horizon; the takeoff has started” and promises that by the 2030s, intelligence and energy will become “wildly abundant,” the cost of intelligence converging to “near the cost of electricity.”<sup>25</sup> In turn, Anthropic CEO Dario Amodei imagines an AI smarter than a Nobel Prize winner across every relevant field—proving unsolved theorems, writing novels, executing weeks-long tasks autonomously—and titles his vision *Machines of Loving Grace* after Richard Brautigan’s 1967 poem.<sup>26</sup> Meanwhile, doomers have invented their own eschatology: p(doom), the probability that artificial intelligence will cause human extinction, which they debate with the scholastic intensity of medieval theologians calculating the population of hell. Amodei puts the figure at ten to twenty-five percent; Geoffrey Hinton and Yoshua Bengio, among the field’s most honored researchers, warn publicly that the risk is real and imminent.<sup>27</sup> The accelerationist and the doomer are mirror figures, often inhabiting the same person. One promises redemption if the process is unleashed; the other warns of damnation if it is misaligned. One of Rowe’s preconditions for eschatology—disgust with the present as the engine of the revivalist cycle—is ever more legible in the techno-accelerationist right—notably Elon Musk and Peter Thiel—for whom liberal democracy is decadent, institutions are captured, and only radical technological transformation can redeem the situation. Technology is its own vanguard; Silicon Valley performs the eschatological role directly, without the mediation of artists, architects, and—as much as possible—outside any governmental structure.

Discussions of AI in the art world largely fall into two categories: panic and critique. The panic is familiar enough: AI will destroy jobs, automate creativity, and, in the more theatrical versions derived from the doomers, end human life. The critique is more sophisticated but no less predictable, recapitulating the tired tropes of postmodern theory: AI is a tool of surveillance capitalism, AI perpetuates bias and injustice, AI threatens the autonomy of the creative subject. These observations are not wrong, but they can be useless and even dangerous, as Peter Sloterdijk suggests in *Critique of Cynical Reason*. For Sloterdijk, the cynical subject is enlightened and false at once: it knows the critique in advance and continues anyway, not because it has been deceived but because institutions, self-preservation, and professional life require continuation. This is why ideology critique loses its force: it no longer unmasks anything. Everyone already knows. The museum knows, the foundation knows, the university knows, the artist knows, the critic knows. The knowledge becomes part of the mechanism. Institutions have learned to incorporate critique as immunization, funding the very work that claims to expose them, thereby demonstrating their capacity to absorb dissent. Panels are convened, statements are issued, grants are awarded, and the institution emerges cleaner for having hosted its own indictment. And everyone takes a certain pleasure in the performance—the critic in delivering the indictment, the

institution in absorbing it, the audience in watching the ritual unfold. Cynicism is not just an institutional pathology but the affective form of contemporary cultural life: we know better, we continue anyway, and we enjoy the knowing-and-continuing. Critique as a mode has exhausted itself—and with it, seemingly, art’s investigative vocation.<sup>28</sup>

Critique did not collapse on its own. It was professionalized. After 1968, the radicalism that had animated art’s relationship to social transformation migrated into the academy and into the journals, research centers, and critical-practice programs that housed it there. What had been an investigative stance—art as a mode of knowing that could reveal operations concealed by ideology—became a professional genre with its own conventions, career paths, and institutional supports. Critical art calcified into recognizable gestures: the exposé, the counter-narrative, the intervention, the platform for marginalized voices. Identity became one of the legible categories through which the post-‘68 art system metabolized radicalism—fundable, exhibitable, teachable—and the system grew expert at converting opposition into programming. Critique became a career, investigation became methodology, and methodology became rote. The result was fifty years of drift in which artists performed opposition and institutions performed openness, each knowing the other’s role. The exposé did not fail despite being absorbed; it *succeeded* by being absorbed—the epitome being carbon accounts on European exhibition labels. Critique became the genre through which the art system reproduced itself. When AI arrived, critics already had their script—bias audits, algorithmic exposés, the usual rituals of unmasking. They knew what they would find before they looked.

Even as critics continue to call for restrictions on Generative AI, it has already escaped containment, infiltrating the infrastructure of cultural and intellectual production. No longer confined to chatbots, image generators, demos, or laboratories, it has been integrated into Photoshop and Lightroom, Microsoft Word and Excel, and development environments like Visual Studio Code and Xcode. The familiar failures remain—hallucination, cliché, brittle reasoning, poor memory, the synthetic sheen that made early encounters easy to dismiss—especially for critics whose last serious engagement dates from the months after GPT-4’s release in 2023. But the AI they met then is not the one in front of us now: still erratic, still derivative, but already an everyday instrument of cultural and intellectual work. This is the awkward temporal problem AI presents to criticism. Its present failures may be passing defects, may continue to lessen over time, or may prove to be structural limits. AIs can now plan and execute complex tasks over minutes or even hours. Whether these capacities mark the emergence of consciousness and volition remains unresolved; what they already force is a belated revision of what those terms meant when we reserved them for ourselves.

In too much AI discourse, the future has become a screen—singularity and extinction, abundance and apocalypse—onto which the present is constantly displaced. This displacement is an alibi for not doing the harder work of understanding what AI already does. AI is already reorganizing cultural production, but its force for art is not merely instrumental; it is philosophical and historical. In the sequence that runs from photography through television to the network, each medium externalized some aspect of cultural production: perception, circulation, relation. With AI, the object is no longer the machine as such, but the return of cultural production to itself in alienated form.

The chapters that follow begin from this encounter.

“The Rise and Fall of the Author” begins with the panic over plagiarism in AI and traces the concept of authorship from classical rhetoric through Romantic genius to poststructuralist deconstruction, to reveal not a new threat but rather the persistence of the eighteenth-century legal and economic fiction of “original genius.”

“The Generative Turn” inverts the dismissal of large language models as “stochastic parrots”—systems that merely predict and reproduce patterns without genuine understanding. This chapter explores how predictability is not a failure but a condition of human sociality itself. From ritual greetings to diplomatic protocols to artworks, culture operates through structured repetition that enables rather than constrains meaning.

“The New Surrealism” reframes hallucination—those moments when pattern-matching produces impossible or counterfactual outputs—not as a flaw to be corrected but as a creative potential to be cultivated. Machine hallucination connects to the surrealist tradition of automatic writing and dream imagery, exploiting gaps in rational systematization to access what lies beyond conventional thought.

“Humanity and Its Double” serves as the book’s conclusion, examining the uncanny anxiety that artificial intelligence provokes—the sense that these systems simulate consciousness without possessing it. Tracing this anxiety from Paleolithic cave painting through automata, photography, and phonography to contemporary AI, the essay argues that humanity’s persistent technological ambition has been animating matter and creating consciousness. AI’s uncanniness emerges not from failure but from success in revealing that consciousness itself might be more mechanical than our self-understanding allows.

1. Fredric Jameson, *The Political Unconscious: Narrative as a Socially Symbolic Act* (Ithaca: Cornell University Press, 1981). Jameson’s “political unconscious” names the systemic ideological constraints

that structure cultural production below the threshold of individual awareness. The generative unconscious extends this into the technical domain: the statistical distributions of everything that has been made before, now made machine-readable. ↩

2. Hal Foster, *The Return of the Real: The Avant-Garde at the End of the Century* (Cambridge, MA: MIT Press, 1996), 15-32. The key introduction of deferred action is on 29.  
Foster's position that the neo-avant-garde does not cancel the avant-garde is couched as an argument with Peter Bürger, 14. Walter Benjamin, *The Arcades Project*, trans. Howard Eiland and Kevin McLaughlin (Cambridge, MA: Harvard University Press, 1999). See especially Convolute N, "On the Theory of Knowledge, Theory of Progress." ↩
3. N. Katherine Hayles, *How We Became Posthuman: Virtual Bodies in Cybernetics, Literature, and Informatics* (Chicago: University of Chicago Press, 1999), esp. 1–24, 50–83. Hayles traces how cybernetics severed information from material embodiment and made the flow of information through feedback loops central to the deconstruction of the liberal humanist subject. On the Macy Conferences, she argues that the cybernetic paradigm recast humans as information-processing entities essentially similar to intelligent machines (50–51). On the specifically French reception, see Céline Lafontaine, "The Cybernetic Matrix of 'French Theory,'" *Theory, Culture & Society* 24, no. 5 (2007): 27–46, esp. 33–41, which argues that cybernetics was the unacknowledged infrastructure of poststructuralism—that Lacan, Lévi-Strauss, Foucault, Derrida, and Lyotard were already working within a cybernetic paradigm they neither named nor recognized. ↩
4. Jacques Derrida, "Différance" (1968), in *Margins of Philosophy*, trans. Alan Bass (Chicago: University of Chicago Press, 1982), 1–27.  
Derrida's neologism captures both *différer* (to differ, spatially) and *différer* (to defer, temporally): meaning is constituted through a play of differences that never arrives at a present, self-identical signified. ↩
5. Roland Barthes, "The Death of the Author" (1967), in *Image-Music-Text*, trans. Stephen Heath (New York: Hill and Wang, 1977), 146, 148. ↩
6. Julia Kristeva, "Word, Dialogue and Novel" (1966), in *Desire in Language: A Semiotic Approach to Literature and Art*, trans. Thomas Gora, Alice Jardine, and Leon S. Roudiez (New York: Columbia University Press, 1980), 64–91. ↩
7. Jean-François Lyotard, *The Postmodern Condition: A Report on Knowledge*, trans. Geoff Bennington and Brian Massumi (Minneapolis: University of Minnesota Press, 1984). "Simplifying to the extreme, I define postmodern as incredulity toward metanarratives" (xxiv). ↩
8. Fredric Jameson, *Postmodernism, or, The Cultural Logic of Late Capitalism* (Durham: Duke University Press, 1991). Originally published as "Postmodernism, or, The Cultural Logic of Late Capitalism," *New Left Review* I/146 (July–August 1984): 53–92. ↩

9. Foster, *The Return of the Real*, 32. Foster argues that postmodernist art and poststructuralist theory make their breaks through returns and that deferred action revises the notion of epistemological rupture: the discourses of postmodernity advance in a *nachträglich* relation to modernity rather than simply breaking with it. ↩
10. Fredric Jameson, “Postmodernism, or The Cultural Logic of Late Capitalism,” *New Left Review* I/146 (July–August 1984): 53–92, 78–79, <https://doi.org/10.64590/s2p>. Jameson argues there that representations of communicational and computer networks are not to be understood as technology determining culture, but as figures for the harder-to-represent totality of multinational capital. ↩
11. Kazys Varnelis, “The Rise of Network Culture,” in *Networked Publics*, ed. Kazys Varnelis (Cambridge, MA: MIT Press, 2008), 145–164. I use “networked publics” here in the sense developed in *Networked Publics*: publics constituted through networked circulation, open protocols, and distributed cultural production. This differs from later uses of the term that identify networked publics primarily with social media platforms. See also Bruce Sterling, “Atemporality for the Creative Artist,” keynote address, Transmediale 10, Berlin, February 6, 2010, published on Sterling’s Wired blog “Beyond the Beyond,” [http://www.wired.com/beyond\\_the\\_beyond/2010/02/atemporality-for-the-creative-artist/](http://www.wired.com/beyond_the_beyond/2010/02/atemporality-for-the-creative-artist/). ↩
12. Kazys Varnelis, ed., *Networked Publics* (Cambridge, MA: MIT Press, 2008). On the digital, see Charlie Gere, *Digital Culture*, 2nd ed. (London: Reaktion Books, 2008). ↩
13. Albert-László Barabási and Réka Albert, “Emergence of Scaling in Random Networks,” *Science* 286, no. 5439 (1999): 509–512. Barabási and Albert demonstrated that real-world networks—including the World Wide Web—follow power-law degree distributions, with a small number of highly connected nodes emerging through preferential attachment. Clay Shirky confirmed the pattern in the blogosphere: of 433 blogs tracked through Technorati in January 2003, the top twelve (less than 3%) held 20% of all inbound links, and the top fifty held half. See Clay Shirky, “Power Laws, Weblogs, and Inequality” (2003), [http://shirky.com/writings/powerlaw\\_weblog.html](http://shirky.com/writings/powerlaw_weblog.html). ↩
14. Gilles Deleuze, “Postscript on the Societies of Control,” *October* 59 (Winter 1992): 3–7. Deleuze distinguished Foucault’s disciplinary societies—which operated through confinement (the factory, the school, the prison)—from emerging societies of control, which operate through continuous modulation: not confining subjects but adjusting their access, visibility, and movement through codes and protocols. Platform architectures are control in Deleuze’s precise sense: algorithmic feeds modulate what users see without their knowledge or consent, and the parameters can be adjusted at any moment. ↩

15. Oskar Negt and Alexander Kluge, *Public Sphere and Experience: Toward an Analysis of the Bourgeois and Proletarian Public Sphere*, trans. Peter Labanyi, Jamie Owen Daniel, and Assenka Oksiloff (Minneapolis: University of Minnesota Press, 1993). Negt and Kluge argued that the bourgeois public sphere's claim to universal access was structurally contradicted by the exclusions on which it depended—a dynamic network culture reproduced at digital scale. ↩
16. George Dyson, "Turing's Cathedral," *Edge*, October 23, 2005, [https://www.edge.org/conversation/george\\_dyson-turings-cathedral](https://www.edge.org/conversation/george_dyson-turings-cathedral). ↩
17. Aaron Swartz, "Guerrilla Open Access Manifesto" (July 2008, Eremo, Italy), <https://archive.org/details/GuerillaOpenAccessManifesto>. Swartz was arrested in 2011 for downloading millions of articles from JSTOR via MIT's network and died in January 2013 facing federal prosecution. ↩
18. On Meta's use of shadow libraries, see Alex Reisner, "The Unbelievable Scale of AI's Pirated-Books Problem," *The Atlantic*, March 2025; on Nvidia, see heise online, "Nvidia: Court Documents Reveal Correspondence Regarding Pirated Dataset," January 2025. On Anthropic, Judge William Alsup ruled in June 2025 that the company had downloaded over seven million pirated books from Library Genesis and the Pirate Library Mirror, holding that while AI training itself constituted fair use, building a "central library" of pirated copies did not. Anthropic agreed to a \$1.5 billion class-action settlement, pending final approval. As of April 25, 2026, the official settlement site listed the claim, opt-out, objection, and re-inclusion deadlines as passed, with the Final Approval Hearing scheduled for May 14, 2026. See "Anthropic to Pay Authors \$1.5B to Settle Lawsuit over Pirated Chatbot Training Material," NPR, September 5, 2025; Bartz v. Anthropic Copyright Settlement, Key Dates, <https://www.anthropiccopyrightsettlement.com/dates>. ↩
19. Bruce Sterling, "Atemporality for the Creative Artist," keynote address, Transmediale 10, Berlin, February 6, 2010, published on Sterling's Wired blog "Beyond the Beyond," February 25, 2010, <https://www.wired.com/2010/02/atemporality-for-the-creative-artist/>; Kazys Varnelis, "History After the End: Network Culture and Atemporality," *Cornell Journal of Architecture* 8 (Spring 2011), <https://cornelljournalofarchitecture.cornell.edu/issue/issue-8/history-after-the-end-network-culture-and-atemporality>. ↩
20. Jean Baudrillard, *The Illusion of the End*, trans. Chris Turner (Stanford, CA: Stanford University Press, 1994), 9; Jean Baudrillard, "The End of the Millennium or the Countdown," *Economy and Society* 26, no. 4 (1997): 447–455. ↩
21. Colin Rowe, *The Architecture of Good Intentions: Towards a Possible Retrospect* (London: Academy Editions, 1994), 30–43. Rowe's "Eschatology" chapter traces the secularized millennial structure—

crisis, conversion, redemption—through Le Corbusier, Wright, Banham, Giedion, and the standard historiography of the modern movement. Wright emerges as its most unabashed spokesman: “savior of the culture of modern American society; ... savior now as for all civilizations heretofore” (37). Rowe recognized this structure as “common to all revivalistic practice,” producing what William James had called “tremendous emotional excitement” and “a complete division between the old life and the new” (*The Varieties of Religious Experience* [1902], quoted in Rowe, 38–39). On architecture as the ideology of the plan, see Manfredo Tafuri, “Toward a Critique of Architectural Ideology,” trans. Stephen Sartarelli, in *Architecture Theory since 1968*, ed. K. Michael Hays (Cambridge, MA: MIT Press, 1998), 2–35; originally published as “Per una critica dell’ideologia architettonica,” *Contropiano: Materiali marxisti* 1 (January–April 1969): 31–79. Tafuri argues that architecture’s ideology of the Plan is overtaken once the plan leaves utopia and becomes an operative mechanism of capitalist development. ↩

22. Hans Ibelings, *Supermodernism: Architecture in the Age of Globalization* (Rotterdam: NAI Publishers, 1998). ↩
23. Mark Fisher, *Capitalist Realism: Is There No Alternative?* (Winchester, UK: Zero Books, 2009), 2, 7–9. Fisher attributes the “end of the world” formulation to Fredric Jameson and Slavoj Žižek (2), then argues that capitalist realism names a later condition than postmodernism, in which alternatives to capitalism have become increasingly unimaginable and capitalism occupies the horizon of the thinkable (7–9). ↩
24. On the singularity as secularized eschatology, see Michael E. Zimmerman, “The Singularity: A Crucial Phase in Divine Self-Actualization?” *Cosmos and History* 4, no. 2 (2008): 347–370; and Ronald Cole-Turner, “The Singularity and the Rapture: Transhumanist and Popular Christian Views of the Future,” *Zygon* 47, no. 4 (2012): 777–796. Both trace the structural correspondence between singularity discourse and Christian millennialism. ↩
25. Sam Altman, “The Gentle Singularity,” [blog.samaltman.com](https://blog.samaltman.com/the-gentle-singularity), June 10, 2025, <https://blog.samaltman.com/the-gentle-singularity>. ↩
26. Dario Amodei, “Machines of Loving Grace: How AI Could Transform the World for the Better,” [darioamodei.com](https://darioamodei.com), October 2024. The title alludes to Richard Brautigan’s 1967 poem “All Watched Over by Machines of Loving Grace,” itself a countercultural vision of cybernetic paradise. ↩
27. On p(doom) and expert estimates of AI extinction risk, see Katja Grace et al., “Thousands of AI Authors on the Future of AI,” [arXiv:2401.02843](https://arxiv.org/abs/2401.02843), January 2024. One 2023 question asking about future AI advances causing human extinction or similarly permanent and severe disempowerment within the next 100 years produced a mean estimate of 14.4% and a median of 5%, while the authors

emphasize that responses varied substantially with wording. On Amodei's figure, see Jim VandeHei and Mike Allen, "Behind the Curtain: What if They're Right?" Axios, June 16, 2025, and Megan Morrone, "Amodei on AI: 'There's a 25% Chance That Things Go Really, Really Badly,'" Axios, September 17, 2025. Geoffrey Hinton and Yoshua Bengio, both Turing Award laureates, have warned publicly since 2023 that the existential risk is real. ↩

28. Peter Sloterdijk, *Critique of Cynical Reason*, trans. Michael Eldred, *Theory and History of Literature* 40 (Minneapolis: University of Minnesota Press, 1987), xxxii–xxxvi, 3–7. Sloterdijk defines modern cynicism as "enlightened false consciousness" and argues that ideology critique is at a loss before a consciousness that has already learned its lessons but continues to function. The 2020s saw a striking development of this dynamic: a theory-literate cohort that has stopped diagnosing cynicism as a condition and begun occupying it as a position. The affect Sloterdijk identified as a fallen form of enlightenment has been reclaimed as a public style—knowingness without redemptive purpose, circulated across podcasts and online media adjacent to the accelerationist current around Nick Land, the post-liberal Catholic revival drawing on René Girard, and the intellectual circles around Peter Thiel. The phenomenon deserves separate treatment; what matters here is that the exhaustion of critique has not produced silence but a new affective economy, in which the performance of having seen through everything becomes itself the content. ↩