

# Strategic misrecognition and speculative rituals in generative AI

Sun-ha Hong

Stanford University, USA

✉ residual@stanford.edu

## Abstract

Public conversation around generative AI is saturated with the ‘realness question’: is the software really intelligent? At what point could we say it is thinking? I argue that attempts to define and measure those thresholds misses the fire for the smoke. The primary societal impact of realness question comes not from the constantly deferred sentient machine of the future, but its present form as *rituals of misrecognition*. Persistent confusion between plausible textual output and internal cognitive processes, or the use of mystifying language like ‘learning’ and ‘hallucination’, configure public expectations around what kinds of politics and ethics of genAI are reasonable or plausible. I adapt the notion of *abductive agency*, originally developed by the anthropologist Alfred Gell, to explain how such misrecognition strategically defines the terms of the AI conversation.

I further argue that such *strategic misrecognition* is not new or accidental, but a central tradition in the social history of computing and artificial intelligence. This tradition runs through the originary deception of the Turing Test, famously never intended as a rigorous test of artificial intelligence, to the present array of drama and public spectacle in the form of competitions, demonstrations and product launches. The primary impact of this tradition is not to progressively clarify the nature of machine intelligence, but to constantly redefine values like intelligence in order to legitimise and mythologise our newest machines – and their increasingly wealthy and powerful owners.

Keywords: generative AI; machine intelligence; agency; ritual; spectacle; history of AI

## 1.1 The real artificial intelligence

Public discourse around generative AI (genAI) is saturated with what we might call the realness question: is the software ‘really’ intelligent, creative, or even, sentient? Does its textual, image, and video output count as ‘really’ artistic or meaningful? The quasi-public releases of Stable Diffusion, Midjourney and DALL-E 2 during the summer of 2022, and then ChatGPT’s sudden launch that November, provoked a cacophony of public spectacles, often explicitly presented as tests of realness. A Midjourney-generated image won an art prize at a local fair (Roose, 2022); GPT-4 allegedly passed bar exams with flying colours (Weiss, 2023); and a United Nations event staged a ‘press conference’ by humanoid robots (AP News, 2023). Yet the realness question does not present any kind of coherent philosophical inquiry; there are rarely clear and consistent criteria for distinguishing genuine intelligence or real creativity from its

facsimiles. Rather, the realness question provides a social ritual for dramatising and renegotiating the boundaries between ourselves and machines – a ritual that can be traced, some cultural histories of automata suggest, back to Western Antiquity (Kang, 2011). In this role, the realness question mediates other dominant public debates around genAI's capabilities and consequences, such as machine intelligence, labour replacement, and 'hallucination'.

These rituals for understanding and evaluating artificial intelligence play out through public dramas. From entrepreneurs' bombastic, self-serving predictions about imminent AI breakthroughs, to enthusiastic users seeking to test genAI tools for signs of intelligence, such spectacles help shape what becomes widely understood as possible, plausible, and inevitable about new technologies. Distinct from the substantive question of whether a given machine is 'really intelligent', public dramas elicit visceral, affective relationships with novel technologies, moulding the sentimental conditions upon which specific debates around regulation or adoption take place. This tendency extends a long historical pattern of how science and technology become publicly understood and justified. Akin to the demonstrative, even theatrical role of experiments in early modern science (Shapin & Schaeffer, 1985) to modern America's fascination with technological 'sublime' (Nye, 1994), public tests of genAI's 'realness' cultivate a collective sense of what proofs or breakthroughs might look and feel like.

Crucially, this dramatisation of genAI repeatedly reinforces a fundamental misrecognition of the technology and its relationship to human qualities like intelligence and creativity. Drawing on Peircian semiotics, Weatherby and Justie (2022) argue that the very question of whether AI 'resembles' human intelligence is a 'metaphysical misstep' based on a 'naïve iconic interpretation' – where visceral resemblance at the output level, such as in LLM-generated news articles or poetry, is taken as proof that machine learning systems process information as human brains do. This misrecognition is not simply an epistemic error, attributable to technological complexity and thus resolvable through literacy or explainability. Rather, it is a strategic misrecognition that calibrates collective expectations around how AI is to be evaluated, and what kinds of political and ethical questions are relevant or plausible – a process that is frequently led by and for a narrow elite of entrepreneurs, venture capitalists, and celebrity scientists. Generative AI output and its various misadventures thus populate public debates with anecdotal folklore, media framing, and other tropes toward the broader myth that genAI is on the cusp of an imminent and inevitable transformation of society (e.g. Bender, 2022; Sadowski, 2018; Vinsel, 2021). Below, I make two arguments: first, that public dramas around genAI's realness question are a central medium for perpetuating this strategic misrecognition; second, that such misrecognition establishes a common set of tests, thresholds, and indicators for evaluating AI and its promises – providing a tautological cycle of justification for the tech elites.

This approach helps clarify the strategic function of the speculative hype and panic around successive generations of technological futures, including generative AI. Much of the current debate around AI's social impact has been structured around the apparent opposition between 'accelerationist' and 'doomerist' positions, and the ostensible feud between the two. Thus Marc Andreessen (2023), one of Silicon Valley's most influential venture capitalists, insists that AI will "save the world" through a virtuous cycle of productivity growth, trickle-down economics and 'skyrocketing' consumer spending. Meanwhile, the 'doomers' issue high-profile calls to pause AI development, such as the longtermist Future of Life Institute's March 2023 letter, featuring star signatories like the computer scientist Yoshua Bengio and Elon Musk (see Gebru et al., 2023). Yet both positions effectively share a foundational commitment to the realness question. Both camps share the foundational assumption that existing genAI systems provide a direct technological ramp to 'really' intelligent machines, i.e. artificial general intelligence [AGI]. It is upon this speculative common ground that both 'camps' argue AGI is imminent and inevitable, and that the incumbent elite of technical AI researchers and industry leaders must continue building these tools at their discretion.

The persistent misrecognition of genAI as 'sparks of AGI' thus supplies pseudo-metaphysical justification for further concentration of power and wealth around AI. These narratives help consolidate

policy debates and investment funding around speculative logics of superintelligence and far future risks (McKelvey & Roberge, 2023, p. 21). Thus Geoffrey Hinton, whose scientific contribution to backpropagation has been central to contemporary machine learning breakthroughs, has since rebranded himself as a public speculator around the existential dangers of AI (e.g. Metz, 2023). He has argued that concerns about AI's amplifying racial disparities or environmental destruction "aren't as existentially serious" and that we should rather focus on the robot apocalypse (Bender, 2023). This dynamic extends what Meredith Whittaker (2021) has called the capture of AI ethics by industry actors – in which Silicon Valley in particular has used its political and economic capital to direct resources and visibility towards incremental, techno-solutionist renditions of AI ethics, squeezing out more foundational problems around facial recognition and algorithmic decision-making systems. The notion of strategic misrecognition emphasises how accelerationist and doomer predictions, and other public dramas around the realness question, work to "[re]capture power and create obfuscation by making new myths and legends" (Coldicutt, 2023) around speculative future technologies and their incumbent elites.

## 1.2 Dramas of justification

Public dramas around generative AI echo a far longer pattern of metaphysical and teleological myths in the history of technology. In the early United States, technology was quickly enlisted toward national origin stories of 'second creation' through settler mastery of nature (Nye, 2003). Post-colonial nation-building often mobilises grand technological projects – dams, telecom infrastructure, and so on – into belief-building objects that Paula Chakravartty (2004) calls 'temples of the future'. Silicon Valley's decades-long fascination with posthumanism and the Singularity extends this mythological tapestry and especially its eschatological dimension (Dinerstein, 2006; Hong, 2022; Kneese, 2015; F. Turner, 2016). Speculative discourse around human extinction and radical transcendence, across technologies like AGI and space exploration (Roberts & Hogan, 2019; Smith & Burrows, 2021; Utrata, 2023), carve out bounded spaces for rehearsing the social and political implications of technology.

These are what Bryan Pfaffenberger (1992) has called technological dramas: public events and controversies through which technological change is collectively visualised, alternative interpretations confronted, providing a scaffolding of expectations and attitudes around speculative notions like artificial general intelligence. Technological dramas are strongly ritualistic in the sense that these collective, staged performances work to enact new social realities (e.g. Rappaport, 1999; V. Turner, 1976). It is the formalised roles and scripts in a ritual, and its public and staged quality, that produces a collective sense that the change in question (say, the reclassification of a neophyte into a full member of the community) is being appropriately tested and subsequently resolved. Today, the means of production and propagation of these dramas is largely dominated by the incumbent group of tech megacorporations and investors, from Microsoft through their relationship with OpenAI to figures like Andreessen.

Public dramas around genAI reinforce the primacy of the realness question as an interpretive lens and historical narrative. The design and promotion of tools like ChatGPT and Bard has encouraged the general public to experiment, 'testing' the machines for the appearance of intelligence or artistry. Casey Fiesler notes that ChatGPT's breakout success owed much to its anthromorphising interface – a deliberate design choice and not driven by any technical necessity – which has the cost of obfuscating the models, data, and human labour behind the front-end (in Bender & Fiesler, 2023). Investment and product development decisions at the corporate level are also driven by the need to generate and steer public dramas in order to bootstrap underdeveloped technologies into cultural and market prominence. David Kampmann (2024) has shown how Silicon Valley's reliance on venture capital funding produces hard economic motivations for such dramatisation: since investment is rooted in expectations of massive 'blitzscaling' success for a few breakout products, spectacles of supposedly imminent breakthroughs become a key collateral through which to gain further funding and contracts. OpenAI's development of ChatGPT in late 2022 was itself frantically rushed (Roose, 2023) specifically to pre-empt competitors in shaping public expectations

around genAI. This speed was rewarded with extended media and policy attention, which OpenAI's leaders maximised with a months-long world tour meeting state leaders, culminating in an early 2024 share sale valuing the company at 80 billion USD. Public dramas, in other words, are key to how nascent and unreliable technologies become embedded into cultural and financial futures.

The realness question thus employs dramatic means (around the misrecognition of 'real' machine intelligence) towards strategic ends (of shaping AI's regimes of justification). Below, I analyse genres of drama around genAI as intelligent machines, and situate them in a longer history of computing and AI – eventually arriving in the long, zombified afterlife of the Turing Test. These dramas around genAI do not constitute a collective, open-ended exploration for the meaning of new technologies, but an elite-driven pursuit of cultural justification for already determined strategies of market domination. In *The New Spirit of Capitalism*, Luc Boltanski and Eve Chiapello (2007) explain that one of capitalism's distinct strengths is at crafting its own regimes of justification: the publicly accepted 'tests' and standards by which new paradigms like neoliberal privatisation or AI-driven automation might be considered proven. For instance, there is mounting evidence that the use of predictive algorithms for human behaviour and social decision-making is beset by fundamental flaws: the target of prediction, such as criminal risk, cannot be robustly defined and operationalised, and the very conceit of 'predicting' social outcomes lacks a coherent theoretical foundation (see overview in Wang et al., 2023). Yet the use of single-figure accuracy scores, for example, has effectively positioned these tools as engaged in a race with the 'average human', yielding a contest in which victory provides an objective seal of legitimacy. This is one example of the justificatory regime built through the incessant testing of generative AI for 'real' intelligence: a collective framework against which other theories and values must make themselves legible – or become relegated to the margins.

This dramatic approach allows us to focus on the concrete societal consequences of genAI, which are driven not by a metaphysical ascent up the ladder of general intelligence, but by the ambiguous, arbitrary, and underdefined ways in which we are encouraged to talk and think about the technology. Researchers have painstakingly shown how present-day generative AI does not meaningfully understand language (Bender, 2023); are stochastic parrots that recombine text and image without meaning (Gebu et al., 2023; Jiang et al., 2023), and exhibits no direct pathway to developing 'general' intelligence (Marcus, 2022) – despite the very label of 'artificial intelligence' imbuing contemporary machine learning techniques with 'magical thinking' (Raley & Rhee, 2023, p. 188) around its capabilities and historical trajectory.<sup>1</sup> This definitional messiness is a feature, not a bug, of the drama. Tech industry and its elites have proven consistently skilled at deflecting scrutiny from regulators, academia and civil society, often by encouraging debates and tests that are predisposed towards incremental, technical solutions: is there a 'tradeoff' between security and privacy in digital systems? Could a more 'optimised' algorithm avoid racial bias and other historical inequalities? (e.g. see critiques in Goldenfein et al., 2020; McKelvey & Neves, 2021; Turow et al., 2015) Accounts of strategic misrecognition can help us think more deliberately and independently about the regimes of justification around generative AI.

Below, I focus on the misrecognition of intelligence in genAI discourse. The field of artificial intelligence has famously lacked a clear and consensual definition of intelligence for its entire history. In public dramas around genAI and the future of AI technology, intelligence appears both as a master criterion for progress and as a messy, inconsistent notion overlapping with qualities like consciousness and creativity. Public dramas around genAI rely on such overlaps: debates around specific capabilities (such as writing generic news articles) or other humanlike qualities (such as performing empathy) quickly become scenes for pseudo-metaphysical speculation around machine intelligence. This allows actors to strategically gesture towards wider cosmological speculations, seeking to endow today's clunky machines with an aura of historical agency. I return below to the history of artificial intelligence as a

<sup>1</sup> Invoking philosopher of science Helen Verran, Lucy Suchman calls it a 'hardening of the categories' – where the fixity of the sign exudes a semiotic confidence belied by the instability of the classificatory work now concealed by the sign (Suchman, 2023, p. 3).

strategically underdefined term. But first: what exactly is being ‘misrecognised’ when genAI is subject to tests of ‘real’ machine intelligence?

## 2.1 The abduction of generative AI

In 1997, the British anthropologist Alfred Gell received a terminal cancer diagnosis – and duly resolved to complete as quickly as possible his final argument: that art objects exhibit agency through a process of *abduction* (Gell, 1998). In this view, many kinds of art – a painted portrait, a finely carved Trobriand canoe, apotropaic emblems like good-luck charms – may all exhibit agency insofar as they are *attributed* causal power in their cultural context. Drawing on Peircian semiotics, Gell argued that art objects are indexes, “chiefly function[ing] as a sign that points to something else” (Chua & Elliott, 2013, p. 8). The objects do not possess agency, but are imputed agentic capacity by others. Thus agency is a relational *effect*, and not an essential *quality*; it is something that “rub[s] off on objects” rather than emanating from within (Abram, 2005, p. 13).<sup>2</sup>

Agency, of course, is not intelligence. But genAI’s misrecognition of intelligence involves a very similar form of abductive attribution. Software *output* is consistently misrecognised as *process*: the appearance of *bildungsroman* form or impressionist technique in the output is taken as proof of some form of underlying skill, intention, or understanding – despite growing evidence of basic methodological flaws behind many high-profile demonstrations that sought to prove LLMs can ‘know’ and ‘think’ (Martínez, 2024; also see Heaven, 2023b).<sup>3</sup> Notably, Gell emphasises that abductive agency does not require free will or consciousness, only a social object with a basic capacity to mediate causal effects. Decoupling agency from metaphysics in this way unclutters much of the category confusion in the realness question. Since agency is abducted, i.e. socially attributed as relational effect rather than an internally possessed quality, we can recognise genAI output’s obvious impact on human beings – to bewilder, to convince, to affect – without needing to jump to logically incoherent claims about machine intelligence. We also need not presume that abductive activity is always a sign of user ignorance or superstition, if we realise that abduction is a common and ordinary form of human sensemaking.

In 2022, Google engineer Blake Lemoine went public with claims that, based on his experience with the company’s LaMDA chatbot, the chatbot must be sentient. The extended global media attention on the drama – heightened by Google’s subsequent termination of Lemoine – further reinforced the dominance of the realness question in public debates around generative AI (e.g. see Johnson, 2022). Here, output is again misrecognised as process, and relational phenomena used to speculate about essential qualities. The Lemoine saga extends a long pattern of such abductive panics in the history of computing. Chatbots, in particular, have been dominated by attributions of humanlike intelligence and consciousness for decades. In 1966, the public debut of the ELIZA chatbot famously provoked similar misrecognition around machine intelligence. Joseph Weizenbaum, its creator, suggested that his own secretary imputed human qualities onto the bot – despite the immediately visible limitations of its Sokratic tendency to repeat the user’s questions back at them (Weizenbaum, 1976, p. 6). The irony was that Weizenbaum had designed ELIZA in part to show how relatively unintelligent software with a small vocabulary could still simulate the appearance of conversational interaction.

Such misrecognitions are often presented as a true/false test: is the computer really intelligent (that is, possessing it as an essential and innate quality), or is the user simply mistaken? Abduction provides an alternative interpretation: that such attributions are pragmatic responses that help establish *useful* fixtures for social and moral responses to new technology, even if they may be technically incorrect. Blaming

<sup>2</sup> More broadly, relational approaches have been central to accounts of information processing and decision-making in complex human-machine assemblages, from drone warfare to biohacking and smart urbanisms, where human perception, cognition and judgment are thoroughly interwoven with the machines around them (Hayles, 2017; Sampson, 2017).

<sup>3</sup> Consider OpenAI’s widely publicised claim that GPT-4 scored in the 90<sup>th</sup> percentile in the Uniform Bar Exam. An independent evaluation by Martínez (2024) has shown that the percentile figure was estimated using unspecified methods, which appear to draw from a skewed dataset of human scores.

algorithms for bias, for example, can be a (temporarily) effective way to seek some accountability for a harmful result where traditional attributions of blame have been obfuscated by new decision-making systems. For tech elites, misrecognition has been a profitable way to stoke controversy around the possible sentience of AI systems, persuading investors, policymakers and the public that AI development must be a grand civilisational priority above all other concerns.

Gell had insisted that a proper anthropology of art needed to be strictly divorced from the question of aesthetics. There were two reasons for such a move. First, the causal effects attributed to art objects were not always dependent on aesthetic appreciation. Second, this would help us consider abduction in objects that don't easily fit into modern Western norms around art – the primacy of painting, the hierarchy of artistic value via museums and curatorial control, the norms of contemplative appreciation, and so on. Similarly, setting aside aesthetic considerations helps clarify what is happening in an abductive process: genAI output's ability to fool or amaze human observers is not any necessary indicator of intelligence, creativity or consciousness, but only of its capacity for inviting abduction. Such capacity emerges from, above all, genAI systems' ability to ingest massive amounts of human art and communication – and the webs of referentiality, intention, and meaning invested into them – and to break down and reprocess it all into an industrial swill for mass production (Salvaggio, 2023). It is from this cacophony of fragments that we infer meaning, relating a meaningless image or text back into our learned sense of art history of literary genre. As scholars have already noted of 'big data' systems, when there is enough noise, it becomes easy enough to 'discover' signal apophenically (e.g. Steyerl, 2016). This distinction puts a stopper on the slippery slope of reasoning, where it is assumed *if* genAI were to produce 'real' art (or legal opinion, or scientific paper), then the machine would need to be treated as an *artist* – and the ensuing speculations about machines' ability to replace human labour wholesale or their eligibility for various human rights. As Gell (1988, chapter 7) pointed out, many objects across world cultures are abductively attributed with powerful causal effects without conferring virtuosic status on their creators.

Thus, if a genAI-produced poem should unsettle us or bring us to tears, this says nothing much about essential humanlike qualities, but points to a long tradition of abductively artistic objects that have made humans laugh and cry without necessarily moving the bar on sentience or existential risk. Automata, in particular, have historically incited public engagement specifically in the form of 'detective work', provoking speculations of realness: does this machine reason? Does it simulate something of life? And if so, what does it suggest about our own reasoning and living? (Geoghegan, 2020) One emblematic case is Jacques de Vaucanson's 18<sup>th</sup> century invention, *Le Canard Digérateur* [The Digesting Duck]. Vaucanson claimed that the machine had mastered the ability to eat, and subsequently, shit – and showed off the routine in popular demonstrations in and around 18<sup>th</sup> century Paris. All available evidence suggests that the duck had simply stored pre-made shit pellets inside; but what is telling is that well into the 19<sup>th</sup> century, commentators often assumed that it must have possessed artificial anatomical features, such as a mechanical intestine (Riskin, 2003). Machines have been objects of abduction for a long time, and misrecognition of this abductive relation has long fuelled speculation around what makes *us* intelligent.

## 2.2 Uncanny capture

Why, though, does genAI output prove so fertile for abductive attribution? Let us remain a little longer with Gell, who provides the label of the causal milieu (1998, p. 37): a cultural, interpretive context that shapes what kinds of agency is likely to be anticipated and recognised. From sorcery and cursework to carnivalesque effigies, the milieu regulates cultural norms around how some forms of abductive attribution become acceptable, while others become categorised as superstitious or transgressive. Gell himself cites the case of the *Rokeby Venus* in 1914 (Gell, 1998, Chapter 4), in which suffragette Mary Richardson's slashing a venerable Velázquez is condemned as a barbaric attack rather than being seen as an 'artistic' experiment of defacement. The difference between what constitutes 'genuine' acts of creativity and skill, and what constitutes mere copy or even vandalism, depends on this quasi-

Bourdieuian milieu. In their very design as well as the promotional discourses around them, genAI tools like Midjourney seek to shape a milieu for themselves, offering statistical recombination of image data as the new benchmark of intelligence.

Experiences of mystery, awe, and captivation are central to shaping these causal milieus. Here we return to technology as myth, as temples of the future. Gell suggests that abduction of agency often occurs when we perceive ordinary means achieving extraordinary results. The brushwork on a Vermeer or the carvings on a Trobriand canoe is not so alien as to defy *any* interpretation, but sufficiently extraordinary to evoke a sense of mystery as well – especially if we’ve been exposed to modern art museums and other causal milieus in our lives (Gell, 1998, p. 71). GenAI output, especially when designed to respond to iterative and experimental human prompts, incites similar activities of speculation and wonder through the aforementioned ‘detective work’. Such abductive activity is further encouraged by the tech elite-driven narrative that AI is fundamentally ‘unexplainable’ even to its own human creators – an air of mystique that is then used to justify both utopian and dystopian claims about the technology’s imminently superintelligent status. Microsoft researchers’ much-publicised paper, ‘Sparks of AGI’ (Bubeck et al., 2023), thus argued that GPT-4 exhibits some forms of ‘general intelligence’, even while employing definitions and measures of intelligence that they acknowledged were rather arbitrary and subjective – and, as other researchers quickly pointed out, relied on an infamous 1997 editorial that was written as an activist endorsement of innate racial differences in human intelligence (Bender, 2023).

The mystery around genAI output thus provokes a sense of the *uncanny*, furnishing space for speculative meaning-making and abductive attribution. In computing, the uncanny usually means the uncanny valley: the zone in which synthetic images are realistic enough to invite recognition *as* real, but not quite realistic enough to convince. Yet both in Masahiro Mori’s original 1970 formulation and contemporary use, the uncanny valley functions as a useful label for clearly observable phenomena, but there is no robust theory of what makes computer graphics uncanny under what conditions. More broadly, the uncanny has a rich tradition in philosophy and literary theory, and has often been central to cultural responses to automation throughout history. The surrealists turned to cyborgs and other uncanny intersections of humans and machines to process the ‘shock of industrialisation’ (Foster, 1991) – as the situationists would later do in response to the dominant myth of cybernetic futures (Routhier, 2023). Many moral panics around new media technologies can also be understood as a reaction to the perceived unsettling of boundaries and categories (also see Ravetto-Biagioli, 2019; Kang, 2011). We may think of the ambiguous gap between my ‘offline’ and ‘online’ self, and later, my ‘data double’ or ‘data doppelganger’ (e.g. Cheney-Lippold, 2017; Watson, 2014), or between bots and ‘real’ users – and the discomfort that arises from the eventual recognition that there can often be no firm separation between the two.

The uncanny is thus an important site for inciting wonder and terror around genAI objects, and in turn, practices of abductive attribution. Susan Lepselter (2016) writes that the uncanny is a central affect in many everyday myths about politics and society. It is a percolating sense, never quite concretised into official forms, but recurring in the margins, in the patterns of the stories we tell: for instance, “a sense that life in America [today] is shaped by some ineffable, enormous power, a power that can be seen only in the patterns of its effects” (p. 1) – from relatively discredited powers like aliens and the Illuminati, to relatively normalised ones like ‘the economy’ or ‘the (deep) state’. The key is that the uncanny is not simply an experience of the destruction of the scaffolding of meaning, but also part of a constant, everyday effort to rebuild that scaffolding. Drawing on Freud’s *unheimlich* – usually read as ‘unhomely’ or ‘unfamiliar’, but also, in his own writing, involving forms of ‘involuntary repetition’ – Lepselter (pp. 22–23) argues that there is a habitual element to the uncanny: that we scramble to *unearth* (also see Vidler, 1999, p. 27) some uncanny gap, and then to see it again and again, so that we may once again seek to close or ‘complete’ it. This provocative, restless dimension comes to the fore in the abductive speculations around generative AI.

In September 2022, multimedia artist Supercomposite reported ‘discovering’ a memorably creepy, deformed face while experimenting with negative prompt weights in genAI tools (in which the software seeks to generate output that is *opposite* to the description) (e.g. Rose, 2022). It was an emergent effect: since this particular face was not specifically summoned by a positive prompt or name, the human user could only indirectly control its appearance in generated output. But reappear it did across many different generated images, with a consistently macabre stylisation. (It is apropos that the face often takes a doll-like aesthetic: Freud examined dolls as a key example of *unheimlich*, drawing from Ernst Jentsch’s earlier reading of the doll Olympia in Romantic writer E.T.A. Hoffmann’s 1817 story, *Der Sandmann*.) Supercomposite promptly named the face ‘Loab’ – observing that even when Loab-producing prompts were combined with prompts for entirely different shapes or styles, human viewers could easily recognise an enduring ‘Loabness’ in the resulting output. The mystery around Loab’s provenance, and its memorable visual features, provoked familiar speculations and panic around genAI’s intelligence, memory, and even spirituality. Once again, the realness question had provided a shared blueprint for collective practices of speculation around machine intelligence.

### 3.1 The spectacle of the test

How are causal milieus formed and maintained for generative AI, cultivating a wider public ‘literacy’ on how to read the technology for ostensible signs of intelligence? One key medium is the *test* as spectacle. From art competitions to bar exams, the realness question (and its uncanny affects) is often operationalised into public tests, presenting putatively objective reference points for tracking the progress of artificial intelligence. Predictions of the imminent arrival of AGI thus faithfully recycle an old tradition of spectacles, tests, and predictions. In February 2022, Sam Altman opined that “techno-optimism is the only good solution to our current problems” – and that this supposedly ‘radical’ and unpopular position will deliver ‘limitless, cheap, clean energy’, as well as space colonisation, fusion, and ‘cure all human disease’ (Sam Altman [@sama], 2022). He has continued to leverage public attention around OpenAI’s generative tools to spread the teleological rhetoric of abundance. Such speculations are not lackadaisical asides, but a central part of the emerging regime of justification around AI. The tech elite’s more ‘pessimistic’ public declarations have tended to exactly mirror Altman’s rhetoric, with only the outcomes inverted: a Center for AI Safety statement in June 2023 argued that “mitigating the risk of extinction from AI should be a global priority” precisely because its superhuman prowess and universal adoption is virtually guaranteed (*Statement on AI Risk*, 2023; also see Wong, 2023). It is no contradiction, in that sense, that Altman was one of the signatories for the letter: AI optimism and pessimism are joined at the hip of strategic interest, protecting a tightly bounded sphere of technological futures.

Such predictions liberally remix religious tropes out of context to craft eschatological messages of urgency and inevitability (also see Cheney-Lippold, 2024; Gebru & Torres, 2024). Thus Anthony Levandowski, one of the most prominent engineers in the self-driving car research community, is also the founder of the short-lived “Way of the Future,” a self-proclaimed “church of AI” – founded in 2015 and shuttered in 2020 – having sought to prepare humanity for the imminent arrival of superintelligent machines. Ray Kurzweil, a serial prophet, has now spent several decades predicting the Singularity as the impending next great chapter of civilisation – a vision of total technological rewriting of society. Implicit in these tales is the conceit that only the Valley’s unique community of transgressive geniuses can see into this future and prepare the rest of us for its arrival (also see Hong, 2022).

Predicting some form of ‘genuine’, all-purpose artificial intelligence in the near future, quietly stepping over the expiry date, and then predicting it again, has been a regular ritual throughout the history of AI research. In 1970, Marvin Minsky, widely hailed as a founding figure of AI, told *Life* magazine that a machine will possess ‘general intelligence of an average human being’ in ‘three to eight years’ (Darrach, 1970) – despite the field’s perennial lack of a robust, consensual definition of intelligence throughout its entire existence (e.g. Dick, 2019). In many ways, “the field [of AI] has always involved fantasy in search



of a practical method” – a conjectural bet, made for promotional reasons as well as intellectual, that machine intelligence might be possible (Stark, 2023a, p. 367). Minsky and his colleagues specifically chose the term artificial intelligence to cast the broadest possible net for funding and significance. Yarden Katz describes how John McCarthy, in proposing the 1956 Dartmouth workshop today mythologised as the founding event for AI, eventually settled on the term in order to avoid “embracing MIT mathematician Norbert Wiener, who coined [cybernetics], as the [new] field’s ‘guru’.” (Katz, 2020, p. 22) Although Claude Shannon had initially argued against the undefined and loaded use of ‘intelligence’, he eventually signed on, along with Minsky, as the workshop’s co-organiser. The new moniker quickly achieved its aims: ARPA enthusiastically funded artificial intelligence projects, sidelining cyberneticians despite their complaints of the AI researchers’ ‘con trick’ (Katz, 2020, pp. 24–25).

This pattern of speculation and prediction is not restricted to a discrete ‘marketing’ dimension, hermetically sealed away from the more serious work of actually doing science and technology. Communities of AI and computing have for decades organised their collective sense of technological progress and its direction through genres of public tests and spectacle. One obvious correlate here are performance competitions, such as ImageNet’s Large Scale Visual Recognition Challenge (ILSVRC) or the Netflix Prize, which organise a formal conflict through which specific models may emerge victorious, and more importantly, a common sense of technical progress may be coordinated. Such tests extend a wider technoscientific culture of justification, from benchmarks (Raji et al., 2021) to prototypes (Corsín Jiménez, 2014; F. Turner, 2016) and variants of design thinking (Irani, 2018). Similarly, a ‘test-bed’ approach is central to the smart technology paradigm: populations and cities become sites of constant, real-time adjustments and experiments in which models are always already rolling out changes into the environment and thus requiring indefinite adaptation (Halpern & Mitchell, 2022). In short, the very notion of artificial intelligence emerged in part as a bid to shape the public conditions of justification. The regular production of speculative predictions about transformative futures, and more formalised rituals of tests, challenges and prototypes, both serve as what Latour (1988) called the ‘theatre of the proof’, cultivating common expectations around what can realistically be demanded of technology in the first place.

### 3.2 Turing’s mirage

One of AI’s most popular origin myths, of course, is a test. In 1950, Alan Turing proposed ‘the imitation game’. Its details are well known: X converses with Y and Z, physically separated and without somatic markers like voice. If X cannot distinguish which of Y and Z is human or machine, the machine has passed the test. This fairly specific exercise was then glossed with a rather speculative estimate: that ‘in about fifty years’ time’, computers would have sufficient storage capacity to pass the test (Turing, 1950).

Crucially, Turing did *not* address the question, ‘can machines think?’, which he explicitly dismissed as “too meaningless to deserve discussion” (p. 442). The imitation game was instead about *substituting* the problem of intelligence with a pragmatic and provisional task. In the 1950 paper, Turing entertained some familiar objections – for example, that machines cannot produce novel thought, or fall in love – but the aim was not to defeat these objections, only to suggest that these objections themselves are not conclusive enough to dismiss the possibility of machine intelligence in the future (pp. 442–454). Simply put, the Turing Test was never a *test*; in Turing’s own words, it remained a game, a demonstration (also see Mitchell, 2024).

The central role of *imitation* in the Turing Test sabotages any retroactive effort to formalise the test as a watertight proof of intelligence, as many have discovered over the decades. After all, imitation points to computers as fundamentally abductive agents, rather than possessing any essentialised agency. Yet the imitative dimension was not accidental; it was rooted in Turing’s lifetime of work. Well before the 1950 paper, the Turing Machine had proposed a universal mathematical definition of computability, and its iconic tape mechanism was designed to mimic human mathematical cognition (Hodges, 2009). One might speculate an alternative history of AI where the field seriously engaged the relational implications of

imitation in the Turing Test – leading, perhaps, to more realistic theories of machine interactivity unburdened of the metaphysical hot air around the sentience of text and pixel recombination software.

Computing took the other road. Since 1950, the Turing Test has taken on a mythological life across both the lay public and expert communities, often taken as an ur-test, a final frontier which offers the possibility of absolutely proving machine intelligence. Although many researchers across philosophy, computer science and other domains have dissected the test's limitations, it has remained prominent both as a thought experiment and as a concrete research objective. Some competitions, such as the Loebner Prize (1990-2020), took a fairly literal path, each year inviting participants in a standardised Turing Test. Some observers were enthusiastic, others more skeptical (e.g. Floridi et al., 2009; Marcus, 2014) – while, somewhat ironically given his own historical promotion of utopian predictions about AI, Marvin Minsky loudly decried the Prize as a publicity stunt. James Moor, a prominent moral philosopher in computing and AI who wrote extensively on the Test for over four decades, admitted that the test fails to provide an operational definition of AI – while also insisting that the test can somehow provide a philosophical foundation for the question of machine intelligence (Moor, 2001).

This is not to paint a history in which original nuance is diluted over time. Simone Natale (2021) argues that since its inception, the Turing Test helped establish deception – or, in my terms here, misrecognition – as a constitutive element of artificial intelligence. Turing had readily admitted that the 1950 piece was written as a “form of propaganda that would stimulate philosophers, mathematicians, and scientists to engage more seriously with the machine intelligence question.” (Natale, 2021, pp. 19–20) This ‘deceptive’ strategy determined the test's initial design as well as its subsequent legacy. Consider the actual parameters of the test. Turing's design requires an ‘average interrogator’ without special knowledge (of, say, linguistics or anthropology), without additional contextual knowledge of how machines of this kind tend to be designed, and tasked with guessing the machine after ‘five minutes of questioning’. There is little theoretical justification as to why any of these parameters are necessarily or ideal. Instead, Turing chose to present a kind of test which can provide an effective and immediate theatre of proof.<sup>4</sup>

The legacy of the Turing Test, too, belongs primarily in the annals of technoscientific demonstrations towards constructing public regimes of justification. Researchers have suggested that competitions like the Netflix Prize taken on a sport-like dimension in terms of its coverage and cultural impact (e.g. Orr & Kang, 2023); we may also point to more literal sport events staged for AI justification, such as AlphaGo's duel against Go master Lee Sedol in 2016. As computer scientist Delip Rao (2023) has argued, these benchmarks and competitions also exert a gravitational pull for research and funding, creating perverse incentives in what kinds of ‘breakthroughs’ are pursued and celebrated.

There has been no shortage of candidates for dethroning the Turing Test with other tests. But most efforts play with the parameters while retaining its dramatic structure. Some suggest a Lovelace Test, for instance. Taking up Ada Lovelace's comment on Babbage's Analytical Engine that computers cannot ‘originate’, it is argued that a test for more creative content, such as poetry, would provide a satisfactory test of intelligence (e.g. Bringsjord et al., 2003). Yet the misrecognition of an abductive relation as an essential one remains. The search for the ‘right’ or better test thus extends rather than breaks from this wider culture of ritualised, spectacular demonstrations. Leif Weatherby (2023) relays Lydia Liu's observation that “all ‘Turing-style’ tests reduce intelligence to a competition between humans and machines” – a sport in which the winner takes home the prized status of public proof and recognition.

What might a qualitatively different test do for public conversations around generative AI? Jack Stilgoe (2023) argues for a Weizenbaum Test: if a technology *can* be built, *should* it? Who will benefit, and who will bear the costs? Such a test steps outside the myopic laboratory conditions imposed by the Turing Test and its heirs, in which an individual human specimen is locked inside a (Chinese) room for experimental proof. Outside this strange and unnecessary confinement, it becomes clear that where

<sup>4</sup> Turing recommended that when given an arithmetic question in an imitation game, the computer should “pause about 30 seconds” before giving the answer (Turing, 1950, p. 434) – a parlour trick quite comparable to how present-day genAI systems take to folksy turns of phrase to perform humanlikeness.

computers are attributed by (some) human beings with qualities like intelligence or consciousness, that practice of attribution is itself shaped by arcs of political and economic interest that have cleared the social stage for the arrival of the new technology in the first place (Marvin, 1988; Pfaffenberger, 1992).

#### 4. Conclusion: Devaluation machines

The misrecognition surrounding generative AI is ritually reproduced abduction, rather than incidental forms of human ignorance. It serves strategic objectives for shaping the terms of the AI conversation in favour of the incumbent tech elites, rather than simply a stumbling block on the road to sentient machines. GenAI, in other words, should be evaluated primarily as an abduction machine with a social and political strategy. Tools like Midjourney and ChatGPT are deliberately designed to produce a particular form of abduction, derived from a long history of anthromorphisation and of a fascination with uncanny machines.

The most serious consequence of such abduction – and one in urgent need of critical analysis and reframing today – is that these machines enact a widespread *devaluation* of the human qualities that the technology purports to reproduce. The question of whether a machine can ‘replace’ a human at a task is a misleading one, because it reproduces the realness question as a substantive problem, and implies that the answer depends on how ‘smart’ the machine is. Rather, the abduction view asks how machines dramatise and redefine qualities like intelligence – and in doing so, justify the pernicious suggestion that whatever can cheaply and fleetingly deceive as intelligence, or creativity, or empathy, must be just as good as the real thing.

Consider the use of genAI to undermine the social valuation of creative and artistic labour. Luke Stark turns to 20<sup>th</sup> century animation as a parallel shift in this regard. Animation, too, was a technological transformation of image production strongly predicated on the abductive function. “All animation simplifies, and so is implicitly dependent in the human ability to make meaningful heuristic inference” (2023b). In this view, animation and generative AI become effective by *provoking* abductive attribution. Yet this is only half of the process. As the animated output becomes misrecognised as agentic, they “pus[h] the the actual living labour of their animators into the background or offstage” (ibid.). Optimising genAI’s abductive capacity involves relentlessly calculating *whatever output might provoke desired user inference at the lowest cost*. This optimisation occurs to a degree in the training phase, where models are honed to select whatever combination of textual and image data that ‘looks real enough’ to the low-paid, subcontracted human coder (e.g. Dzieza, 2023). But it also occurs in the public presentation of recombinatory bullshit as ‘real’ art or communication. The corollary to the argument that genAI output could replace human workers is the premise that nothing deeper than the rough appearance of humanlike qualities has any value.

This politics of devaluation thus corresponds to a rather retrograde essentialism that I described as a strategic misrecognition about abduction: the belief that rather than being complex and relational things, intelligence or creativity is an essence that an AI might ‘have’ or not have. This misrecognition is not simply an epistemic confusion, but a crucial instrument for the ongoing politics of justification around data-driven technologies as mechanisms of extraction (e.g. Thatcher et al., 2016). Generative AI as a market product grows by attaching itself to existing infrastructures of production and then offering a lower-cost alternative. From illustration and digital journalism to essay mills and bot farms, genAI takes a pre-existing milieu for assigning value to labour, and by decoupling output from process, promises to discard the bulk of that labour and its associated costs. Where Amazon Mechanical Turk and successive crowdwork platforms once achieved this arbitrage largely by slicing employment into smaller chunks of time in order to minimise wages and benefits (Irani, 2015; Shapiro, 2018), generative AI offers a sort of Turing Test for the precariat: that if a generated text or image is *just* good enough to pass as human-made for a non-discerning customer in the short term, then the human skill and labour for that work must not have much value at all.

The public dramas around genAI thus shape the likely path and extent of the technology's political economic impacts. Anticipating and critically evaluating those 'disruptions' require tracing how dramatised myths till the soil of public and industry expectation. What researchers have called fauxtimation or 'potemkin AI' (e.g. Sadowski, 2018; Taylor, 2018) is not a temporary stepping stone, but the most likely endpoint of an optimisation function that prioritises minimally convincing output above meaning, intention, or process. As genAI relies on abductive effects to justify and promote itself, it is likely to incentivise other economic and social decisions that strip away such 'extraneous' processes. In this sense, an image generator software does not 'replace' artists or designers; rather, it executes a scorched earth policy of devaluation, depressing the social value of that work for all parties.

In 1979, Lyotard opened *The Postmodern Condition* by noting that in computerised societies, we may thus expect a thorough exteriorisation of knowledge with respect to the "knower," at whatever point he or she may occupy in the knowledge process. The old principle that the acquisition of knowledge is indissociable from the training (*Bildung*) of minds, or even of individuals, is becoming obsolete and will become ever more so (Lyotard, 1984, p. 5; also see Beardon, 1994).

The postwar rise of computing technologies corresponded with a broader shift from reason to rationality: at the broadest strokes, a move away from liberal subjects endowed with relatively self-sufficient capabilities for reasoning and knowledge, toward impersonal and calculative systems as custodians of knowledge (Erickson et al., 2013). The very notion of an 'artificial' intelligence extends this epistemic project, and its premise that intelligence can be decoupled from human subjects. Yet generative AI and its public dramas are also deeply reliant on the enduring belief that in order to think, to know, to write, one must be a subject. This vestigial imagination is now being exploited to advance the strategic misrecognition of machine intelligence, where the appearance of thinking, knowing, and writing, provokes speculation that these machines should perhaps be treated as subjects. For Kittler, the phonograph had already inaugurated the age of "writing without a subject" (1986, p. 44); today, the spectre of the writing subject still haunts the fantasy of the autonomous machine.

## References

- Abram, S. (2005). Introduction: Science/technology as politics by other means. *Focaal—European Journal of Anthropology*, 45, 3–20.
- Andreessen, M. (2023, June 6). *Why AI Will Save the World*. Andreessen Horowitz. <https://a16z.com/ai-will-save-the-world/>
- Beardon, C. (1994). Computers, postmodernism and the culture of the artificial. *AI & Society*, 8(1), 1–16. <https://doi.org/10.1007/BF02065174>
- Bender, E. (2022). *On NYT Magazine on AI: Resist the Urge to be Impressed*. Medium. <https://medium.com/@emilymenonbender/on-nyt-magazine-on-ai-resist-the-urge-to-be-impressed-3d92fd9a0edd>
- Bender, E. M. (2023, July 6). Talking about a 'schism' is ahistorical. *Medium*. <https://medium.com/@emilymenonbender/talking-about-a-schism-is-ahistorical-3c454a77220f>
- Bender, E. M., & Fiesler, C. (2023, March 1). *The Limitations of ChatGPT with Emily M. Bender and Casey Fiesler*. The Radical AI Podcast. <https://www.radicalai.org/chatgpt-limitations>
- Boltanski, L., & Chiapello, E. (2007). *The New Spirit of Capitalism*. Verso.
- Bringsjord, S., Bello, P., & Ferrucci, D. (2003). Creativity, the Turing Test, and the (Better) Lovelace Test. In J. H. Moor (Ed.), *The Turing Test: The Elusive Standard of Artificial Intelligence* (pp. 215–239). Springer Netherlands. [https://doi.org/10.1007/978-94-010-0105-2\\_12](https://doi.org/10.1007/978-94-010-0105-2_12)
- Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., & Zhang, Y. (2023). *Sparks of Artificial General Intelligence: Early experiments with GPT-4* (arXiv:2303.12712). arXiv. <https://doi.org/10.48550/arXiv.2303.12712>
- Chakravartty, P. (2004). Telecom, National Development and the Indian State: A Postcolonial Critique. *Media, Culture & Society*, 26(2), 227–249. <https://doi.org/10.1177/016344370404117>
- Cheney-Lippold, J. (2017). *We are Data—Algorithms and the making of our digital selves*. New York University Press.
- Cheney-Lippold, J. (2024). The silicon future. *New Media & Society*, OnlineFirst. <https://doi.org/10.1177/14614448241234864>
- Chua, L., & Elliott, M. (Eds.). (2013). *Distributed objects: Meaning and mattering after Alfred Gell*. Berghahn.

- Coldicutt, R. (2023, June 1). *On Understanding Power and Technology*. Medium. <https://rachelcoldicutt.medium.com/on-understanding-power-and-technology-1345dc57a1a>
- Corsín Jiménez, A. (2014). Introduction: The prototype: More than many and less than one. *Journal of Cultural Economy*, 7(4), 381–398. <https://doi.org/10.1080/17530350.2013.858059>
- Dick, S. (2019). Artificial Intelligence. *Harvard Data Science Review*, 1(1). <https://doi.org/10.1162/99608f92.92fe150c>
- Dinerstein, J. (2006). Technology and Its Discontents: On the Verge of the Posthuman. *American Quarterly*, 58(3), 569–595. <http://dx.doi.org/10.1353/aq.2006.0056>
- Dzieza, J. (2023, June 20). *Inside the AI Factory*. The Verge. <https://www.theverge.com/features/23764584/ai-artificial-intelligence-data-notation-labor-scale-surge-remotasks-openai-chatbots>
- Erickson, P., Klein, J. L., Daston, L., Lemov, R., Sturm, T., & Gordin, M. D. (2013). *How Reason Almost Lost Its Mind*. University of Chicago Press.
- Floridi, L., Taddeo, M., & Turilli, M. (2009). Turing's Imitation Game: Still an Impossible Challenge for All Machines and Some Judges—An Evaluation of the 2008 Loebner Contest. *Minds and Machines*, 19(1), 145–150. <https://doi.org/10.1007/s11023-008-9130-6>
- Foster, H. (1991). Exquisite Corpses. *Visual Anthropology Review*, 7(1), 51–61. <https://doi.org/10.1525/var.1991.7.1.51>
- Gebru, T., Bender, E. M., McMillan-Major, A., & Mitchell, M. (2023, March 31). *Statement from the listed authors of Stochastic Parrots on the "AI pause" letter*. Distributed AI Research Institute. <https://www.dair-institute.org/blog/letter-statement-March2023/>
- Gebru, T., & Torres, É. (2024). The TESCREAL bundle: Eugenics and the promise of utopia through artificial general intelligence. *First Monday*, 29(4). <https://doi.org/10.5210/fm.v29i4.13636>
- Gell, A. (1998). *Art and Agency: An Anthropological Theory*. Clarendon Press.
- Geoghegan, B. D. (2020). Orientalism and Informatics: Alterity from the Chess-Playing Turk to Amazon's Mechanical Turk. *Ex-Position*, 43, 45.
- Goldenfein, J., Green, B., & Viljoen, S. (2020, April 17). *Privacy Versus Health Is a False Trade-Off*. Jacobin. <https://jacobinmag.com/2020/4/privacy-health-surveillance-coronavirus-pandemic-technology>
- Halpern, O., & Mitchell, R. (2022). *The smartness mandate*. MIT Press.
- Hayles, K. N. (2017). *Unthought: The Power of the Cognitive Nonconscious*. University of Chicago Press.
- Heaven, W. D. (2023, August 30). *AI hype is built on high test scores. Those tests are flawed*. MIT Technology Review. <https://www.technologyreview.com/2023/08/30/1078670/large-language-models-arent-people-lets-stop-testing-them-like-they-were/>
- Hodges, A. (2009). Alan Turing and the Turing Test. In R. Epstein, G. Roberts, & G. Beber (Eds.), *Parsing the Turing Test: Philosophical and Methodological Issues in the Quest for the Thinking Computer* (pp. 13–22). Springer Netherlands. [https://doi.org/10.1007/978-1-4020-6710-5\\_2](https://doi.org/10.1007/978-1-4020-6710-5_2)
- Hong, S. (2022). Predictions Without Futures. *History and Theory*, 61(3), 371–390. <https://doi.org/10.1111/hith.12232>
- Irani, L. (2015). The cultural work of microwork. *New Media & Society*, 17(5), 720–739. <https://doi.org/10.1177/1461444813511926>
- Irani, L. (2018). "Design Thinking": Defending Silicon Valley at the Apex of Global Labor Hierarchies. *Catalyst: Feminism, Theory, Technoscience*, 4(1), Article 1. <https://doi.org/10.28968/cftt.v4i1.29638>
- Jiang, H. H., Brown, L., Cheng, J., Khan, M., Gupta, A., Workman, D., Hanna, A., Flowers, J., & Gebru, T. (2023). AI Art and its Impact on Artists. *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 363–374. <http://dx.doi.org/10.1145/3600211.3604681>
- Johnson, K. (2022, June 14). LaMDA and the Sentient AI Trap. *Wired*. <https://www.wired.com/story/lamda-sentient-ai-bias-google-blake-lemoine/>
- Kampmann, D. (2024). Venture capital, the fetish of artificial intelligence, and the contradictions of making intangible assets. *Economy and Society*, 0(0), 1–28. <https://doi.org/10.1080/03085147.2023.2294602>
- Kang, M. (2011). *Sublime Dreams of Living Machines: The Automaton in the European Imagination*. Harvard University Press.
- Katz, Y. (2020). *Artificial whiteness: Politics and ideology in artificial intelligence*. Columbia University Press.
- Kittler, F. A. (1986). *Gramophone, Film, Typewriter* (G. Winthrop-Young & M. Wutz, Trans.). Stanford University Press.
- Kneese, T. (2015, October). Spiritual Machines: Transhumanist Temporalities and Startup Culture. *Humanism and Its Prefixes*.
- Latour, B. (1988). The Pasteurization of France. In *Practical Synthetic Organic Chemistry: Reactions, Principles, and Techniques*. Harvard University Press.
- Lepselter, S. (2016). *The Resonance of Unseen Things: Poetics, Power, Captivity, and UFOs in the American Uncanny*. University of Michigan Press.
- Liotard, J.-F. (1984). *The Postmodern Condition: A Report on Knowledge*. University of Minnesota Press.
- Marcus, G. (2014, June 9). What Comes After the Turing Test? *The New Yorker*. <https://www.newyorker.com/tech/annals-of-technology/what-comes-after-the-turing-test>
- Marcus, G. (2022, December 25). What to Expect When You're Expecting ... GPT-4 [Substack]. *Marcus on AI*. <https://garymarcus.substack.com/p/what-to-expect-when-youre-expecting>

- Martínez, E. (2024). Re-evaluating GPT-4's bar exam performance. *Artificial Intelligence and Law*.  
<https://doi.org/10.1007/s10506-024-09396-9>
- Marvin, C. (1988). *When Old Technologies Were New: Thinking About Electric Communication in the Late Nineteenth Century*. Oxford University Press.
- McKelvey, F., & Neves, J. (2021). Introduction: Optimization and its discontents. *Review of Communication*, 21(2), 95–112.  
<https://doi.org/10.1080/15358593.2021.1936143>
- McKelvey, F., & Roberge, J. (2023). Recursive power: AI governmentality and technofutures. In *Handbook of Critical Studies of Artificial Intelligence* (pp. 21–32). Edward Elgar Publishing.
- Metz, C. (2023, May 1). 'The Godfather of A.I.' Leaves Google and Warns of Danger Ahead. *The New York Times*.  
<https://www.nytimes.com/2023/05/01/technology/ai-google-chatbot-engineer-quits-hinton.html>
- Mitchell, M. (2024). The Turing Test and our shifting conceptions of intelligence. *Science*, 385, eaqd9356.
- Moor, J. H. (2001). The Status and Future of the Turing Test. *Minds and Machines*, 11(1), 77–93.  
<https://doi.org/10.1023/A:1011218925467>
- Natale, S. (2021). *Deceitful Media: Artificial Intelligence and Social Life after the Turing Test*. Oxford University Press.
- Nye, D. (2003). *America as Second Creation*. MIT Press.
- Nye, D. (1994). *American Technological Sublime*. MIT Press.
- Orr, W., & Kang, E. (2023, November 9). *AI as a sport: On the competitive epistemologies of benchmarking*. Society for the Social Studies of Science [4S] Annual Meeting, Honolulu.
- Pfaffenberger, B. (1992). Technological Dramas. *Science, Technology & Human Values*, 17(3), 282–312.  
<https://doi.org/10.1177/016224399201700302>
- Raji, I. D., Bender, E. M., Paullada, A., Denton, E., & Hanna, A. (2021). AI and the Everything in the Whole Wide World Benchmark. *NeurIPS 2021*.
- Raley, R., & Rhee, J. (2023). Critical AI: A Field in Formation. *American Literature*, 95(2), 185–204.  
<https://doi.org/10.1215/00029831-10575021>
- Rao, D. (2023, December 20). The Great AI Weirding [Substack newsletter]. *AI Research & Strategy*.  
<https://deliprao.substack.com/p/the-great-ai-weirding>
- Rappaport, R. (1999). *Ritual and Religion in the Making of Humanity*. Cambridge University Press.
- Ravetto-Biagioli, K. (2019). *Digital Uncanny*. Oxford University Press.
- Riskin, J. (2003). The Defecating Duck, or, the Ambiguous Origins of Artificial Life. *Critical Inquiry*, 29(4), 599–633.  
<https://doi.org/10.1086/377722>
- Roberts, S. T., & Hogan, M. (2019). Left Behind: Futurist Fetishes, Prepping and the Abandonment of Earth. *B2o: An Online Journal*, 4(2).
- Roose, K. (2022, September 2). An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy. *The New York Times*.  
<https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html>
- Roose, K. (2023, February 3). How ChatGPT Kicked Off an A.I. Arms Race. *The New York Times*.  
<https://www.nytimes.com/2023/02/03/technology/chatgpt-openai-artificial-intelligence.html>
- Rose, J. (2022, September 7). Why Does This Horrifying Woman Keep Appearing in AI-Generated Images? *Vice*.  
<https://www.vice.com/en/article/g5vjw3/why-does-this-horrifying-woman-keep-appearing-in-ai-generated-images>
- Routhier, D. (2023). *With and Against: The Situationist International in the Age of Automation*. Verso.
- Sadowski, J. (2018, August 6). *Potemkin AI*. Real Life. <http://reallifemag.com/potemkin-ai/>
- Salvaggio, E. (2023, October 29). *The Hypothetical Image*. Cybernetic Forests.  
<https://www.cyberneticforests.com/news/social-diffusion-amp-the-seance-of-the-digital-archive>
- Sam Altman [@sama]. (2022, February 4). *Techno-optimism is the only good solution to our current problems. Unfortunately, somehow expressing optimism about the future has become a radical act.* [Tweet]. Twitter.  
<https://twitter.com/sama/status/1489648210842816516>
- Sampson, T. D. (2017). *The Assemblage Brain: Sense Making in Neuroculture*. University of Minnesota Press.
- Shapin, S., & Schaeffer, S. (1985). *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life*. Princeton University Press.
- Shapiro, A. (2018). Between autonomy and control: Strategies of arbitrage in the “on-demand” economy. *New Media & Society*, 20(8), 2954–2971. <https://doi.org/10.1177/1461444817738236>
- Smith, H., & Burrows, R. (2021). Software, Sovereignty and the Post-Neoliberal Politics of Exit. *Theory, Culture & Society*, 38(6), 143–166. <https://doi.org/10.1177/0263276421999439>
- Stark, L. (2023a). Breaking Up (with) AI Ethics. *American Literature*, 95(2), 365–379. <https://doi.org/10.1215/00029831-10575148>
- Stark, L. (2023b, March 16). *ChatGPT is Mickey Mouse*. Daily Nous. <https://dailynous.com/2023/03/16/philosophers-on-next-generation-large-language-models/>
- Statement on AI Risk. (2023, June 2). Center for AI Safety. <https://www.safe.ai/statement-on-ai-risk#signatories>
- Stilgoe, J. (2023). We need a Weizenbaum test for AI. *Science*, 381(6658), eadk0176. <https://doi.org/10.1126/science.adk0176>

- Suchman, L. (2023). The uncontroversial ‘thingness’ of AI. *Big Data & Society*, 10(2).  
<https://doi.org/10.1177/20539517231206794>
- Steyerl, H. (2016). A sea of data: apophenia and pattern recognition. *e-flux*, 72. <http://www.e-flux.com/journal/a-sea-of-data-apophenia-and-pattern-mis-recognition/>
- Taylor, A. (2018). The Automation Charade. *Logic*, 5. <https://logicmag.io/failure/the-automation-charade/>
- Thatcher, J., O’Sullivan, D., & Mahmoudi, D. (2016). Data colonialism through accumulation by dispossession: New metaphors for daily data. *Environment and Planning D: Society and Space*, 34(6), 990–1006.  
<https://doi.org/10.1177/026377581663319>
- Turing, A. (1950). Computing Machinery and Intelligence. *Mind, New Series*, 59(236), 433–460.  
<https://doi.org/10.1093/mind/LIX.236.433>
- Turner, F. (2016). Prototype. In B. Peters (Ed.), *Digital Keywords: A Vocabulary of Information, Society & Culture* (pp. 256–268). Princeton University Press.
- Turner, V. (1976). Social Dramas and Ritual Metaphors. In R. Schechner & M. Schumann (Eds.), *Ritual, Play, and Performance: Readings in The Social Sciences / Theatre* (pp. 97–120). The Seabury Press.
- Turow, J., Hennessy, M., & Draper, N. (2015). *The Tradeoff Fallacy: How Marketers Are Misrepresenting American Consumers And Opening Them Up To Exploitation*. Annenberg School for Communication, University of Pennsylvania.
- UN tech agency rolls out human-looking robots for questions at a Geneva news conference. (2023, July 7). AP News.  
<https://apnews.com/article/humanoid-robots-better-leaders-ai-geneva-486bb2bad260454a28aaa51ea31580a6>
- Utrata, A. (2023). Engineering Territory: Space and Colonies in Silicon Valley. *American Political Science Review*, 1–13.  
<https://doi.org/10.1017/S0003055423001156>
- Vidler, A. (1999). *The Architectural Uncanny: Essays in the Modern Unhomely*. MIT Press.
- Vinsel, L. (2021). *You’re Doing It Wrong: Notes on Criticism and Technology Hype*. Medium. <https://sts-news.medium.com/youre-doing-it-wrong-notes-on-criticism-and-technology-hype-18b08b4307e5>
- Wang, A., Kapoor, S., Barocas, S., & Narayanan, A. (2023). Against Predictive Optimization: On the Legitimacy of Decision-Making Algorithms that Optimize Predictive Accuracy. *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 626. <https://doi.org/10.1145/3593013.3594030>
- Watson, S. M. (2014, June 16). *Data Doppelgängers and the Uncanny Valley of Personalization*. The Atlantic.  
<http://www.theatlantic.com/technology/archive/2014/06/data-doppelgangers-and-the-uncanny-valley-of-personalization/372780/>
- Weatherby, L. (2023, November 13). *ChatGPT broke the Turing test*. The Boston Globe.  
<https://www.bostonglobe.com/2023/11/13/opinion/turing-test-ai-chatgpt/>
- Weatherby, L., & Justie, B. (2022). Indexical AI. *Critical Inquiry*, 48(2), 381–415. <https://doi.org/10.1086/717312>
- Weiss, D. C. (2023, March 16). *Latest version of ChatGPT aces bar exam with score nearing 90th percentile*. ABA Journal.  
<https://www.abajournal.com/web/article/latest-version-of-chatgpt-aces-the-bar-exam-with-score-in-90th-percentile>
- Weizenbaum, J. (1976). *Computer Power and Human Reason: From Judgment to Calculation*. W.H. Freeman and Company.
- Whittaker, M. (2021). The Steep Cost of Capture. *ACM Interactions*, November-December, 51–55.  
<https://doi.org/10.1145/3488666>
- Wong, M. (2023, June 2). *AI Doomerism Is a Decoy*. The Atlantic.  
<https://www.theatlantic.com/technology/archive/2023/06/ai-regulation-sam-altman-bill-gates/674278/>