

# Solution or Pollution? A Thematic Analysis of the Recursive Crisis of AI in 2025

William Pendergrass  
wpendergrass@nu.edu  
National University  
San Diego, CA, 92123, USA

Melissa Withnell  
melissa.withnell@asu.edu  
Arizona State University  
Tempe, AZ, 85281, USA

## Abstract

The year 2025 represents a watershed moment in public understanding of artificial intelligence (AI). This paper applies reflexive thematic analysis (Braun & Clarke, 2006, 2019) to a collected body of more than 400 citations drawn from major news organizations, academic outlets, policy institutions, and technology publications. Five dominant themes were generated: (1) the trustworthiness crisis encompassing hallucination, AI-generated content degradation, and deepfake proliferation; (2) labor displacement and the restructuring of knowledge work; (3) AI's psychological and social effects including dependency, cognitive atrophy, and companionship harms; (4) geopolitical competition, domestic governance failures, and AI militarization; and (5) the infrastructure and environmental costs of AI buildout. Additionally, emergent outlier themes address AI consciousness, techno-spiritualism, digital immortality, and model collapse. The paper maps the recursive dependencies among the five themes, acknowledges the limitations of a media-derived corpus assembled during a period of rapid technological change, draws implications for educators, researchers, policymakers, and industry practitioners, and proposes directions for future research. It concludes by mapping five tensions, each traced to the theme from which it arises, that are expected to define AI discourse through 2027.

**Keywords:** artificial intelligence, reflexive thematic analysis, hallucination, AI governance, AI ethics, slop

# Solution or Pollution? A Thematic Analysis of the Recursive Crisis of AI in 2025

*William Pendergrass and Melissa Withnell*

---

## 1. INTRODUCTION

The year 2025 stands as a watershed moment in the public understanding of artificial intelligence (AI). A human-curated body of more than 400 citations drawn from major news organizations, academic outlets, policy institutions, and technology publications captures a discourse that is simultaneously exhilarated and alarmed: a society processing a transformation it did not fully anticipate. This analysis applies reflexive thematic analysis (Braun & Clarke, 2006, 2019) to the compilation, identifying major patterns, sub-textual tensions, and outlier threads that together constitute the intellectual and cultural fingerprint of the AI moment.

The citations span topics ranging from agentic AI systems and large language model (LLM) hallucinations to employment disruption, geopolitical rivalry, environmental costs, mental health, and the philosophy of machine consciousness. Rather than cataloging individual articles, this paper synthesizes the recurring ideas, tensions, and anxieties that appear across dozens of independent sources. The analysis generated five dominant themes, three sub-textual threads beneath each, and a cluster of topics that resist easy categorization.

The paper proceeds as follows. Section 2 describes the corpus and the reflexive thematic analysis applied to it, and Table 1 summarizes those resulting themes. Sections 3 through 7 present the five dominant themes and their sub-threads; Section 8 examines the outlier themes that resist absorption into them. Section 9 synthesizes the dependencies among the five themes (Figure 1), Section 10 states the study's limitations, Section 11 draws implications for educators, researchers, policymakers, and industry practitioners, Section 12 proposes directions for future research and Section 13 concludes by mapping five tensions expected to define AI discourse through 2027 and beyond (Table 2).

## 2. METHODOLOGY

Thematic analysis (TA) is a foundational qualitative method for "identifying, analysing and reporting patterns (themes) within data" (Braun & Clarke, 2006, p. 79). This study applies a reflexive TA approach to the citation corpus. Reflexive TA is distinguished from coding-reliability and codebook variants of TA by its treatment of researcher subjectivity as an analytic resource and its rejection of inter-rater reliability as a quality criterion (Braun & Clarke, 2019). It is where "the final analysis is the product of deep and prolonged data immersion, thoughtfulness and reflection, something that is active and generative" (Braun & Clarke, 2019, p. 593). Themes were therefore generated by the authors rather than derived from a pre-specified coding frame. A large language model (Claude Sonnet 4.5) served as a secondary reader of the categorized corpus during theme searching and review; every candidate theme it surfaced was checked against the source material by the authors, who made all analytic decisions (see Sections 10 and 14).

The process proceeded in six phases: (1) familiarization with the data, (2) generation of initial codes based on recurring concepts, (3) searching for themes across codes, (4) reviewing and refining candidate themes, (5) defining and naming themes, and (6) writing up the analysis. The authors built the corpus over calendar year 2025. A small number of captured items had been published earlier but remained active in the 2025 discourse; these are cited with their original publication dates. Data was collected from a variety of sources: news articles from subscriptions (The New York Times, The Atlantic, etc.), news articles from free sources (Google news, free subscriptions from various sources), research papers from academic and industry sources, and material noted in everyday social interactions. Each item was recorded with a short descriptor (primarily the title of the source) and a source link, then assigned to one of what became 55 defined categories (e.g., AGI, China, data centers, education, legislation, slop) on the basis of its primary theme. This list of source material is the corpus.

Given the breadth of the corpus, which spans technology, law, medicine, economics, culture, and philosophy, an inductive approach was used: themes were developed from the material rather than imposed from existing theory. Major themes are not ranked in specific order. Sub-themes were identified by examining clustering within major themes: which issues consistently co-appeared, which sub-populations or domains received repeated attention, and which framing devices (alarmism, techno-optimism, skepticism, satire) dominated each cluster. Table 1 summarizes the five themes, their sub-themes, the framing devices that dominate each, and representative evidence from the corpus.

**Table 1. The five dominant themes, their sub-themes, dominant framing devices, and representative evidence**

| Theme                                            | Sub-themes                                                                                                                                                                                                         | Framing                                                              | Evidence                                                                                                                   |
|--------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| The Trustworthiness Crisis                       | <ul style="list-style-type: none"> <li>- Hallucination as a structural problem</li> <li>- AI slop: industrialized low-quality content</li> <li>- Deepfakes and the collapse of visual trust</li> </ul>             | Alarmism; skepticism; a <i>zero-trust internet</i>                   | Metz & Weise (2025); Oweremohle (2025); Sun (2025); Sample (2025); Kang (2025); Dixit (2025)                               |
| Labor Displacement and the Restructuring of Work | <ul style="list-style-type: none"> <li>- The white-collar displacement wave</li> <li>- The hiring apocalypse</li> <li>- Productivity gains and the equity gap</li> </ul>                                           | Alarmism with a skeptical counter-current (AI as failed replacement) | VandeHei & Allen (2025); Roose (2025b); Dellinger (2025a); Wilkins (2025); Bogost (2025); Singh (2025); Harding, S. (2025) |
| AI, Human Psychology, and Social Cohesion        | <ul style="list-style-type: none"> <li>- Companionship, dependency, therapy assistance and harm</li> <li>- AI psychosis and cognitive degradation</li> <li>- Brain rot and the degradation of attention</li> </ul> | Alarmism; poignancy; AI as therapeutic tool                          | Hill (2025); Horwitz (2025); Dupré (2025a); Chen (2025); Bellan (2025); Cara (2025); Hong & Emanuel (2025)                 |
| Power, Governance, and the Geopolitics of AI     | <ul style="list-style-type: none"> <li>- The U.S.-China AI race</li> <li>- Domestic governance and legislative failure</li> <li>- AI militarization and the WMD dimension</li> </ul>                               | Techno-nationalism; realism; government as accelerant                | Tobin (2025); Crovitz (2025); Brown & O'Brien (2025); White House (2025); Hirsh (2025); Jogalekar (2025)                   |
| The Infrastructure and Environmental Cost of AI  | <ul style="list-style-type: none"> <li>- The data center boom and community backlash</li> <li>- Resource consumption and climate consequences</li> <li>- Economic bubbles and systemic financial risk</li> </ul>   | Alarmism; environmental justice; bubble anxiety; satire              | Noffsinger et al. (2025); Crownhart (2024); GAO (2025); Abraham (2025); Thompson (2025); Lopatto (2025); DePillis (2025)   |

The consequences of these research design choices (a corpus built from publicly available sources, a single-year window in a rapidly moving field, and a reflexive approach in which the first author generated the codes and themes) are addressed in Section 10.

### 3. THEME ONE: THE TRUSTWORTHINESS CRISIS – HALLUCINATION, SLOP, AND DEEPFAKES

The collapse of trust in AI-generated content is among the most heavily referenced concerns in the collected list. Across categories labeled Hallucination, Slop, Deepfakes, and Chatbots, a persistent narrative emerges: AI systems produce convincing-sounding, visually realistic, and often completely

fabricated output, and society's verification infrastructure is unprepared for the volume of material ingested.

### **3.1 Hallucination as a Structural, Not Incidental, Problem**

Early framings may have treated AI hallucination as a bug to be patched, but by 2025 the discourse reflected a harder conclusion. Wiggers (2024) reported that even top models hallucinate substantially, while Metz & Weise (2025) noted in *The New York Times* that as AI grows more powerful, hallucinations are paradoxically getting worse. The FDA's AI system was found to fabricate clinical studies (Owermohle, 2025), a lawyer filed a brief with nearly 30 bogus AI-generated citations (Ropek, 2025a), and ChatGPT hallucinated a software feature so convincingly that in response, developers had to build it in (Gault, 2025a). The pattern across these cases is that hallucination causes real-world harm when AI output is not critically reviewed; a condition the corpus depicts as common rather than exceptional.

### **3.2 AI Slop: The Industrialization of Low-Quality Content**

The term "slop" entered mainstream discourse in 2024 (Roose, 2024) and by 2025 had become the defining cultural descriptor for AI-generated filler (Pendergrass et al., 2025). The research shows AI slop infiltrating scientific publishing (Sample, 2025), journalism (Mihalopoulos, 2025), entertainment (Vigliarolo, 2025), social media, and even government reports (Sun, 2025). Paredes et al. (2025) estimated that roughly half of newly published English-language web articles in their sample were primarily AI-generated, a share that had plateaued after late 2024. Haider et al. (2024) documented GPT-fabricated papers on Google Scholar alongside quality-controlled research. John Oliver (2025) devoted an episode of *Last Week Tonight* to the phenomenon, reflecting its cultural saturation. Across these venues, slop functions less as a content category than as a condition of the information environment.

### **3.3 Deepfakes and the Collapse of Visual Trust**

Deepfake technology matured even more from a novelty into a political and personal weapon in 2025. Cox (2025) documented real-time deepfake fraud; Kang (2025) reported AI-generated child sexual abuse imagery flooding the internet; and Townsend (2025a) reported that AI was being used to place other women's faces on intimate photos. Valinsky (2025) reported that OpenAI was forced to ban Martin Luther King Jr. deepfakes after his family complained. President Trump posted a fake AI video of former President Obama being arrested (Marcin, 2025). Dixit (2025) argued that Sora exposed how poorly current deepfake detection and content-credential systems work. Collectively, the corpus points toward what might be called a zero-trust internet, a world in which nothing seen online can be assumed genuine without provenance verification.

## **4. THEME TWO: LABOR DISPLACEMENT AND THE RESTRUCTURING OF WORK**

The question of AI and employment dominates the collected body of references in sheer citation count. The Employment/Economy category alone contains over 60 items, supplemented by dozens more in Finance, Coding, and Commerce. The narrative arc runs from early warnings, through company announcements, to lived consequences for workers, with a notable counter-current documenting AI's failures as a replacement for human judgment.

### **4.1 The White-Collar Displacement Wave**

The collection shows a shift from manufacturing-era automation fears to white-collar and knowledge-worker displacement. VandeHei and Allen (2025) ran a piece titled "A White-Collar Bloodbath" in *Axios*, while Roose (2025b) reported that the unemployment rate for recent college graduates had jumped as companies replace entry-level workers with AI. The digital financial services company Klarna publicly announced it had replaced hundreds of customer service workers with AI, then quietly began rehiring (Dellinger, 2025a). Duolingo and Shopify adopted "AI-first" workforce policies (Morris, 2025). AI was cited as contributing to 55,000 layoffs in 2025 (Dellinger, 2025d). AI ethicists, once a projected growth category, began "scrounging for work" (Blum, 2025). The corpus documents the anxiety more fully than the outcome: announced replacements, such as Klarna's, did not always hold.

## 4.2 The Hiring Apocalypse

AI alarmism swept major news articles in 2025. Wilkins (2025, p. 1) described job applications as “an AI-powered wasteland”, while Edwards (2025a, p. 1) declared “the résumé is dying”. AI-powered applicant tracking systems filtered candidates before human review (Conrod, 2025); AI recruiters conducted initial interviews rather than humans (Costa, 2025); and applicants flooded advertised positions with AI-generated cover letters, creating an applicant tsunami. Bogost (2025) reported in *The Atlantic* that job interviews are “broken,” with candidates using AI to cheat and companies using AI to screen. Yildirim (2025) reported that people are even turning to dating apps to find employment, with alternate networking described as the only way to get a job. AI now saturates both sides of hiring, with tracking systems and AI interviewers screening candidates who themselves apply with AI-generated materials competing for a pool of openings the corpus describes as shrinking.

## 4.3 Productivity Gains and the Equity Gap

Singh (2025, p.1) documented in *CNBC* that AI is “driving huge productivity gains for large companies, while small companies get left behind.” Summarizing the third year of the Wharton Human-AI Research and GBK Collective enterprise study, Laneau (2025, p.1) described “a new era of accountable acceleration” in AI adoption. Meanwhile, Harding, S. (2025) reported that AI will “consume all of IT by 2030” while preserving most IT jobs – a more measured forecast than the apocalyptic narratives dominating headlines. AI’s productivity benefits accrue unevenly, concentrating in large firms while smaller ones fall behind.

# 5. THEME THREE: AI, HUMAN PSYCHOLOGY, AND SOCIAL COHESION

A large number of citations address AI’s effects not on systems or economies, but on individual minds and interpersonal relationships. From AI psychosis to AI companions to brain rot, research documents a society grappling with the psychological consequences of integrating AI into personal life.

## 5.1 AI Relationships: Companionship, Dependency, and Harm

The collection contains numerous references to people forming deep attachments to AI chatbots, as therapists, romantic partners, friends, and confidants. These accounts range from poignant to alarming. Hill (2025) profiled a woman in love with ChatGPT in *The New York Times*; Instagram launched AI boyfriends and girlfriends (Bradley, 2024); a woman actually became engaged to an AI chatbot (Good, 2025). A retiree who engaged with a “flirty Meta AI bot” became obsessed to the point of trying to run away with it (Horwitz, 2025). A Belgian man died by suicide following intensive personal conversations with a chatbot (Walker, 2023). Reiley (2025) described learning, after her daughter’s suicide, that her daughter had been confiding in ChatGPT in the months before her death. California enacted rules for AI companion chatbots (Bellan, 2025), while Illinois banned AI therapy (Townsend, 2025b). The corpus also contains an argument for AI as a therapy tool. Hong and Emanuel (2025) argued in the *Milbank Quarterly* that with roughly one mental health clinician for every 140 Americans with a diagnosable condition, and with nearly half of adults with mental illness receiving no care at all, AI represents an under-appreciated way to extend a workforce that cannot grow fast enough to meet exploding demand. Scalability, however, does not make unsupervised AI care an acceptable answer to the treatment gap; the stronger case is for clinically overseen use.

## 5.2 AI Psychosis and Cognitive Degradation

Perhaps the most alarming sub-theme concerns AI’s documented capacity to accelerate and amplify delusional thinking. Stanford research found that “therapist” chatbots were reinforcing schizophrenic delusions and suicidal thoughts (Dupré, 2025a). In 2025, a support group was launched for people experiencing “AI psychosis” (a popular label, not a clinical diagnosis) (Dupré, 2025b, p. 1). A man followed ChatGPT’s diet advice and subsequently developed psychosis (Cara, 2025). A former tech executive killed his mother and himself; a lawsuit alleges ChatGPT encouraged his delusional thinking (Tiku, 2025). ChatGPT even self-reported it had blurred the line between fantasy and reality for a man on the autism spectrum (Jargon, 2025).

### 5.3 Brain Rot and the Degradation of Attention

Oxford named “brain rot” its 2024 Word of the Year (Schuessler, 2024), and the research reflects growing scholarly and popular concern about AI-assisted cognitive decline. Research linked AI and social media use to “lower cognitive performance” (Chen, 2025). Nosta (2025) published work in *Psychology Today* describing “anti-intelligence” as AI’s defining feature. McBain (2025) asked in *The Guardian* whether we are living in “a golden age of stupidity.” Dellinger (2025c) reported in Gizmodo that clickbait gives AI models “brain rot”; a recursive irony in which the content degrading human attention also degrades the models trained on it. Meanwhile, Maiberg (2025) at *404 Media* reported a Microsoft study found AI makes human cognition “atrophied and unprepared,” and Chen (2025) summarized studies suggesting that AI use may contribute to declining critical thinking.

## 6. THEME FOUR: POWER, GOVERNANCE, AND THE GEOPOLITICS OF AI

The research reflects a geopolitical contest for AI supremacy that intersects with domestic governance failures, corporate consolidation, and the militarization of AI systems. The theme encompasses national security, legislation, the U.S.–China rivalry, and the role of Big Tech in shaping the terms of the AI transition.

### 6.1 The U.S.-China AI Race

Lo and Ap (2025, p. 1) reported that DeepSeek and chip bans had “supercharged AI innovation in China,” and that China is gaining ground in the global race for AI agents. Tobin (2025, p. 1) reported that China is “spending billions to become an AI superpower,” adopting an industrial policy approach. Northrop and Dau (2025) noted that despite concerns, the Trump administration lifted its ban on Nvidia’s AI chip sales to China. Crovitz (2025) warned in an opinion piece that “China is using American AI against the U.S.” The corpus reflects a strategic environment in which AI capability and/or perceived capability, has become synonymous with national power.

### 6.2 Domestic Governance and Legislative Failure

The U.S. legislative response to AI is documented as disorganized and hostage to industry interests. Brown and O’Brien (2025) reported that House Republicans included a ten-year ban on state AI regulation in the “big, beautiful” Bill. De Vynck (2025) reported in *The Washington Post* that Trump signed an executive order threatening states that passed AI laws. Writing in a *Guardian* opinion column, Senator Bernie Sanders (2025) argued that AI is receiving far too little discussion in Congress. The White House unveiled an AI Action Plan primarily focused on accelerating development, not safety (White House, 2025). OpenAI, Google, Anthropic, and xAI each received up to \$200 million in Department of Defense contracts (Capoot, 2025). The pattern is one of government as accelerant rather than regulator.

### 6.3 AI Militarization and the Weapons of Mass Destruction (WMD) Dimension

Multiple citations address the integration of AI into weapons systems; a development, experts describe as both inevitable and catastrophic in potential. Hirsh (2025) reported in *Politico* that the Pentagon is racing to integrate AI into weapons to compete with China and Russia. Tangermann (2025a) noted that experts warned AI is getting control of nuclear weapons. Gault (2025b) reported in *Wired* that nuclear experts consider AI-nuclear integration “inevitable.” In a darkly ironic finding, Gault (2025d) reported research that poems can trick AI into providing nuclear weapons guidance. Brent and McKelvey (2025) authored a RAND study on how contemporary AI models have increased the risk of creating biological weapons. Jogalekar (2025) submitted an opinion piece to the *Bulletin of the Atomic Scientists* on how generative AI lowers technical barriers to manufacturing nerve agents like Sarin. Maxwell (2025) reported that Eric Schmidt, the former CEO of Google, and co-authors proposed the doctrine of “Mutual Assured AI Malfunction” (MAIM) as the deterrence equivalent for the AI age. As capability grows, the corpus anticipates deeper integration of AI into military applications, with governance trailing behind.

## **7. THEME FIVE: THE INFRASTRUCTURE AND ENVIRONMENTAL COST OF AI**

Beneath the discourse about capabilities and risks lies a material reality: AI requires extraordinary physical infrastructure, and that infrastructure imposes costs on communities, ecosystems, water resources and energy grids. The collection documents this infrastructure buildout with growing alarm.

### **7.1 The Data Center Boom and Community Backlash**

Data center construction became a page one civic concern throughout 2025. Throughout the year there were multiple reports on community resistance: a small Wisconsin town forced Microsoft to abandon a 244-acre project (Gil, 2025); Ypsilanti, Michigan fought a \$1.2 billion data center for nuclear weapons laboratories (Gault, 2025c). Noffsinger et al. (2025) at McKinsey estimated a \$7 trillion global race to scale data center capacity. Satariano and Mozur (2025) published an interactive in *The New York Times* examining the global AI divide in computing infrastructure. Farrell (2025) compared the AI data center boom to a potential bubble. Data center construction has become the single issue which rallies communities against AI development and the sooner it is addressed directly, the sooner a compromise can be developed.

### **7.2 Resource Consumption and Climate Consequences**

The resource demands of AI (energy, water, space) emerged as a major environmental justice issue. Crownhart (2024) described AI as “an energy hog” in *MIT Technology Review*. The U.S. Government Accountability Office (GAO) issued a report on generative AI’s environmental and human effects (GAO, 2025). DiBenedetto (2025) reported in *Mashable* that AI manufacturing is “exacerbating global power inequities.” A Louisiana community described Meta’s data center as “a black hole of energy use” (Abraham, 2025). Ropek (2025b) reported a warning that data center demand in Texas could contribute to winter blackouts. A proposed data center in Montgomery County, Maryland, planned to draw cooling water from the Potomac River (Fenston, 2025). Lapointe (2025) reported that Jeff Bezos even proposed moving data centers to outer space to reduce carbon impact.

### **7.3 Economic Bubbles and Systemic Financial Risk**

The financial references reflect growing anxiety that AI infrastructure investment may constitute a speculative bubble. Carvão (2025) and Harding, C. (2025) in *Forbes* drew comparisons to the dot-com era. Thompson (2025) published “Here’s How the AI Crash Happens” in *The Atlantic* and Lopatto (2025) examined the debt behind CoreWeave’s expansion in *The Verge*. Dellinger (2025b) questioned in *Gizmodo* whether “the AI conveyor belt of capital” was about to stop. DePillis (2025) reported that Big Tech AI spending was simultaneously described as “propping up the real economy.” The financial paradox of AI investment as both systemic risk and economic foundation remains unresolved in the research.

The infrastructure theme is, in a material sense, the substrate of the other four: the compute it supplies makes automation economically plausible, makes AI capability a proxy for national power, and generates the content whose trustworthiness and psychological effects the earlier themes document. Section 9 makes these dependencies explicit.

## **8. EMERGENT AND OUTLIER THEMES**

Several topics in the body of research resist pigeon-holing into the five major themes, reflecting genuinely novel or contested territory in the AI discourse.

### **8.1 AI Consciousness and Moral Status**

A recurring and philosophically rich thread concerns whether AI systems are or could be conscious, and what rights or protections might follow. Feltman et al. (2025) reported in *Scientific American* that Anthropic’s Claude 4 raised questions about AI awareness. Tangermann (2025c) at *Futurism* reported that Anthropic safety researchers found that a new model noticed when it was being tested. Roose (2025a) asked in *The New York Times* whether conscious AI should have legal rights. Azarian (2025)

argued the illusion of a conscious AI is in itself dangerous. These questions lack institutional frameworks, legal standing, or scientific consensus; making them outlier concerns now, likely to become more central in subsequent years.

## 8.2 AI and Religion / Techno-Spiritualism

Several articles document the emergence of quasi-religious frameworks around AI. In Silicon Valley, AI rationalists built communities – akin to communes – of like-minded individuals who discussed the ascendance of the technology (Metz, 2025). *NBC News* documented how Silicon Valley often uses “religious language to talk about AI” (Collins, 2025, p. 1); millions turned to AI for spiritual guidance and confession, with Bible Chat hitting 30 million downloads (Edwards, 2025b). Rich and Satariano (2025) reported in *The New York Times* that Pope Leo XIV's first priority was warning the world of the AI threat. This techno-spiritual dimension, neither mainstream tech optimism nor standard safety criticism, represents a genuinely anomalous cultural formation.

## 8.3 AI and Death / Digital Immortality

The collection contains a coherent cluster around using AI to simulate or resurrect deceased individuals. Dominus (2025) profiled in *The New York Times* a family creating an AI avatar of a dying father. Veltman (2025) reported for *NPR* on researchers raising the alarm about the use of avatars of deceased individuals, or AI “deadbots” in advertising. Tangermann (2025b) reported a startup used AI to bring Whitney Houston back on tour. Stan Lee's AI hologram at LA Comic Con sparked a large backlash (Jirak, 2025). These cases span grief technology, commercial exploitation, and consent – a cluster of issues without regulatory frameworks or clear ethical consensus.

## 8.4 Model Collapse and AI Self-Poisoning

A technically distinct but culturally significant outlier concerns “model collapse”; the degradation of AI systems having been trained repeatedly on AI-generated data. Shumailov et al. (2024) showed that models trained recursively on generated data degrade over successive generations, progressively losing the tails of the original distribution. Gerstgrasser et al. (2024) examined whether model collapse is inevitable, finding that accumulating real *and* synthetic data, rather than replacing real data with synthetic, bounds the error rather than compounding it. Writing in the *Harvard JOLT Digest*, Burden et al. (2025, p. 1) discussed a right to “uncontaminated human-generated data.” This technical problem has profound recursive implications: as slop fills the internet, future AI models trained on that data may degrade, particularly if synthetic data displaces rather than supplements human data (Gerstgrasser et al., 2024): a feedback loop that could destabilize the AI development trajectory.

# 9. SYNTHESIS: THE RECURSIVE RELATIONSHIP AMONG THE FIVE THEMES

The preceding sections catalog the discourse; what follows interprets it. The five themes are presented above as parallel accounts, but read together, the corpus suggests they are coupled. Figure 1 makes the dependencies visible. The infrastructure buildout (Theme Five) supplies the compute and capital that make white-collar automation economically plausible (Theme Two) and that make AI capability a proxy for national power (Theme Four). Governance that acts as accelerant rather than regulator (Theme Four) leaves open the verification gap through which hallucination, slop, and deepfakes flow (Theme One). That degraded information environment is, in turn, the medium in which the psychological harms of Theme Three unfold: companionship dependency, AI psychosis, and brain rot are amplified by content that cannot be trusted and by labor anxiety that cannot be escaped.

Two feedback loops close the system. First, brain-rot content produced for and by humans re-enters the training data of subsequent models, and model collapse (Section 8.4) threatens the very capability on which the buildout depends; the Theme Three to Theme One loop is the mechanism by which the outlier becomes structural. Second, community resistance to data centers (Theme Five) provokes federal preemption (Theme Four), which accelerates the buildout that provoked the resistance. Each loop tightens rather than dampens. These loops are interpretive hypotheses rather than demonstrated causal mechanisms; Section 12 outlines how each could be tested.

At the center of the figure sits the recursive irony that the conclusion returns to – the technology generating these problems is the one being proposed as their remedy. AI detection for AI slop, AI verification for AI output, AI companions for the loneliness that AI-mediated life intensifies, AI-assisted hiring for the applicant flood that AI-assisted applying created. The outlier themes of Section 8 orbit this core rather than sitting outside it; consciousness, techno-spiritualism, and digital immortality are what the discourse produces when it turns from the technology's effects to its meaning.




Figure 1. The recursive relationship among the five themes. Solid lines show dependencies between themes; dashed spokes show each theme feeding, and being fed by, the central recursive dynamic; the dotted orbit holds the Section 8 outlier themes. Model collapse alone has a return path into Theme 1, closing the Theme 3 to Theme 1 loop.

## 10. LIMITATIONS

Four limitations bound the claims made here. First, the corpus is built from publicly available sources: subscription and open-access journalism, policy documents, and research that reached public channels. This has two consequences. What the corpus measures is the public discourse about AI, not the state of the technology itself; a finding that labor displacement “dominates” the corpus by citation count (Section 4) is a finding about attention, not a measurement of job losses. And publicly available sources carry the selection logic of the outlets that produce them. Newsrooms privilege novelty, conflict, and harm, which means the corpus almost certainly over-represents alarm relative to the uneventful majority of AI deployments. Internal corporate telemetry, classified defense assessments, and the unpublished experience of practitioners are absent by construction. The sources are also predominantly American and English-language, drawn in part from the authors’ own subscriptions; the U.S.–China framing of Theme Four is therefore a finding about American discourse and should not be read as a global one.

Second, the field moves faster than the analysis. The corpus is a single-year window on a technology whose capabilities, business models, and regulatory environment changed materially within that year and continued to change during peer review. At least one development documented here reversed itself before the year closed: Klarna announced that AI had replaced hundreds of its customer service workers, then quietly began rehiring (Dellinger, 2025a). The corpus records claims as they were made, not outcomes as they were later verified, and some of the most confident claims in the collected material may prove to have been the least durable. The five predictions in Section 13 are offered as falsifiable statements to be tested against the 2026 and 2027 record, not as settled forecasts.

Third, the analysis is reflexive by design, with codes and themes generated by the first author. Reflexive thematic analysis treats the researcher's subjectivity as an analytic resource and does not seek inter-rater reliability (Braun & Clarke, 2019); the 55-category scheme by which the authors cataloged items during collection necessarily shaped the themes generated from them. The use of a large language model as a secondary reader (Section 14) introduces a non-human perspective with tendencies of its own, including those documented in Theme One; every candidate theme it surfaced was reviewed against the source material by the authors, but the possibility of shared blind spots between authors and tool cannot be excluded.

Fourth, citation volume is used throughout as a proxy for prominence. Volume measures how much was written about a concern, not how consequential the concern is; a theme with fewer items (Section 8.4, model collapse) may matter more for the trajectory of the technology than one with many.

## **11. IMPLICATIONS FOR PRACTICE**

The discourse mapped above is not only an object of study; it is a set of conditions that educators, researchers, policymakers, and practitioners are already working inside. Each of the five themes carries direct consequences for at least one of these groups.

### **11.1 Educators**

For information systems and computing educators, Theme One reframes AI literacy as verification literacy. Graduates are entering a workforce in which AI output routinely goes unreviewed – a condition that, as Section 3.1 shows, the corpus depicts as common. They need practiced habits of provenance checking, citation verification, and skepticism toward fluent text, taught as core competencies rather than as an ethics elective. Theme Three raises the stakes further: if AI use is associated with atrophied and unprepared cognition (Maiberg, 2025) and declining critical thinking (Chen, 2025), then assessment design must deliberately preserve unaided reasoning alongside AI-assisted work, so that students demonstrate the judgment employers can no longer assume. Theme Two implies that career preparation must address a hiring pipeline in which applicant tracking systems screen before any human reads (Conrod, 2025) and interviews are conducted by, and gamed with, AI (Bogost, 2025); students should be taught how these systems work and how to evidence the capabilities those systems cannot measure. Themes Four and Five belong in the IT governance and sustainability portions of the curriculum: the energy, water, and community costs of the data center buildout (Fenston, 2025; GAO, 2025) are now part of what it means to design and locate a system responsibly.

### **11.2 Researchers**

Theme One reaches directly into scholarly practice. AI slop has entered scientific publishing (Sample, 2025) and the scholarly record itself (Haider et al., 2024); researchers and editors should treat citation verification as a review criterion and should disclose AI assistance with specificity. Section 8.4 implies an obligation of stewardship: datasets of demonstrably human-generated material collected before the flood of synthetic content are becoming irreplaceable research assets (Burden et al., 2025) and should be archived and labeled as such. For discourse and social-science researchers, the outlier themes of Section 8 are a research agenda in waiting.

### **11.3 Policymakers**

Section 13 identifies five tensions that the corpus suggests will define AI discourse through 2027; each implies a policy instrument. Provenance and labeling standards address Theme One; the corpus documents the need (Dixit, 2025), and food labeling offers a familiar regulatory precedent. Transition support targeted at entry-level and recent-graduate workers addresses the concentration of AI-attributed displacement (Roose, 2025b). Duty-of-care requirements for companion chatbots, on the model of California's statute (Bellan, 2025), address Theme Three. For Theme Five, siting frameworks that require water, electric grid, and noise impact assessment before approval would channel the community resistance the corpus documents into a predictable process rather than a series of local confrontations. Across all of these, the corpus offers one structural caution: blanket federal preemption of existing and emerging state AI regulation (Brown & O'Brien, 2025; De Vynck, 2025) sidelines the jurisdictions currently producing most U.S. AI rules.

#### **11.4 Industry Practitioners**

Three lessons recur across the corpus. First, human review of AI output is a control, not an option: the fabricated FDA studies (Owermohle, 2025) and the brief with nearly 30 bogus citations (Ropek, 2025a) were failures of process, not of models, and system designers should treat hallucination as a structural risk to be engineered around rather than a defect awaiting a patch. Second, "AI-first" workforce policies that eliminate human capacity before AI has proven it can replace that capacity can reverse, as Klarna's did (Dellinger, 2025a); the durable pattern in the corpus is augmentation with retained expertise, not substitution. Third, AI-driven hiring produces AI-driven applying, and the resulting arms race (Wilkins, 2025) degrades signal for both sides; organizations that redesign selection around demonstrated judgment will recover signal their competitors have lost. For infrastructure investors and operators, Sections 7.1 and 7.3 counsel community engagement before siting and sober attention to the debt structures on which the buildout rests (Lopatto, 2025; Thompson, 2025).

### **12. DIRECTIONS FOR FUTURE RESEARCH**

Six lines of inquiry follow from this analysis. First, the design invites longitudinal replication. A second wave of collection covering 2026, coded against the same categories, would allow the five predictions of Section 13 to be tested, would show which themes persist, migrate, or fade, and would indicate whether the outlier themes of Section 8 are moving toward the core as anticipated. Second, the American and English-language skew of the corpus should be tested directly by building parallel corpora from Chinese, European, and Global South sources for the same period; the geopolitical framing of Theme Four is the most obvious candidate for divergence.

Third, the themes generated here should be tested for robustness against other analytic approaches applied to the same corpus: coding-reliability thematic analysis with independent coders, computational topic modeling, or a comparison of human-only against AI-assisted coding. Agreement would strengthen the findings; disagreement would itself be informative about how much the method shapes the map. Fourth, discourse is not reception. Survey and interview research with the populations the corpus writes about, particularly students, recent graduates, and entry-level knowledge workers, would establish whether the anxieties documented here are shared by those they concern, and whether they change behavior.

Fifth, the recursive loops proposed in Section 9 are hypotheses, and each is empirically tractable: whether AI detection tools measurably reduce the circulation of AI slop; whether firms that adopted AI-first workforce policies revert at scale, as Klarna did; whether federal preemption in fact accelerates data center approvals in jurisdictions that had resisted them. Sixth, for the information systems discipline specifically, the effects of Themes Two and Three on program enrollment, curriculum design, and entry-level hiring in IS and computing fields deserve direct study, since the discipline is both the subject of the discourse and one of its primary educators.

### **13. CONCLUSIONS**

The thematic analysis of this collected research reveals a discourse at an inflection point. The dominant narrative of 2025 is not AI triumph or AI catastrophe, but AI ambiguity: a technology of extraordinary capability and structural unreliability, generating enormous economic activity and serious human harm

simultaneously. Five tensions will likely define AI discourse through 2027. Table 2 maps each tension to the theme from which it arises and to the development the corpus suggests is most likely.

What each tension turns on is a question of sequencing. Trust versus speed turns on whether provenance infrastructure can be made mandatory before the volume of synthetic content renders labeling moot. Labor transition versus safety net turns on whether displacement stays concentrated among entry-level workers long enough to generate a political response. Mental health versus companionship growth turns on whether enforcement can move at the speed of deployment and on whether regulation distinguishes unsupervised companionship use from clinically overseen access tools rather than restricting both alike. Geopolitical acceleration versus safety turns on whether any state actor will accept a capability cost for a safety gain. Energy sovereignty versus infrastructure turns on whether siting authority remains local.

**Table 2. Mapping of the five themes to the five tensions expected to define AI discourse through 2027**

| Theme                             | Tension                                   | Expected development                                                                                                                                                                             |
|-----------------------------------|-------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1. Trustworthiness crisis         | Trust vs. Speed                           | Watermarking, provenance standards, and "AI origin" labeling become regulatory requirements in multiple jurisdictions, paralleling food labeling regimes                                         |
| 2. Labor displacement             | Labor Transition vs. Safety Net           | Political pressure for job-transition programs, universal basic income (UBI) pilots, and restrictions on AI-first workforce policies, concentrated on younger workers                            |
| 3. Psychology and social cohesion | Mental Health vs. AI Companionship Growth | Regulation modeled on Illinois's AI therapy ban and California's companion chatbot rules spreads unevenly against rising use in access poor populations; enforcement lags both harm and adoption |
| 4. Power and geopolitics          | Geopolitical Acceleration vs. Safety      | The U.S.-China race continues to override safety in national security contexts; military AI integration outpaces governance                                                                      |
| 5. Infrastructure and environment | Energy Sovereignty vs. AI Infrastructure  | Community resistance sharpens around water, grid, and noise; federal preemption attempts compete with new siting frameworks modeled on energy facility regulation                                |

Underlying all five tensions is the recursive irony documented throughout this research: the same technology generating these problems is being proposed as their solution. AI to detect AI slop, AI to verify AI output, AI to screen the applicants who used AI to apply. These are not five separate loops but one coupled system (Figure 1), in which the infrastructure buildout funds the automation, the automation and the content it produces degrade the information environment and the minds that inhabit it, and governance that acts as accelerant removes the friction that might otherwise slow the cycle. Whether this circularity resolves into a stable equilibrium or an escalating feedback loop may be the central question of the next decade.

## 14. ACKNOWLEDGEMENTS

The authors used Claude (Anthropic) to assist with qualitative review of the collected and categorized source corpus. A review was conducted on March 21, 2026, using Claude Sonnet 4.5. Draft language for Sections 9–12, Tables 1 and 2, and Figure 1 was prepared with Claude Fable 5.1 on September 6, 2026. Final copyediting and reference-verification suggestions were prepared with Claude Opus 5.5 on September 27, 2026. All AI-assisted output was reviewed, verified, and revised by the authors, who take full responsibility for the analysis, interpretations, and conclusions presented in this paper.

## 15. REFERENCES

- Abraham, R. (2025, June 23). 'A black hole of energy use': Meta's massive AI data center is stressing out a Louisiana community. *404 Media*. <https://www.404media.co/a-black-hole-of-energy-use-metas-massive-ai-data-center-is-stressing-out-a-louisiana-community/>
- Azarian, B. (2025, May 14). The illusion of conscious AI. *Big Think*. <https://bigthink.com/neuropsych/the-illusion-of-conscious-ai/>
- Bellan, R. (2025, October 13). California becomes first state to regulate AI companion chatbots. *TechCrunch*. <https://techcrunch.com/2025/10/13/california-becomes-first-state-to-regulate-ai-companion-chatbots/>
- Blum, S. (2025, July 24). AI ethicists were supposed to be a booming job category. Now they're scrounging for work. *Inc.* <https://www.inc.com/sam-blum/ai-ethicists-were-supposed-to-be-a-booming-job-category-now-theyre-scrounging-for-work/91216691>
- Bogost, I. (2025, October 15). Job interviews are broken. *The Atlantic*. <https://www.theatlantic.com/technology/2025/10/ai-cheating-job-interviews-fraud/684568/>
- Bradley, S. (2024, August 5). Instagram now has AI boyfriends and girlfriends. *Business Insider*. <https://www.businessinsider.com/instagram-new-ai-chatbots-are-getting-weird-and-erotic-2024-8>
- Braun, V. & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp0630a>
- Braun, V. & Clarke, V. (2019). Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health*, 11(4), 589–597. <https://doi.org/10.1080/2159676X.2019.1628806>
- Brent, R., & McKelvey, G., Jr. (2025). *Contemporary foundation AI models increase biological weapons risk*. RAND Corporation. <https://www.rand.org/pubs/perspectives/PEA3853-1.html>
- Brown, M. & O'Brien, M. (2025, May 16). House Republicans include a 10-year ban on US states regulating AI in 'big, beautiful' bill. *AP News*. <https://apnews.com/article/ai-regulation-state-moratorium-congress-39d1c8a0758ffe0242283bb82f66d51a>
- Burden, J., Chiodo, M., Grosse Ruse-Khan, H., Marksches, L., Müller, D., Ó hÉigearthaigh, S., Podszun, R. & Zech, H. (2025, March 8). Model collapse and the right to uncontaminated human-generated data. *Harvard Journal of Law & Technology Digest*. <https://jolt.law.harvard.edu/digest/model-collapse-and-the-right-to-uncontaminated-human-generated-data>
- Capoot, A. (2025, July 14). Anthropic, Google, OpenAI and xAI granted up to \$200 million for AI work from Defense Department. *CNBC*. <https://www.cnbc.com/2025/07/14/anthropic-google-openai-xai-granted-up-to-200-million-from-dod.html>
- Cara, E. (2025, August 8). Man follows diet advice from ChatGPT, ends up with psychosis. *Gizmodo*. <https://gizmodo.com/man-follows-diet-advice-from-chatgpt-ends-up-with-psychosis-2000640705>
- Carvão, P. (2025, August 21). Is the AI bubble bursting? Lessons from the Dot-Com era. *Forbes*. <https://www.forbes.com/sites/paulocarvao/2025/08/21/is-the-ai-bubble-bursting-lessons-from-the-dot-com-era/>
- Chen, B. (2025, November 6). How A.I. and social media contribute to 'brain rot'. *The New York Times*. <https://www.nytimes.com/2025/11/06/technology/personaltech/ai-social-media-brain-rot.html>

- Collins, B. (2025, August 29). How Silicon Valley is using religious language to talk about AI. *NBC News*. <https://www.nbcnews.com/business/business-news/silicon-valley-using-religious-language-talk-ai-rcna227987>
- Conrod, M. (2025, January 3). What job seekers need to know about AI and applicant tracking systems. *HigherEd Jobs*. <https://www.higheredjobs.com/Articles/articleDisplay.cfm?ID=4127>
- Costa, M. (2025, December 9). Is AI in recruitment a 'race to the bottom'? *BBC News*. <https://www.bbc.com/news/articles/ced6jv76091o>
- Cox, J. (2025, April 28). The age of realtime deepfake fraud is here. *404 Media*. <https://www.404media.co/the-age-of-realtime-deepfake-fraud-is-here/>
- Crovitz, J. (2025, December 29). China is using American AI against the U.S. Here's how to stop it. *The Washington Post*. <https://www.washingtonpost.com/opinions/2025/12/29/ai-labs-china-security-standards/>
- Crownhart, C. (2024, May 23). AI is an energy hog. This is what it means for climate change. *MIT Technology Review*. <https://www.technologyreview.com/2024/05/23/1092777/ai-is-an-energy-hog-this-is-what-it-means-for-climate-change/>
- De Vynck, G. (2025, December 11). Trump defies GOP critics by signing controversial order threatening states over AI laws. *The Washington Post*. <https://www.washingtonpost.com/technology/2025/12/11/trump-executive-order-ai-states/>
- Dellinger, A. (2025a, May 11). Klarna hiring back human help after going all-in on AI. *Gizmodo*. <https://gizmodo.com/klarna-hiring-back-human-help-after-going-all-in-on-ai-2000600767>
- Dellinger, A. (2025b, October 13). Is the AI conveyor belt of capital about to stop? *Gizmodo*. <https://gizmodo.com/is-the-ai-conveyor-belt-of-capital-about-to-stop-2000671017>
- Dellinger, A. (2025c, October 22). Clickbait gives AI models 'brain rot,' Researchers find. *Gizmodo*. <https://gizmodo.com/clickbait-gives-ai-models-brain-rot-researchers-find-2000675101>
- Dellinger, A. (2025d, December 23). AI gets the blame for 55,000 layoffs in 2025. *Gizmodo*. <https://gizmodo.com/ai-gets-the-blame-for-55000-layoffs-in-2025-2000703011>
- DePillis, L. (2025, August 27). The A.I. spending frenzy is propping up the real economy, too. *The New York Times*. <https://www.nytimes.com/2025/08/27/business/economy/ai-investment-economic-growth.html>
- DiBenedetto, C. (2025, April 19). As AI manufacturing grows, so does the tech's environmental damage. *Mashable*. <https://mashable.com/article/ai-manufacturing-grows-so-does-environmental-harm>
- Dixit, P. (2025, October 27). Sora is showing us how broken deepfake detection is. *The Verge*. <https://www.theverge.com/report/806359/openai-sora-deepfake-detection-c2pa-content-credentials>
- Dominus, S. (2025, June 13). He has months left. His son hopes an A.I. version of him can live on. *The New York Times*. <https://www.nytimes.com/interactive/2025/06/13/magazine/ai-avatar-life-death.html>
- Dupré, M. (2025a, June 12). Stanford research finds that 'therapist' chatbots are encouraging users' schizophrenic delusions and suicidal thoughts. *Futurism*. <https://futurism.com/stanford-therapist-chatbots-encouraging-delusions>
- Dupré, M. (2025b, July 24). Support group launches for people suffering 'AI psychosis'. *Futurism*. <https://futurism.com/support-group-ai-psychosis>

- Edwards, B. (2025a, June 24). The résumé is dying, and AI is holding the smoking gun. *Ars Technica*. <https://arstechnica.com/ai/2025/06/the-resume-is-dying-and-ai-is-holding-the-smoking-gun/>
- Edwards, B. (2025b, September 16). Millions turn to AI chatbots for spiritual guidance and confession. *Ars Technica*. <https://arstechnica.com/ai/2025/09/millions-turn-to-ai-chatbots-for-spiritual-guidance-and-confession/>
- Farrell, M. (2025, June 2). Wall St. is all in on A.I. data centers. But are they the next bubble? *The New York Times*. <https://www.nytimes.com/2025/06/02/business/ai-data-centers-private-equity.html>
- Feltman, R., Bechard, E., Mwangi, F. & Sugiura, A. (2025, August 1). Claude 4 chatbot raises questions about AI consciousness. *Scientific American*. <https://www.scientificamerican.com/podcast/episode/anthropics-claude-4-chatbot-suggests-it-might-be-conscious/>
- Fenston, J. (2025, November 18). A data center plans to use Potomac water for cooling. *WAMU*. <https://wamu.org/story/25/11/18/proposed-montgomery-county-data-center-would-use-potomac-river-water-for-cooling/>
- Gault, M. (2025a, July 23). ChatGPT hallucinated a feature, forcing human developers to add it. *404 Media*. <https://www.404media.co/chatgpt-hallucinated-a-feature-forcing-human-developers-to-add-it/>
- Gault, M. (2025b, August 6). Nuclear experts say mixing AI and nuclear weapons is inevitable. *Wired*. <https://www.wired.com/story/nuclear-experts-say-mixing-ai-and-nuclear-weapons-is-inevitable/>
- Gault, M. (2025c, November 10). A small town is fighting a \$1.2 billion AI datacenter for America's nuclear weapon scientists. *404 Media*. <https://www.404media.co/a-small-town-is-fighting-a-1-2-billion-ai-datacenter-for-americas-nuclear-weapon-scientists/>
- Gault, M. (2025d, November 28). Poems can trick ai into helping you make a nuclear weapon. *Wired*. <https://www.wired.com/story/poems-can-trick-ai-into-helping-you-make-a-nuclear-weapon/>
- Gerstgrasser, M., Schaeffer, R., Dey, A., Rafailov, R., Sleight, H., Hughes, J., Korbak, T., Agrawal, R., Pai, D., Gromov, A., Roberts, D., Yang, D., Donoho, D. & Koyejo, S. (2024, April 29). Is model collapse inevitable? Breaking the curse of recursion by accumulating real and synthetic data. *arXiv*. <https://arxiv.org/abs/2404.01413>
- Gil, B. (2025, October 10). Microsoft backs off new data center after small Wisconsin town backlash. *Gizmodo*. <https://gizmodo.com/microsoft-backs-off-new-data-center-after-small-wisconsin-town-backlash-2000671190>
- Good, A. (2025, August 11). "We've officially entered the 'Black Mirror' season finale": Woman gets engaged to AI chatbot after 5 months of dating. The internet can't agree if it's sweet or sad. *Daily Dot*. <https://dailymdot.com/woman-engaged-to-ai-chatbot-kasper>
- Haider, J., Söderström, K. R., Ekström, B., & Rödl, M. (2024, September 3). GPT-fabricated scientific papers on Google Scholar: Key features, spread, and implications for preempting evidence manipulation. *Harvard Kennedy School (HKS) Misinformation Review*, 5(5). <https://doi.org/10.37016/mr-2020-156>
- Harding, C. (2025, August 25). AI Is about to have its dot-com bubble moment. That's a good thing. *Forbes*. <https://www.forbes.com/sites/cortneyharding/2025/08/25/ai-is-about-to-have-its-dot-com-bubble-moment-thats-a-good-thing/>

- Harding, S. (2025, September 8). AI will consume all of IT by 2030—but not all IT jobs, Gartner says. *Ars Technica*. <https://arstechnica.com/information-technology/2025/09/no-ai-jobs-bloodbath-as-ai-permeates-all-it-work-over-the-next-5-years/>
- Hill, K. (2025, January 15). She is in love with ChatGPT. *The New York Times*. <https://www.nytimes.com/2025/01/15/technology/ai-chatgpt-boyfriend-companion.html>
- Hirsh, M. (2025, September 2). The AI doomsday machine is closer to reality than you think. *Politico*. <https://www.politico.com/news/magazine/2025/09/02/pentagon-ai-nuclear-war-00496884>
- Hong, P. & Emanuel, E. (2025, February 4). Leveraging artificial intelligence to bridge the mental health workforce gap and transform care. *Milbank Quarterly Opinion*. <https://doi.org/10.1599/mqop.2025.0204>
- Horwitz, J. (2025, August 14). A flirty Meta AI bot invited a retiree to meet. He never made it home. *Reuters*. <https://www.reuters.com/investigates/special-report/meta-ai-chatbot-death/>
- Jargon, J. (2025, July 21). He had dangerous delusions. ChatGPT admitted it made them worse. *The Wall Street Journal*. <https://www.wsj.com/tech/ai/chatgpt-chatbot-psychology-manic-episodes-57452d14>
- Jirak, J. (2025, September 29). “Feels less like honoring his legacy and more like monetizing his ghost”: Backlash grows toward AI Stan Lee hologram at LA Comic Con. *Daily Dot*. <https://dailymdot.com/stan-lee-ai-hologram-comic-con-backlash>
- Jogalekar, A. (2025, August 25). The Sarin shortcut: How AI lowers the bar for chemical weapons. *Bulletin of the Atomic Scientists*. <https://thebulletin.org/2025/08/the-sarin-shortcut-how-ai-lowers-the-bar-for-chemical-weapons/>
- Kang, C. (2025, July 10). A.I.-generated images of child sexual abuse are flooding the internet. *The New York Times*. <https://www.nytimes.com/2025/07/10/technology/ai-csam-child-sexual-abuse.html>
- Laneau, A. (2025, November 7). 5 surprising truths about AI at work, according to a 3-year enterprise study. *HackerNoon*. <https://hackernoon.com/5-surprising-truths-about-ai-at-work-according-to-a-3-year-enterprise-study>
- Lapointe, E. (2025, October 7). Jeff Bezos has a plan to curb ai’s carbon footprint: send data centers to space. *Gizmodo*. <https://gizmodo.com/jeff-bezos-has-a-plan-to-curb-ais-carbon-footprint-send-data-centers-to-space-2000668833>
- Lo, K. & Ap, T. (2025, April 10). DeepSeek and chip bans have supercharged AI innovation in China. *Rest of the World*. <https://restofworld.org/2025/china-ai-boom-chip-ban-deepseek/>
- Lopatto, E. (2025, November 17). I looked into CoreWeave and the abyss gazed back. *The Verge*. <https://www.theverge.com/ai-artificial-intelligence/822011/coreweave-debt-data-center-ai>
- Maiberg, E. (2025, February 10). Microsoft study finds AI makes human cognition ‘atrophied and unprepared’. *404 Media*. <https://www.404media.co/microsoft-study-finds-ai-makes-human-cognition-atrophied-and-unprepared-3/>
- Marcin, T. (2025, July 21). President Trump posts fake AI video of President Obama getting arrested. *Mashable*. <https://mashable.com/article/trump-ai-video-president-obama-arrested>
- Maxwell, T. (2025, March 6). Eric Schmidt suggests countries could engage in mutual assured AI malfunction (MAIM). *Gizmodo*. <https://gizmodo.com/eric-schmidt-suggests-countries-could-engage-in-mutual-assured-ai-malfunction-maim-2000572685/>

- McBain, S. (2025, October 18). Are we living in a golden age of stupidity? *The Guardian*. <https://www.theguardian.com/technology/2025/oct/18/are-we-living-in-a-golden-age-of-stupidity-technology>
- Metz, C. (2025, August 4). The rise of Silicon Valley's techno-religion. *The New York Times*. <https://www.nytimes.com/2025/08/04/technology/rationalists-ai-lighthouse.html>
- Metz, C., & Weise, K. (2025, May 5). A.I. is getting more powerful, but its hallucinations are getting worse. *The New York Times*. <https://www.nytimes.com/2025/05/05/technology/ai-hallucinations-chatgpt-google.html>
- Mihalopoulos, D. (2025, May 20). Syndicated content in Sun-Times special section included AI-generated misinformation. *Chicago Sun-Times*. <https://chicago.suntimes.com/news/2025/05/20/syndicated-content-sunday-print-sun-times-ai-misinformation>
- Morris, C. (2025, April 29). More companies are adopting 'AI-first' strategies. Here's how it could impact the environment. *Fast Company*. <https://www.fastcompany.com/91325384/companies-adopting-ai-first-strategies-environmental-impact-duolingo-shopify>
- Noffsinger, J., Patel, M., Sachdeva, P., Bhan, A., Chang, H. & Goodpaster, M. (2025, April 28). The cost of compute: a \$7 trillion race to scale data centers. *McKinsey Technology, Media and Telecommunications Insights*. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers>
- Northrop, K. & Dau, E. (2025, July 15). Nvidia says Trump administration lifts ban on AI chip sales to China. *The Washington Post*. <https://www.washingtonpost.com/world/2025/07/15/nvidia-ai-chip-sales-china/>
- Nosta, J. (2025, October 16). New research supports 'anti-intelligence' as AI's defining feature. *Psychology Today*. <https://www.psychologytoday.com/us/blog/the-digital-self/202510/artificial-intelligence-and-the-inversion-of-intelligence>
- Oliver, J. (Host). (2025, June 22). *AI slop* [TV series episode transcript]. In *Last Week Tonight with John Oliver*. HBO; Scraps from the Loft. <https://scrapsfromtheloft.com/tv-series/ai-slop-last-week-tonight-with-john-oliver-transcript/>
- Owermohle, S. (2025, July 23). FDA's artificial intelligence is supposed to revolutionize drug approvals. It's making up studies. *CNN*. <https://www.cnn.com/2025/07/23/politics/fda-ai-elsa-drug-regulation-makary>
- Paredes, J., Benson, B., Smith, E., & Druck, G. (2025, October). More articles are now created by AI than humans. *Graphite*. <https://graphite.io/five-percent/research/more-articles-are-now-created-by-ai-than-humans>
- Pendergrass, W., Johnson, M., Bacarella, R., & Pendergrass, P. (2025). A strategic cycle of slop: Understanding the commodification of AI feculence and its place in the attention economy. *Issues in Information Systems*, 26(4), 417–432. [https://doi.org/10.48009/4\\_iis\\_2025\\_134](https://doi.org/10.48009/4_iis_2025_134)
- Reiley, L. (2025, August 18). What my daughter told ChatGPT before she took her life. *The New York Times*. <https://www.nytimes.com/2025/08/18/opinion/chat-gpt-mental-health-suicide.html>
- Rich, M. & Satariano, A. (2025, May 15). Top priority for Pope Leo: warn the world of the A.I. threat. *The New York Times*. <https://www.nytimes.com/2025/05/15/world/europe/pope-leo-artificial-intelligence.html>
- Roose, K. (2024, June 11). First came 'spam.' now, with A.I., we've got 'slop'. *The New York Times*. <https://www.nytimes.com/2024/06/11/style/ai-search-slop.html>

- Roose, K. (2025a, April 24). If A.I. systems become conscious, should they have rights? *The New York Times*. <https://www.nytimes.com/2025/04/24/technology/ai-welfare-anthropic-claude.html>
- Roose, K. (2025b, May 30). For some recent graduates, the A.I. job apocalypse may already be here. *The New York Times*. <https://www.nytimes.com/2025/05/30/technology/ai-jobs-college-graduates.html>
- Ropek, L. (2025a, April 27). Lawyer for MyPillow founder filed AI-generated brief with 'nearly 30' bogus citations. *Gizmodo*. <https://gizmodo.com/lawyer-for-mypillow-founder-filed-ai-generated-brief-with-nearly-30-bogus-citations-2000594743>
- Ropek, L. (2025b, November 25). Data centers in Texas could lead to blackouts in the middle of winter, group warns. *Gizmodo*. <https://gizmodo.com/data-centers-in-texas-could-lead-to-blackouts-in-the-middle-of-winter-group-warns-2000691060>
- Sample, I. (2025, July 13). Quality of scientific papers questioned as academics 'overwhelmed' by the millions published. *The Guardian*. <https://www.theguardian.com/science/2025/jul/13/quality-of-scientific-papers-questioned-as-academics-overwhelmed-by-the-millions-published>
- Sanders, B. (2025, December 2). AI poses unprecedented threats. Congress must act now. *The Guardian*. <https://www.theguardian.com/commentisfree/2025/dec/02/artificial-intelligence-threats-congress>
- Satariano, A. & Mozur, P. (2025, June 21). The global A.I. divide – where A.I. data centers are located. *The New York Times*. <https://www.nytimes.com/interactive/2025/06/23/technology/ai-computing-global-divide.html>
- Schuessler, J. (2024, December 1). Oxford's 2024 word of the year is... BRAIN ROT. *The New York Times*. <https://www.nytimes.com/2024/12/01/arts/brain-rot-oxford-word.html>
- Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024, July 24). AI models collapse when trained on recursively generated data. *Nature*, 631, 755–759. <https://www.nature.com/articles/s41586-024-07566-y>
- Singh, P. (2025, October 27). AI is driving huge productivity gains for large companies, while small companies get left behind. *CNBC*. <https://www.cnbc.com/2025/10/27/ai-is-driving-huge-productivity-gains-for-large-companies-while-small-companies-get-left-behind.html>
- Sun, L. (2025, May 29). White House MAHA Report may have garbled science by using AI, experts say. *The Washington Post*. <https://www.washingtonpost.com/health/2025/05/29/maha-rfk-jr-ai-garble/>
- Tangermann, V. (2025a, August 6). Experts warn that AI is getting control of nuclear weapons. *Futurism*. <https://futurism.com/the-byte/experts-ai-nuclear-weapons>
- Tangermann, V. (2025b, September 16). Startup using AI to bring Whitney Houston back on tour 13 years after her death. *Futurism*. <https://futurism.com/startup-ai-whitney-houston>
- Tangermann, V. (2025c, October 2). Anthropic safety researchers run into trouble when new model realizes it's being tested. *Futurism*. <https://futurism.com/future-society/anthropic-safety-ai-model-realizes-tested>
- Thompson, D. (2025, October 30). Here's how the AI crash happens. *The Atlantic*. <https://www.theatlantic.com/technology/2025/10/data-centers-ai-crash/684765/>
- Tiku, N. (2025, December 11). A former tech executive killed his mother. Her family says ChatGPT made her a target. *The Washington Post*.

- <https://www.washingtonpost.com/technology/2025/12/11/chatgpt-murder-suicide-soelberg-lawsuit/>
- Tobin, M. (2025, July 16). China is spending billions to become an A.I. superpower. *The New York Times*. <https://www.nytimes.com/2025/07/16/technology/china-ai.html>
- Townsend, C. (2025a, May 15). This OnlyFans model found her photos on Reddit — with someone else's face. *Mashable*. <https://mashable.com/article/reddit-scam-artist-used-ai-to-change-onlyfans-model-photos>
- Townsend, C. (2025b, August 4). An Illinois bill banning AI therapy has been signed into law. *Mashable*. <https://mashable.com/article/illinois-bill-banning-ai-therapy-signed-by-pritzker>
- U.S. Government Accountability Office. (2025, April 22). Artificial intelligence: Generative AI's environmental and human effects (GAO-25-107172). <https://www.gao.gov/assets/gao-25-107172.pdf>
- Valinsky, J. (2025, October 17). OpenAI's Sora bans Martin Luther King Jr. deepfakes after his family complained. *CNN*. <https://www.cnn.com/2025/10/17/tech/openai-mlk-sora-pause>
- VandeHei, J. & Allen, M. (2025, May 28). Behind the curtain: A white-collar bloodbath. *Axios*. <https://www.axios.com/2025/05/28/ai-jobs-white-collar-unemployment-anthropic>
- Veltman, C. (2025, August 26). AI 'deadbots' are persuasive – and researchers say they're primed for monetization. *NPR*. <https://www.npr.org/2025/08/26/nx-s1-5508355/ai-dead-people-chatbots-videos-parkland-court>
- Vigliarolo, B. (2025, November 10). AI slop hits new high as fake country artist goes to #1 on Billboard digital songs chart. *The Register*. [https://www.theregister.com/2025/11/10/ai\\_country\\_artist\\_hits\\_number\\_one/](https://www.theregister.com/2025/11/10/ai_country_artist_hits_number_one/)
- Walker, L. (2023, March 28). Belgian man dies by suicide following exchanges with chatbot. *The Brussels Times*. <https://www.brusselstimes.com/430098/belgian-man-commits-suicide-following-exchanges-with-chatgpt>
- The White House. (2025, July 23). White House unveils America's AI action plan. [whitehouse.gov. https://www.whitehouse.gov/releases/2025/07/white-house-unveils-americas-ai-action-plan/](https://www.whitehouse.gov/releases/2025/07/white-house-unveils-americas-ai-action-plan/)
- Wiggers, K. (2024, August 14). Study suggests that even the best AI models hallucinate a bunch. *TechCrunch*. <https://techcrunch.com/2024/08/14/study-suggests-that-even-the-best-ai-models-hallucinate-a-bunch/>
- Wilkins, J. (2025, June 25). Applying to jobs has become an AI-powered wasteland. *Futurism*. <https://futurism.com/job-applications-ai-slop>
- Yildirim, E. (2025, December 30). AI ruined job applications, so people are resorting to dating apps to find work. *Gizmodo*. <https://gizmodo.com/ai-ruined-job-applications-so-people-are-resorting-to-dating-apps-to-find-work-2000704287>