It's Time StratCom Got Strategic about AI

A Review Essay by Louis Brooke

Empire of AI: Inside the Reckless Race for Total Domination
Karen Hao. London, Allen Lane, 2025.

AI First: The Playbook for a FutureProof Business and Brand
Adam Brotman and Andy Sac. Harvard Business Review Press, 2025.

Keywords—*strategic communications, strategic communication, AI, artificial intelligence, technology, influence operations, transformation, disruption*

About the Author
Louis Brooke is a co-founder and CSO of Zinc Network, an agency that helps solve governance and security challenges, with technology, behaviour change, and communications. Zinc has also launched a new AI tool for StratCom and IO, Kora.

The question of what AI means for individuals, organisations, and societies is one of the most urgent of our time. Investors and businesses are already placing their bets. But the world of strategic communications (StratCom) has been sluggish in its embrace of the technology, both in terms of integrating AI into its practice and in grappling with how AI will reshape the context in which it operates.

StratCom sits between two very different worlds. Its practice is drawn from the dynamic, commercially driven world of marketing and technology, while its funding, institutional structures, and human resource are firmly anchored in the analytical but bureaucratic world of government and public policy. AI is pulling these worlds further apart, and fast. Two books this year speak to these alter egos of StratCom. Karen Hao's *Empire of AI* unpacks the political economy of the AI industry, with a focus on

DOI 10.30966/2018.RIGA.16.7

the social and economic harms it causes, while *AI First: The Playbook for a FutureProof Business and Brand*, by Adam Brotman and Andy Sac, is an urgent plea for marketeers and all business leaders to embrace AI as rapidly and fully as possible, or risk being left behind.

And there is a need for urgency. Marketing has been one of the sectors most rapidly disrupted by and fastest to embrace AI. Copywriters, planners, producers, and designers all are seeing their roles, if not destroyed, transformed. The revenues and valuations of major advertising groups are falling, as smaller agencies can now deliver and distribute quality content at unprecedented pace and scale. Each week a plethora of new AI-powered marketing tools, 'mAItech', is released, propagated by a thriving ecosystem of thought leaders, influencers, and evangelists.

Yet the StratCom community, like much of the public sector, is mainly focused on security and ethics, and all too easily retreats into our comfort zone of talking about big ideas. Will AI bring climate catastrophe, mass unemployment, or even the end of humanity? Or will it usher in a new dawn of a three-day working week, exponential scientific breakthroughs, and universal prosperity? While these questions are clearly important, the facts on the ground are changing before they can be cogently formulated, let alone answered. More importantly, they are also the questions the AI labs want us to focus on, while further entrenching their power. StratCom professionals must become better Marxists. Our job is not only to interpret the world, but to change it.

Move *Very* Fast and Break Lots of Things

For those not fully immersed in the world of AI, who don't listen to the podcasts, follow the Substacks, or experiment with every new tool, it can be hard to grasp how fast the technology is reshaping both our sector and the wider information environment. The sheer pace of deals, product launches, and analysis is overwhelming. Deliberately so.

DOI 10.30966/2018.RIGA.16.7

In October and November 2025 alone, OpenAI has announced plans to spend $ 1.4 trillion on computing power over the next five or so years, roughly the GDP of Spain or Australia. Analysts at Barclays Bank estimate the company is building, or planning to build, data centres providing 46 gigawatts of compute power and consuming 55 gigawatts of electricity, equivalent to the entire electricity consumption of Argentina or South Africa. And that's just one lab.

Major product updates follow almost weekly, some with seismic implications but often buried in the torrent of announcements. In April 2025 OpenAI quietly revealed that ChatGPT would now remember user conversations by default, unless users opted out, purportedly to help it provide more bespoke answers.[1] Yet days later, Sam Altman gave a TED talk envisioning a world where people grow up with lifelong AI companions that become extensions of ourselves that make us more productive and guide us through life.[2] Think of a sycophantic, Silicon Valley version of a dæmon from Philip Pullman's *Dark Materials*. Remembering conversations by default is the first step towards that vision. With minimal public debate or policy scrutiny, the labs are trying to reshape what it means to be human.

The same dynamic is driving the race for wearables. Until now, the major labs have focused on securing enough 'compute'—chips, energy, and data centres—to train their frontier models. But their next constraint is data. The models have already devoured nearly all the available online information, regardless of quality or legality. Meta, for instance, downloaded 81.7 terabytes of pirated books from sites like LibGen in 2024, including millions of copyrighted works. Executives knew the legal risks but correctly calculated that ignoring copyright could simply be treated as a cost of doing business.

1 'Memory and New Controls for ChatGPT', *OpenAI*, 13 February 2024, https://openai.com/index/memory-and-new-controls-for-chatgpt.

2 'OpenAI's Sam Altman Talks ChatGPT, AI Agents and Superintelligence', *TED*, April 2025, www.ted.com/talks/sam_altman_openai_s_sam_altman_talks_chatgpt_ai_agents_and_superintelligence_live_at_ted2025/transcript.

DOI 10.30966/2018.RIGA.16.7

To advance further, AI systems need exponentially more data. Not from the Internet, but from the physical world—the visual, aural, and behavioural data that make up human experience. Hence the rush to develop AI wearables like glasses, watches, and other devices that turn wearers into data harvesters to feed ever-hungrier models. Meta's co-branded Ray-Ban sunglasses were first to market, but other labs are close behind with wristbands, visors, and handheld devices.

The true significance of these developments is obscured by the industry's relentless focus on artificial general intelligence (AGI). Nvidia's Jensen Huang, Meta's Yann LeCun, and AI pioneers such as Yoshua Bengio and Geoffrey Hinton recently claimed 'human-level' AI is already here, or that we have entered the AGI 'spectrum'. This framing keeps the debate polarised between doomsters—warning of climate collapse, mass unemployment, and even extinction—and boosters, who foresee an age of superabundance and scientific breakthrough. That Manichean narrative conveniently hides the present reality that the labs are consolidating capital, infrastructure, and political influence at unprecedented speed, entrenching their oligopolistic power while presenting their actions as necessary to either bring about utopia or prevent dystopia.

Silicon Valley's mantra, 'move fast and break things', has entered warp speed, turbocharged by the tech titan's messianic desire to usher in a new world of AGI, fierce competition between the AI labs to get there first, and floods of capital with nowhere else but overpriced gold to flow into. It is exacerbated by the geopolitical context in which Western governments are reluctant to regulate their national champions for fear of ceding ground to China.

This leaves individuals, organisations, and even nations in the position of eighteenth-century yarn spinners confronting the spinning jenny. In the short term there will be no meaningful regulation of AI or mitigation of its harms. We must act quickly to protect ourselves from the worst impacts, and if we can, ride the wave rather than be swept away by it.

DOI 10.30966/2018.RIGA.16.7

Each new product launch brings us closer to a techno-libertarian vision of society, shaped by a handful of the world's richest men in the white heat of Silicon Valley. The speed of execution is driven not only by competition but also by strategy: to outpace regulation, avoid public scrutiny, and set the terms of debate before anyone else can respond.

The only rational response, for individuals and organisations alike, is to use AI wherever possible to further our own interests. One can be sceptical of its wider social consequences, but in our professional and personal lives, it is time to go all in.

Competition on Mount Olympus

It's the power and relationships of these tech titans that Karen Hao's *Empire of AI* captures most compellingly. She charts the rise of OpenAI from a non-profit dedicated to ensuring AGI benefits all of humanity into the fastest-growing company in history. Hao's central figure, and bête noire, is OpenAI's CEO, Sam Altman. She follows his journey from geeky schoolboy to ambitious entrepreneur and his eventual transformation into a self-anointed messiah (or antichrist, depending on your point of view) heralding the new age of AGI.

Reading it, one is reminded less of corporate boardrooms than of Olympian gods; deities neither good nor evil, but power-obsessed, capricious, jealous, and vain. Their internecine struggles atop the mountain determine peace or war, feast or famine, and even life or death for us mortals below. All we can do is offer tribute, flatter, and hope for mercy.

Hao's quote from a younger Altman lays his grandiosity bare: 'The most successful founders do not set out to create companies,' Altman reflected on his blog in 2013. 'They are on a mission to create something closer

DOI 10.30966/2018.RIGA.16.7

to a religion, and at some point, it turns out that forming a company is the easiest way to do so.'

Elon Musk and Altman originally co-founded OpenAI as a non-profit dedicated to sharing cutting-edge AI research to ensure that AGI benefits all of humanity. Musk and Altman were purportedly united by a shared belief that Google could not be trusted with a technology as powerful as AGI, and that OpenAI therefore had to get there first.

Drawing on 260 interviews, Hao takes us on a blow-by-blow account of Altman and Musk's inevitable battle for control before Musk eventually departs to found xAI (read the underworld), allowing Altman to claim the crown of king of the gods, or OpenAI CEO.

The book charts OpenAI's evolution, exploring different forms of artificial intelligence from bots that could master complex online games, to robotic hands that solved Rubik's cubes, before committing fully to large language models (LLMs). These models derive their apparent intelligence from identifying statistical patterns across vast datasets of text, which essentially enables them to predict what word is most likely to come next in a sequence. The result of this focus on LLMs was ChatGPT, and it was a breakthrough. The chatbot gained 1 million users after just five days and 100 million after two months, and today has over 800 million active monthly users. Retelling this journey offers an important reminder that LLMs are only one form that AI can take. While chatbots offer a compelling mass market offer, there is nothing inevitable about the ability to manipulate natural language being the dominant form of AI.

The highs and lows of the story elucidate the incestuous world of tech titans and big money. Even after the spectacular success of GPT, at the instigation of some of Altman's closest colleagues, OpenAI's not-for-profit board tried to defenestrate him. Some were concerned about his lack of accountability and hazy relationship with the truth, and others genuinely believing they were saving humanity. However, within

DOI 10.30966/2018.RIGA.16.7

hours Altman's allies resigned in protest, hundreds of staff signed an open letter demanding his return, and investors, including Microsoft, OpenAI's main backer, panicked. Within days, he was reinstated. The lesson was clear, AI is capital hungry, and those who can bring in the money, retain control.

Hao also tells the story of Altman making a pilgrimage to a Delphian Bill Gates in search of yet more billions for his loss-making venture. Previously, Gates had been ambivalent about OpenAI's demos. But this time was different. GPT aced his challenge of passing an AP Biology exam, missing only a single question. Gates then posed an emotional question: '*What do you say to a father with a sick child?*' According to Gates, GPT's answer was 'probably better than most of the humans in the room could have given'.

Gates later recalled: 'The whole experience was stunning. I knew I had just seen the most important advance in technology in my lifetime.'

Needless to say, the billions kept flowing from Microsoft.

Scaling Laws: Bigger *Is* Always Better

OpenAI attributes its success to its discovery of 'scaling laws', which purport to show that, for LLMs, performance improves smoothly and predictably when you scale (1) the number of parameters, (2) the amount of data used for training, and (3) compute (read chips) in the right proportions. If scaling laws hold—and to date, they largely have—then whichever lab secures the most data, chips, and resources will be the first to reach AGI. This quest for scale at all costs, to get to AGI first, defines the dynamics of the industry, the trajectory of the technology, and potentially the future of our societies.

DOI 10.30966/2018.RIGA.16.7

Hao argues that scaling laws became not merely a descriptive observation about how model performance improves under current approaches, but an organising ideology for OpenAI and other labs—a mythos that, with enough compute and data (and the capital and natural resource required to secure them), intelligence would emerge. An end that can justify any means. GPT-4, for instance, is estimated to have used tens of thousands of Nvidia A100 chips running for months, a training run costing around $ 100 million in cloud compute alone.[3]

Hao explores how this unrelenting quest for scale and the resources it demands creates vast and inequitable economic and social harms. She argues that these harms are not side effects, but central to the political economy of scaling laws.

She further highlights how AI reshapes labour markets not only for creative and knowledge workers, but also by generating a vast underclass of low-paid, insecure jobs as data labellers in the Global South. As OpenAI pursued scale, it needed ever more data, sucking it in from every swampy backwater of the Internet, even if it was inaccurate, violent, and illegal. The decline in quality of training data gave rise to a network of companies specialising in coding and cleaning that data ready to train LLMs. This work depends on cheap but educated labour in the Global South, where firms pay workers fractions of a cent per task to prepare raw data for model training. These workers have no contracts and no protections, and are simply banned from the task platform when they ask for support or attempt to organise.

Hao also foregrounds the ecological impacts of generative AI, arguing that they are both vast and systematically downplayed by the industry. She traces how the exponential growth in model size demands enormous data centres, consuming staggering amounts of electricity and vast volumes of water for cooling, which strains local ecosystems. It is anticipated that global data-centre power demand will hit ~1000 TWh by 2026,

3 Lance Johnson, 'OpenAI Spent $80M to $100M Training GPT-4', BytePlus, 22 August 2025, www.byteplus.com/en/topic/415209?title=openai-spent-80m-to-100m-training-gpt-4&utm_source=chatgpt.com.

DOI 10.30966/2018.RIGA.16.7

about three times the UK's entire 2023 electricity consumption,[4] while AI's water withdrawals will be 4.2–6.6 billion m³ by 2027, roughly equivalent to the UK's annual public water supply.[5] Crucially, Hao argues that this environmental burden is not evenly shared. The benefits are centralised in wealthy tech hubs, while the costs are often borne by developing countries hosting data centres or supplying the raw materials for chips and servers.

Empires and Robber Barons

The weakest part of Hao's book is its central analogy of AI labs as empires. She vividly describes how, in their quest for scale, the labs centralise power, absorb capital, and consume natural resources and people, particularly in the Global South. Most compellingly, she parallels the ideology of imperialism with the eschatological mythos of the AI labs that portrays them as agents of the end of biological history. But ultimately the empire analogy fails. AI labs make no claim to any of the defining features of a state such as sovereignty, territorial control, or monopoly on violence.

The problem is not merely rhetorical. The empire analogy pushes Hao toward a postcolonial, critical-theory lens that shapes her proposed remedies, leaving them mismatched to the scale of the challenge she lays out. While it is true that data coders are paid pennies and that AI labs site data centres in weakly regulated countries, these are symptoms of mobile global capital, not unique to AI. They are egregious examples of corporate malfeasance, but they do not elucidate the distinctive nature of today's AI problem.

4 Andreas Franke, 'Global Data Center Power Demand to Double by 2030 on AI Surge: IEA', *S&P Global*, 10 April 2025, www.spglobal.com/commodity-insights/en/news-research/latest-news/electric-power/041025-global-data-center-power-demand-to-double-by-2030-on-ai-surge-iea?utm_source=chatgpt.com.

5 Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren, 'Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models', *arXiv*, https://arxiv.org/abs/2304.03271?utm_source=chatgpt.com.

DOI 10.30966/2018.RIGA.16.7

The framing of AI as empire leads Hao to focus on resistance efforts defined by identity and interests, such as local communities in South America protesting to protect water supplies, labour groups defending underpaid Kenyan data workers, and indigenous communities building AI tools to preserve their languages. Laudable, and potentially transformational for some people, as these initiatives are, they do not add up to a coherent response to the oligopoly of AI labs. There will always be cash-strapped governments willing to host data centres despite local opposition, and if one group of data coders unionises, companies simply move to the next cheap labour pool. Hao's proposed solutions, though well intentioned, will not regain social control over this epoch-defining technology.

Rather than empires, a more fitting analogy for the AI age may be the robber barons of the Gilded Age—the plutocrats who captured the railroads, oilfields, and telegraph networks of the second industrial revolution, and with them, the political systems and public spheres of their day. In the nineteenth century, railroads were the connective tissue of the economy; whoever owned the rails controlled trade, mobility, and information flow. In the twenty-first, cloud infrastructure and foundation models play that role. Amazon, Microsoft, Google, and OpenAI now sit astride these new rails, the data pipelines and model weights through which modern knowledge, productivity, and communication increasingly move.

Today's AI magnates are repeating the playbook of the robber barons. They are consolidating control over the essential infrastructure of the century, buying the means of mass communication (Musk's purchase of Twitter now seems well worth his $ 44 bn), and capturing political power in an effort to turn their oligopoly into a fait accompli. The late nineteenth-century robber barons were eventually broken by the reform movement. While this drew on grassroots organising, newly assertive unions, and a press demanding accountability, above all it succeeded because of aggressive antitrust action that reintroduced competition, which also broke the robber barons' grasp on political power.

DOI 10.30966/2018.RIGA.16.7

If power over a transformational technology like AI is concentrated in the hands of a few eccentric men from one small corner of the world, the most plausible remedy is competition. Here there is an opportunity for real change. The AI labs' LLMs are the ultimate generalists, trained for competence across the widest possible range of tasks. This need to be excellent in many domains, drives their insatiable appetite for data, energy, and capital. Investors have piled in for fear of missing out, but it is far from clear that these bets will pay off. Speculation is already mounting about a market correction, with huge volatility and a trillion dollars wiped from AI stocks in the first week of November 2025 alone.

A new generation of smaller, domain-specific competitors is emerging, focused on concrete problems such as health, education, and climate, rather than the AGI moonshot. These models require far less data and compute, and therefore less capital, yet may offer stronger returns. Their rise could rebalance power, restoring diversity and dynamism to an industry now dominated by a handful of giants. Those concerned about the power of AI labs should focus on antitrust regulation, not anti-imperialist movements.

Feeling AI-Anxious?

In contrast to Hao's political analysis, Brotman and Sac's *AI First sits* squarely on the marketing end of the StratCom spectrum. For that reason, many in our sector will find it a less enjoyable read (we read too many books on politics, and too few on business management and marketing). The first half is structured around a series of breathy interviews with tech titans like Altman, Gates, and Reid Hoffman (founder of LinkedIn), providing their perspectives on how AI will shape business, marketing, and branding. The second half is a practical playbook, aimed at executives, urging them not to see AI as just another tool, but to be AI first. That is, to think of AI as a core utility like electricity, underpinning your organisation's strategy, workflows, and culture.

DOI 10.30966/2018.RIGA.16.7

The first chapter is anchored on an interview with the eminently quotable Altman, who claims that in three to five years '95 % of what marketers use agencies, strategists, and creative professionals for today, will easily, nearly instantly and at almost no cost, be handled by AI'.

Given AI's current capabilities and its trajectory (and we must remember that AI is the worst it will ever be now), the spirit of what Altman is saying is surely correct. And if it is true for marketing, it is true for StratCom.

Yet there is a real lag between the speed of technological transformation and the pace at which organisations are adapting. The biggest determinant of whether a business is an early adopter or a laggard is unsurprisingly the posture of its leaders. And most leaders and executives have been slow to adopt AI in their own work lives. The authors quote a 2025 Dresner Advisory Services report that found *88 per cent of senior leaders express interest in generative AI, yet 80 per cent are not regularly using it.*

Surely that stat would be even more damning for StratCom. If you are a leader in StratCom, **and you are not personally AI first, then you should feel anxious.** *AI First*, despite being a clunky read, provides the push, and practical support, to get started.

Three Questions for StratCom to Ask about AI

Empire of AI is a good tour of the AI industry and its darker dynamics. But the brute reality is that AI is here to stay, and due to scaling laws and intense geopolitical competition, it is unlikely to face meaningful regulation. StratCom professionals should avoid getting lost in the doomster vs booster debates amplified by the AI labs, put their marketeer hats on, and focus on how AI can make us cheaper, better, and faster than the competition. Only once we have embraced the technology, can we afford to ask the bigger questions.

DOI 10.30966/2018.RIGA.16.7

1. How can StratCom use AI to be more effective and more efficient?

In a contested information environment, our adversaries are already using AI to flood audiences with tailored content at unprecedented pace and scale. We must harness every available tool to match, and surpass, their speed, precision, and reach. The core functions of StratCom—research, strategy, creative development, production, distribution, in-person activation, and evaluation—have already been disrupted by AI. Fortunately, the marketing world has shown what's possible.

Today's researchers use AI tools not only to analyse quantitative and qualitative data in seconds, but to design and run surveys, scrape and segment public conversations, and even build synthetic audience profiles—digital twins that allow us to test concepts and messages instantly, and at almost no cost.

In strategy development, LLMs can rapidly break down complex problem sets, propose priority objectives, draw on established behavioural science to generate evidence-based communications strategies, and then stress-test them with scenario planning. They can develop messaging frameworks and narratives, and even automatically generate responses to adversary messages. Many companies are now releasing AI agents that can complete these processes autonomously, continuously improving through feedback loops, with minimal human input.

Production is where the transformation is most visible. AI-first production studios are already delivering films and campaigns that would once have cost hundreds of thousands of pounds and months to produce, in a matter of days, for a fraction of the cost. Anyone with a story to tell and a subscription to a tool like Midjourney or Veo 3 can now produce broadcast-quality video, audio, and design assets without leaving their desk.

DOI 10.30966/2018.RIGA.16.7

If the tasks change, the organisations and roles delivering them must change too. Most StratCom structures, particularly those attached to governments, multilaterals, and militaries, were built for an analogue or early digital world. They are bureaucratic, hierarchical, and slow. That model will not survive the AI wave. Already, leading marketing and political campaigning organisations are flattening their structures, relying on smaller teams that combine strategic vision with hands-on AI literacy. In the marketing sector, small teams can now deliver in a few days what might have taken a large agency months to do before gen AI. StratCom teams should be doing the same.

Organisational change also requires cultural change. AI introduces asymmetries of capability; a smart junior staff member with mastery of advanced tools can outperform entire teams. Leaders will need to rethink how to build, contract, and scale teams. There's also a hard truth: not everyone will adapt. As with every technological revolution, some roles will disappear, some will transform, and some entirely new ones will emerge. From the top down, everyone in the organisation must be AI first. Those that can't or won't change will need to go.

Here *AI First* provides some simple steps:

1. Start using AI personally: every leader needs to see and experience the reality of the transformation.

2. Set up an AI working group or council that drives AI implementation across the organisation.

3. Get functional team leaders to start running AI pilots to use AI to improve the quality, quantity, or efficiency of their outputs. Each functional lead should identify the key pain-points in their processes and the weak points in their outputs, and challenge themselves to use AI to deliver five times more impact in the same time without compromising quality or integrity.

DOI 10.30966/2018.RIGA.16.7

4. Work with your operations lead to identify the five key processes in the business whether they are back office, like allocating staff hours, or output focused, like production, and start designing an AI, and ideally agentic-led, version.

5. Change the culture. Make clear that you expect your teams to constantly explore and use the technology and think about how it can be applied. And that using AI isn't a cheat, it's a requirement.

2. How will AI change the information environment?

However, if we focus only on using AI to do better what we already do, we risk missing the fact that AI may change *what we are trying to do entirely*. It would be like perfecting a StratCom minidisc on the eve of the iPod launch.

Imagine we have just discovered a vast new continent in the mid-Atlantic, home to a billion tireless PhD students, willing to work for pennies an hour, and with no barriers to hiring them. How would you go about solving your customer need or mission? It probably wouldn't be a tweaked version of what you are doing today. AI is that continent, and the PhDs are only getting smarter, more skilled, and more numerous. We need to start imagining how this near limitless intellectual resource offered by AI will change the fundamental nature of communication. Three key trends are emerging to which StratCom must adapt: superabundance of content, hyper-personalisation, and centralisation of truth.

Given that AI has already drastically lowered the cost and skill barriers to producing and distributing content, there will be much more *stuff*. We are entering a world of information superabundance. Some of it will be low-quality 'AI gloop'— like the clickbait of rabbits on trampolines[6] that went viral on TikTok and YouTube. But much of it will be indistinguishable in quality from human work, making the

6 'This Viral Video…', Lenny The Bunny, *YouTube*, 9 August 2025, www.youtube.com/shorts/vTPx_muLksc.

DOI 10.30966/2018.RIGA.16.7

information environment vastly denser and more competitive. In such a world, the central challenge shifts from *production* to *discovery*: how do people decide what to pay attention to? The information environment will be defined by the strategies and tools that consumers adopt.

As information supply explodes but human attention remains finite, the role of tastemakers—those who can curate and signal what is good, true, and beautiful—becomes central. As *The Atlantic* argued, 'we are moving from a creator economy to a curator economy'.[7] From influencers on social media to the upper echelons of film, art, and photography, audiences will increasingly rely on the reputation and aesthetic sensibilities of specific individuals to tell them what is important or desirable. In this context, trusting and identifying with these tastemakers become paramount. More than ever, it will be credibility of the source, not the content itself, that matters. This will be reinforced by the fact that deepfakes are becoming indistinguishable from the real thing, as we lose the visual cues that help us judge authenticity. With so much content and no way to tell what is real or not real, those distinctions become less meaningful. Media literacy becomes obsolete. In this world, the StratCom primary means of influence will be building coalitions of tastemakers, not delivering messages.

The other strategy for deciding what content to consume will be to use technology. Soon, we will exist in an information environment in which AI agents will talk to each other more than humans talk to humans. Synthetic content will be produced, distributed, and consumed at scale without direct human involvement. At the same time, users will have AI companions and filters, digital assistants that curate, summarise, and triage information on their behalf. This means much of the information flow will be agent to agent, not human to human. Strategic communicators will need to think less about *content creation* and more about *influencing how information is filtered, prioritised, and surfaced* by networks of billions of agents. This will be a technical

7 Katherine Hu, 'The Influencer Economy Is Warping the American Dream', *The Atlantic*, 18 April 2023, www.theatlantic.com/newsletters/archive/2023/04/social-media-influencers-american-economy/673762.

DOI 10.30966/2018.RIGA.16.7

challenge more than a question of persuasion, meaning the sector needs more technical minds, and fewer creative and strategic thinkers.

AI's capacity for hyper-personalisation is the second trend disrupting the information environment. Hundreds of millions of AI agents constantly generating and iterating content based on live data will enable us to move us from targeting *segments* to a targeting *individuals*. Yet StratCom is still adopting psychographic segmentations, a technique first developed by marketeers in the 1960s. AI is now enabling a world where each person can receive their own individual messages, constantly optimised by AI agents based on their responses.

In an environment where no two people receive the same message, does it even make sense to talk of *a narrative* anymore? Or indeed a shared information environment? We may need to shift from thinking of narratives as arcs to thinking of narratives as fabrics. Each thread is a story for an individual, and strategic communicators must find a coherent weave through the weft and the warp to enable audiences to make sense of a fragmented information environment.

This fragmentation could also accelerate existing trends of polarisation, breaking social media tribes down into smaller and more unruly 'clans' clustered around networks of tastemakers. Recent research has shown how the commercial incentives of cable TV to maximise eyeball hours led it to prioritise negative content on culture war issues, and helped set in train the polarisation of American society.[8] Similarly, social media algorithms have been shown to promote angry and negative content that affirms viewers' existing opinions, measurably contributing to polarising attitudes.[9] Shared public narratives, the foundation of democratic discourse, will become even harder to sustain. Those on the extremes of the debate will have ever more influence to shift the Overton window. StratCom will have to evolve from thinking about how to reach the

8 Aakaash Rao, 'The Business of the Culture War', https://sites.harvard.edu/aakaash-rao/job-market-paper.

9 William J. Brady, Joshua Conrad Jackson, Meriel Doyle, and Silvan Baier, 'Engagement-Based Algorithms Disrupt Human Social Norm Learning', *OSF Preprints*, 13 February 2025, https://osf.io/preprints/osf/mgdwq_v1.

DOI 10.30966/2018.RIGA.16.7

persuadable middle, to being focused on how to contain or influence the most disruptive clans.

The third significant impact of AI on the information environment is epistemology. The first casualty of every communications revolution is truth itself. Our ways of knowing are shaped by the technologies through which we acquire and share information. The printing press shattered the monopoly of Church and Crown over knowledge. Newspapers and broadcast media recentralised authority in professional gatekeepers that mediated knowledge production and dissemination on behalf of large interest groups, often class based. The Internet and social media fractured our shared sense of truth again, by disintermediating the production and consumption of knowledge, and thus eroding consensus and increasing polarisation.

AI is becoming an epistemic technology, not just a productivity tool, but a system for *producing, filtering, and legitimising knowledge*. Whoever controls the major AI models effectively controls how truth and meaning are generated: what people see, what information is prioritised, how facts are framed, and how language itself evolves.

At present that infrastructure is highly concentrated in a handful of AI labs (OpenAI, Anthropic, Google DeepMind, Meta, xAI, etc.). These organisations, and in practice the individuals who lead them, have extraordinary discretion over model training data, alignment goals, and the value systems embedded in responses. That means truth will be shaped by a few private actors, with no accountability mechanism, rather than a plural public sphere.

As the *Brookings Institution* warns, 'control over training data and model design is control over epistemology'.[10] Elon Musk's xAI, for example, has explicitly stated that its Grok model is being trained to remove what he calls 'woke bias'—an ideological intervention in how information is

10 Chinasa T. Okolo, 'Examining the Capabilities and Risks of Advanced AI Systems', 10 September 2024, www.brookings.edu/articles/examining-advanced-ai-capabilities-and-risks.

DOI 10.30966/2018.RIGA.16.7

filtered and framed. It is well documented that Musk frequently tweaks Grok's master prompts in the middle of the night according to his whims. Whose truth, then, do we want? Musk's? Altman's?

However, there is an alternative. If open-source AI develops robustly—models whose weights, training data, and methods are public and reproducible—then the production of knowledge could become radically decentralised, more akin to *Wikipedia* than *Google*. In that case, AI could become a democratising epistemic technology, enabling collective participation in what counts as valid, useful, or true.

In either scenario, for strategic communicators, a central part of any intervention must focus on how to get LLMs to imbibe and propagate our articulation of the truth. Again, this question is just as much technical as it is strategic. Strategic communicators should be more concerned with artificial intelligence optimisation (AIO) than message and content development, as this will be a cornerstone of how we set the epistemic parameters of public discourse.

3. Finally, how will AI change the political, institutional, and moral framework within which StratCom exists?

StratCom is the practice of achieving influence without coercion. What distinguishes it from propaganda is that it operates within the institutional and moral boundaries of liberal democracy, which constrains both its ends and means. StratCom legitimacy depends on truth, consent, and respect for individual agency, and it seeks to promote ends consistent with, and furthering, liberal democracy. Yet AI is transforming the very conditions that make liberal-democratic communication possible by reordering geopolitics, eroding the institutional 'mass middle' that underpins democratic stability, and challenging the humanist assumptions on which notions of agency and persuasion rest.

DOI 10.30966/2018.RIGA.16.7

AI is becoming a *new axis of global power*. The states, corporations, and individuals who control the most capable AI systems will not merely influence the global order; they will help define it. At present this landscape is dominated by the United States and China. Both possess the data, capital, and technical expertise to train competitive LLMs, largely because of the logic of scaling laws. This dynamic reinforces their existing advantages across industry, defence, and influence. Europe and most of the rest of the world are, for now, rule-takers, not rule-makers, in the emerging AI order.

Yet the current *bipolarity* may not last. If large-scale models prove economically unsustainable due to their enormous capital demands, then smaller, more efficient, domain-specific systems may overtake them. Open-source models such as Mistral, DeepSeek, and Falcon have already shown that highly capable systems can be built with modest resources. This could allow other technologically sophisticated polities such as the UK, EU, South Korea, Japan, or Israel to join the race, diffusing AI's productivity and power gains more widely. The shape of the coming order depends on whether AI remains centralised and capital-intensive or becomes distributed and modular. Each scenario carries profound implications for the global information environment, and for the strategic communicators who operate within it.

As argued above, StratCom differs from propaganda precisely because it exists *within* a liberal-democratic framework. It assumes that influence can be exercised through persuasion, transparency, and consent, not coercion or deception, and that truth, evidence, and shared norms provide the common ground for public discourse. That framework is under strain. AI is beginning to hollow the broad, educated middle class that has sustained liberal democracy for over a century.

Since the launch of ChatGPT in late 2022, white-collar hiring in the US and Europe has slowed sharply. The think tank IPPR estimates that up to

DOI 10.30966/2018.RIGA.16.7

8 million jobs are at risk from AI in the UK alone.[11] These are precisely the strata—managers, analysts, administrators, journalists—that have historically mediated between elite decision-makers and the broader public. As these roles shrink, societies risk bifurcating into a narrow elite of AI proprietors, those entrepreneurs who have adopted the technology to make high-revenue, low-head-count businesses, and the highly skilled AI-first knowledge workers, and a wider precariat of those who consume what the machines produce. The weakening of this 'mass middle' also weakens the institutional trust infrastructure that StratCom depends upon for functioning media ecosystems, stable bureaucracies, informed publics, and broad-based political partes. In short, not only is AI transforming how we communicate. It could undermine the *socio-economic conditions* that make liberal-democratic communication possible.

At the deepest level, the liberal-democratic framework rests on a humanist article of faith that individuals are unique moral agents, capable of reasoning, making choices, and pursuing their own ends. StratCom exists within this moral universe. It assumes there are autonomous citizens capable of persuasion, deliberation, and consent. AI destabilises this foundation in two ways.

First, cognitively: as machines acquire abilities that appear creative, reflective, and even empathetic, the boundary between human and artificial intelligence blurs. If an AI can simulate human reasoning and emotion, what remains uniquely human, and what moral claims do other individuals hold over us?

Second, behaviourally: big data and machine learning can now anticipate individual behaviour with high levels of accuracy—our preferences, vulnerabilities, even moral intuitions. As the popular philosopher Yuval Harari puts it, 'Once we begin to count on AI to decide what to study, where to work, and whom to date or even marry, human life will cease

11 IPPR, 'Up to 8 million UK Jobs at Risk from AI unless Government Acts, Finds IPPR', 27 March 2024, www.ippr.org/media-office/up-to-8-million-uk-jobs-at-risk-from-ai-unless-government-acts-finds-ippr#:~:text=This%20would%20also%20impact%20non,GDP%20(%C2%A3306bn%20per%20year).

DOI 10.30966/2018.RIGA.16.7

to be a drama of decision making, and our conception of life will need to change. Democratic elections and free markets might cease to make sense.'[12]

This convergence challenges the Enlightenment assumptions underpinning liberal democracy that individuals are rational, self-knowing agents. It raises uncomfortable questions for our field. If persuasion becomes a matter of micro-targeting cognitive vulnerabilities, are we still practising *strategic communications*, or something closer to *algorithmic manipulation*? And if the 'individual' becomes merely a bundle of predictable data points, what happens to consent, legitimacy, and the moral distinction between communication and coercion?

Get Practical, Then Get Praxis

These are big questions—the type that strategic communicators relish. Yet before we start to answer them, we must first master the immediate, practical ones: *how can we use AI to become better at our craft?* Only by engaging hands-on with the technology, experimenting, learning, and adapting will communicators develop the insight necessary to navigate the deeper ethical and institutional shifts it brings.

We need to radically change our workflows, evolve staff skills, and transform organisational structures to look more like technology firms than civil service departments. We must reconceptualise what it is we do—from creating and disseminating messages to building a narrative arc, to influencing the ways audiences filter and prioritise existing content, so we can orientate them in a fragmented information environment. We must create mechanisms to influence knowledge generation through technology and by building coalitions of tastemakers rather than by persuasion. And we must recognise that influence increasingly lies

12 Yuval Noah Harari, 'Why Technology Favors Tyranny', The Atlantic, October 2018, www.theatlantic.com/magazine/archive/2018/10/yuval-noah-harari-technology-tyranny/568330.

DOI 10.30966/2018.RIGA.16.7

with the margins, not the mainstream. If we don't, AI will change the information environment in ways that render much of our sector obsolete, while we pontificate.