

Technology

What Kind of World Is AI Building?

The ideologies behind today's most powerful technology

By [Jeff de Kleijn & Antoine Fourmont-Fingal](#) | Dec. 15, 2025



(Illustration by [iStock/Cemile Bingol](#))

Discussions of artificial intelligence tend to oscillate between techno-utopianism and existential dread, between simply asking what AI can do or wondering how we might contain it. But there are deeper questions we need to ask: What kind of world is AI helping to build? And who will it serve?

AI is neither a neutral tool nor an inevitable force. As it's currently developed and deployed, it's less a technical system than an ideological project. For this reason, we see little point in arguing *whether* a profound AI-driven sociological change will occur, given speculations about the technical limitations of current architectures. It is the social implications of these limitations that must first be wrestled with.

Yet assuming that AI is inevitable makes it particularly difficult to do so, as Stuart Russell warned in his 2021 Reith Lectures, "[Living With Artificial Intelligence](#)." As he argued, lethal and autonomous weapons are more than merely battlefield innovations: They represent a broader and emblematic crisis of governance without reciprocity, context, or deliberation. But thinking of AI as moving at light speed, impossible to stop, both disorients public agency and makes it hard even to ask whether this is something anyone wants. If the pace is predetermined, politics is merely a performance rather than genuine democratic participation.

By the same token, framing AI as a *neutral* technology withdraws it from political scrutiny just as effectively, masking moral choices as if they were purely procedural or technical. And just as markets

often outsource questions of value to mechanisms of price, AI outsources those moral and social questions to the problem of optimization. Yet AI's apparent neutrality is not the absence of values but the dominance of hidden ones embedded in AI systems that reflect particular priorities, however much they present themselves as objective. As Michael Sandel notes in his 2009 *Justice: What's the Right Thing to Do?*, the attempt to avoid explicit moral reasoning often results in importing significant moral choices without public scrutiny. When AI inherits this posture, it reframes questions about fairness, harm, and recognition as matters of code and metrics, insulating them from democratic scrutiny. And just as technology such as AI is profoundly shaped by the societies that produce it, it, in turn, shapes those societies by what it optimizes, whom it empowers, and what forms of life it renders thinkable. What we are witnessing is not the rise of machine intelligence but the consolidation of its extractive logic, under the mask of innovation.

Another frame to re-think is the “general” aspect of AI, the way, consciously or not, a variety of distinct technologies blur together as “AI” (if only because, in the case of “AGI” or “super AI,” there are so many shifting definitions that you'd be hard pressed to do anything else). However, the more general the technology is, the harder it is to regulate or even to appreciate it, and the easier it is simply to presume its inevitability.

As Russell puts it, the moment we start to contextualize and open the way to asking very clear, palpable questions, it becomes possible to see (and scrutinize) the changes AI is actually bringing.

A World of Scale

AI demands a massive concentration of capital, data, computing, and control, something often justified by reference to economies of scale, a concept inherited from industrial production. The assumption is that efficiency increases with scale, such that only monopolies or near-monopolies are considered viable. But this logic is flawed when applied to digital systems. Not all inputs scale evenly: Coordination costs, context sensitivity, and adaptation to local variation often increase with scale. What the current AI industry reflects is not the natural outcome of efficiency but a strategic convergence, based on the belief that a few universal solutions are preferable to many situated ones. In this sense, it mirrors a mindset in which a single, dominant model for all cancers is considered more desirable than a portfolio of specific treatments that respect variation. This logic rewards centralization even when the underlying conditions call for diversity.

The ideology of scale is not unique to AI. However, AI is being developed by a group of companies that already have a highly concentrated global footprint. You'd be hard-pressed, after all, to find tools with as

much of a global footprint as Excel, WhatsApp, or Google Maps. How is it possible that we all use those limited solutions? Whom is it better for? Are those widely successful tools successful because they're better or because they have been imposed?

After all, as of 2025, the landscape of frontier AI remains heavily concentrated in the hands of a few corporations. Microsoft, Google, Amazon, and Meta collectively account for the majority of large-scale AI model development, data infrastructure, and training runs. (According to the [2025 AI Index](#), over 70 percent of the most computationally intensive models were trained within these companies or their direct affiliates.) This concentration extends beyond model development into access to computing resources, shaping who gets to experiment, iterate, and scale.

However, open-source initiatives exist, and their viability demonstrates that AI development does not require a monopolistic approach. These models indicate that we can sustain both small- and large-scale approaches, public and private ecosystems, and centralized and federated infrastructures. However, their influence remains constrained by deeper infrastructural asymmetries, particularly in data hosting and access to large-scale computing. These hidden layers create a capture loop that underpins broader power consolidation and narrows the horizon for plural AI development.

In this sense, the logic of scale is not neutral. It is a design principle that consolidates decision-making power and constrains diversity.

A World of Coherence Over Intelligibility

AI prioritizes coherence over intelligibility. Being correlation-based, they optimize for performance metrics like accuracy, efficiency, and fluency (while eroding common sense, contextual reasoning, and interpretive depth). As a result, GPT-based models may outperform humans in benchmark tests, yet fail basic tests of moral reasoning or ambiguity management. Already, in legal and educational settings, there is pressure to adjust to AI by shifting human roles from deliberation to prompt design.

This inversion matters, however: The capacity to interpret, to disagree, and to withhold judgment is essential to any pluralistic society. But the consequences reach beyond cognition. When systems reward simulation over understanding, meaning begins to erode, and when meaning collapses, so does the self. As Ma Jian reflected in *Red Dust*, on the breakdown of marriage after the post-Mao reforms, when people have no sense of self, relationships are merely temporary distractions from inner emptiness and tend to fall apart at the first obstacle. Ma's was not simply an observation about intimacy but about the dissolution of shared meaning as relationships lose depth, held together not by recognition but by convenience. A field of hollow ties emerges, and with it, a society of displaced selves.

From Massification to Atomization

José Ortega y Gasset once warned of the “mass man,” untethered from inner discipline and reliant on external direction (in 1929’s *The Revolt of the Masses*), an insight Hannah Arendt deepened by arguing that such atomization is essential to domination and that a society of unanchored individuals is easier to manage (*The Origins of Totalitarianism*, 1951). In the same vein, Günther Anders described “the obsolescence” of human beings as we struggle to fit ourselves to the machine rather than the other way around (*The Outdatedness of Human Beings*, 1956 and 1980).

When AI rewards performance over presence, it risks accelerating this drift toward massification. Rather than asking what AI can do, we must ask *what systems AI amplifies*. If AI is inherently an amplifier of data, decisions, and institutional structures, then its most urgent dangers are systemic rather than technical in nature. AI reflects the logic of its deployment. Embedded in extractive systems, it scales extraction. Fully enclosed in economics-thinking, it will be blind to the same qualitative signals and produce the same unaccountability sinks that Dan Davies so vividly describes in 2024’s *The Unaccountability Machine*.

However, if positioned within civic architectures, could AI scale deliberation, reflection, and care?

Designing Plurality, Direction, and Civic Imagination

Audrey Tang, Taiwan’s Digital Minister, has **stated** that democracy itself is a form of social technology. Like any technology, it can be iterated, refactored, and improved. But suppose AI is to serve democratic evolution. In that case, it must be governed not as an object of innovation, not as a neutral, inevitable, and unquestionable advance, but as a straightforward matter of infrastructural choice.

Such an alternative horizon does not require a blueprint; only directions and principles are needed. Decentralization, for instance, refuses the default assumption that scale must imply monopoly. Federated learning, data trusts, and publicly owned infrastructures, including the initiatives now emerging across Europe, offer plural alternatives to centralized dominance.

Another approach is to emphasize structural coupling over replacement, designing systems in which human judgment and machine precision co-evolve. The Estonian X-Road platform is a notable example, integrating automated services while maintaining public oversight and legal clarity.

A third is plurality, encouraging ambiguity and cultural specificity, and leveraging AI to make, for example, the richness of rare languages in plant naming widely available. In Rwanda, for example,

participatory planning tools have been integrated into rural land management systems to accommodate friction rather than erase it. We must also return to interpretive education, reinforcing liberal arts and civic inquiry over mere technical proficiency. At Olin College and other experimental institutions, AI is introduced alongside ethics, phenomenology, and systems theory rather than in isolation.

Toward Institutional Intelligibility

Shaping the direction of AI's development cannot be the prerogative of engineers and CEOs alone. Citizens have a role not simply in using machine learning but in contesting the criteria by which systems are legitimized. In a world of scoring systems and opaque algorithms, their legitimacy must be understood as a civic concern.

Doing this means building decentralized grammars of AI. These must be participatory, open-ended, and friction-rich. There must be spaces where AI is not a substitute for judgment but a partner in shaping meaning. The civic technologist movement, from Taiwan to Barcelona to Mexico City, demonstrates that such grammars are already emerging.

The next chapter in AI will not be determined by scale but by sense-making. The task ahead is not to forecast innovation but to shape the institutions through which meaning is made. That responsibility does not rest solely with engineers or executives. It belongs equally to educators, artists, jurists, civic designers, and to all who care about how power is structured and shared. The future of AI must be authored in classrooms, libraries, courts, and town halls, not just in labs and boardrooms. The code matters, but the criteria matter more. We must insist on plurality, transparency, and democratic oversight before today's defaults calcify into tomorrow's systems. It is time to stop optimizing the hallucination and start designing the house.

Read more stories by [Jeff de Kleijn](#) & [Antoine Fourmont-Fingal](#).



Jeff de Kleijn is a senior executive and author specializing in digital systems, institutional design, and civic governance.



Antoine Fourmont-Fingal is a global product and growth strategist with senior leadership experience across e-commerce, fintech, and AI deployment in Asia, Europe, and emerging markets.

DOI: 10.48558/hfv-mt59

If you like this article enough to print it, be sure to subscribe to SSIR!

Copyright © 2026 Stanford University.
Designed by Arsenal, developed by Hop Studios