SBS SWISS BUSINESS SCHOOL

Artificial Intelligence After Scaling: Why Slowdown May Signal Maturity, Not Failure

By Prof. Dr. Bert Wolfs, SBS Swiss Business School.

The debate around artificial intelligence has entered a more reflective phase. Cognitive scientist Gary Marcus and investor Steve Eisman, approaching the subject from different intellectual and professional backgrounds, have both suggested that improvements in current AI models may soon reach a plateau. Their shared skepticism challenges a dominant narrative that assumes AI will continue to advance rapidly through sheer scale.

For academic practitioners, those operating at the intersection of research, leadership, and organizational decision-making, this moment calls for analysis rather than anxiety.

Over the past decade, AI progress has been driven largely by scaling, marked by collecting more data, developing larger models, and increasing computational power. This approach has produced striking results, particularly in natural language processing and multimodal generation. Empirical research on scaling laws shows predictable performance gains as model size and data volume increase (Kaplan et al., 2020). However, the same literature also highlights diminishing returns, rising costs, and practical constraints.

Marcus's critique addresses a deeper issue. He argues that current AI systems, despite their fluency, lack core elements of human cognition, including causal reasoning, symbolic abstraction, and stable world models (Marcus, 2020; Marcus & Davis, 2019). From a cognitive science perspective, this criticism aligns with long-standing debates about the limits of purely statistical learning. Systems trained primarily to predict the next token may succeed at imitation without achieving understanding, a distinction that has been emphasized for decades in the philosophy of mind and AI research.

These limitations are not merely theoretical. Persistent problems, such as hallucinated facts, logical inconsistencies, and weak generalization in unfamiliar contexts, suggest architectural constraints rather than temporary technical issues. As Marcus and others have argued, scaling alone does not resolve these issues; it may merely mask them under increasingly sophisticated outputs.

Eisman's skepticism approaches the question from an economic angle, aligning with academic concerns in recent studies. Training and operating large AI models requires substantial capital investment in hardware, energy, and infrastructure. At the same time, evidence of broad-based productivity gains remains mixed. Research on previous general-purpose technologies, ranging from electricity to information technology, shows that the economic impact often lags technological invention by years or even decades because of organizational and institutional frictions (Brynjolfsson, Rock, & Syverson, 2021).

From this perspective, a slowdown in model-level improvements would not be surprising. It may indicate that the current paradigm is approaching diminishing marginal returns. Importantly, this does not imply stagnation of AI as a socio-technical system.

Historically, periods of rapid technological advance are often followed by phases of consolidation and integration. In AI, this shift is already visible. The most consequential developments today occur not in marginal benchmark improvements, but in system-level design: the integration of models with tools, workflows, governance mechanisms, and human oversight.

For academic institutions, businesses, and public-sector organizations, this distinction matters. The value of AI increasingly depends on organizational rather than computational questions: When should AI be used? Under what constraints? What accountability structures? These concerns align with long-standing work in socio-technical systems theory, which emphasizes that technology outcomes are shaped as much by institutional design as by technical capabilities.

The implications for education and leadership development are equally significant. The growing need is not only for technical specialists but also for professionals who understand both AI's limitations and its strengths, individuals capable of critical evaluation, ethical judgment, and strategic deployment. This aligns with emerging research on responsible AI and human-centered design, which emphasizes transparency, robustness, and context-awareness over raw performance.

In this light, skepticism about the endless acceleration of AI may be constructive. It encourages a shift from speculative narratives about artificial general intelligence toward grounded discussions of reliability, governance, and value creation. It also brings expectations closer to what organizations can realistically absorb and manage.

Marcus and Eisman may well be correct that the era of easy gains from scaling is nearing its end. Where their view deserves refinement is in how to interpret what comes next. For academic practitioners, the more compelling conclusion is that artificial intelligence is entering a phase of maturity, in which progress becomes slower, more incremental, and more tightly coupled to institutional capacity.

In this sense, a plateau is not a failure. It is often the point at which technology stops trying to impress and begins the harder work of becoming dependable.

References

Brynjolfsson E., Rock D., & Syverson C. (2021). "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," American Economic Journal: Macroeconomics, American Economic Association, vol. 13(1), pages 333–372, January.

Marcus, G. (2020). *The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence*. arXiv preprint.

Marcus, G., & Davis, E. (2019). *Rebooting AI: Building Artificial Intelligence, We Can Trust*. Pantheon Books.

Kaplan, J. et al. (2020). *Scaling laws for neural language models*. arXiv preprint.