Don't Imitate, Innovate: The Case for Useful AI

JP Vasseur, Sr Distinguished Engineer
Email: jvasseur@nvidia.com
September 30, 2025

Executive Abstract

For too long, Artificial Intelligence has pursued the wrong goal: building machines that think like humans and judging them by their proximity to human intelligence, rather than their usefulness. This paper argues that this anthropocentric goal is a strategic error. The true value of AI lies not in imitation, but in innovation that leads to usefulness. Intelligence is not a singular, human-like essence, but a plural and contextual capability. The obsession with mimicking human cognition has distracted from what really matters: building systems that solve urgent problems and open new frontiers.

It is time to change our metrics for success. While building them requires heavy engineering, the benchmarks we need are not Turing Tests, but measures of real-world contribution: scientific discoveries accelerated, diseases detected earlier, energy systems optimized, quality of code produced, complex autonomous systems built, to name a few. We must stop asking, "How close is AI to thinking like a person?" and instead ask, "What new capabilities can we create to augment human potential?" The goal is not to imitate minds, but to build tools that change the world—through systems that are not human-like, but profoundly, undeniably useful.

Introduction

The field of Artificial Intelligence stands at a crossroads. One path leads toward creating mirrors of human intelligence, driven by the long-held dream of a machine that thinks like a person. The other path leads toward creating tools: systems that possess novel forms of intelligence designed to solve problems beyond human capacity. This paper argues for the latter, making the case that the future of AI does not lie in mimicry, but in functional, targeted, and ultimately useful innovation.

1. Intelligence: The Myth of a Singular Measure

1.1. Intelligence, General Intelligence, and Superintelligence

The concept of intelligence has long been the subject of extensive debate across disciplines such as psychology, philosophy, neuroscience, and computer science. Despite over a century of research, there is still no universally accepted definition of intelligence. In the context of artificial intelligence, this lack of consensus has profound implications. When we speak of Artificial General Intelligence (AGI) we risk importing flawed assumptions and overly narrow definitions of what it means to be intelligent. This section unpacks the multidimensional nature of intelligence, critiques anthropocentric metrics, and argues that our framing of intelligence deeply shapes the goals and (potential) limitations of AI development.

1.2. The Multi-Dimensional Nature of Intelligence

The purpose of this paper is not to resolve the long-standing debate over the definition of intelligence. Instead, it acknowledges that intelligence is not a monolith. Decades of research in psychology and biology support this pluralistic view, from Howard Gardner's theory of multiple intelligences to the diverse, domain-specific abilities seen in the animal kingdom (Gardner, 1983; Shettleworth, 2010). Intelligence is better understood as a spectrum of cognitive skills that vary by context, species, and task.

This diversity complicates any effort to encode or emulate intelligence in machines. Most contemporary AI systems excel in narrow tasks but falter when expected to generalize across contexts. That failure is often viewed as a deficit when compared to human cognition—but such a view presumes that human-like generalization is the gold standard. Instead, we should ask: is this even the right benchmark?

1.3. Intelligence as a Human-Centric Illusion

Artificial General Intelligence is often framed as an endpoint: a system that can perform any intellectual task that a human can. But this framing assumes that the spectrum of human cognition represents the pinnacle—or at least the template—of what intelligence should be. This anthropocentric bias not only limits innovation but also misrepresents the nature of intelligence itself. AGI becomes a mirror we hold up to ourselves, rather than a tool to explore forms of cognition unlike our own.

A chess-playing AI is not inferior because it cannot fold laundry nor is a cat less intelligent because it cannot play chess. In this light, the pursuit of AGI as abstract ideals may be less important than the design of systems that are useful, complementary, and functionally aligned with human needs.

1.4. Toward a Functional Framing

If we accept that intelligence is inherently plural and situated, then it follows that AI should not be judged by how well it mimics humans, but by how effectively it performs in specific roles. The next sections argue for a functional, task-oriented approach to AI that prioritizes usefulness over imitation. Rather than chasing an ill-defined notion of generality, we should build systems that excel precisely because they are not like us. The cost of chasing the mirage of human-like AI is significant, diverting research attention and funding from efforts that could yield measurable, real-world benefits.

In short: the question is not whether machines can think like humans, but whether they can do things that matter.

2. From Analogy to Application: A Survey of Biologically-Inspired AI

Despite the critique of anthropocentric metrics, it is essential to acknowledge that drawing inspiration from biology has led to valuable innovations in AI. To be clear, mimicking human cognition is not in my opinion inherently flawed—it remains a scientifically fruitful approach when framed correctly. This section provides a few examples of ongoing efforts to replicate and integrate principles from neural and biological systems, highlights recent advancements, and critically examines their benefits and limitations as a means to an end, not the end itself.

2.1. Neuromorphic Computing: Hardware Inspired by the Brain

Neuromorphic computing exemplifies a bio-inspired strategy, designing hardware that emulates spiking neural networks to enable real-time, energy-efficient computation. The field includes a wide range of prominent initiatives:

- **Major Technology Firms:** Intel's Loihi 2 and IBM's TrueNorth are foundational research chips demonstrating on-chip learning and low-power pattern recognition.
- **Large-Scale Academic Projects:** The SpiNNaker (University of Manchester) and BrainScaleS (Heidelberg University) platforms focus on large-scale brain emulation and modeling cortical dynamics.
- **Specialized Commercial Hardware:** Companies like Prophesee, SynSense, and Innatera are developing low-power neuromorphic sensors and accelerators for edge AI applications like robotics and real-time sensor fusion.
- **Research Institutions & Consortia:** Stanford's Neurogrid and work at national labs like LANL and universities like the University at Buffalo continue to push the boundaries of brain-like efficiency and complex processing.

While still an emerging field, research in neuromorphic computing continues to advance, with some systems scaling to over one billion neurons for specialized tasks like edge computing and sensor fusion (Quantum Zeitgeist, 2025). This approach shows promise for energy efficiency, and novel materials mimicking synaptic behavior are being developed with the goal of overcoming traditional von Neumann bottlenecks (Wang et al., 2025).

2.2. Bio-Inspired Models and Algorithms

Beyond direct hardware emulation, another strand of AI draws inspiration from higher-level biological principles like evolution and sensory learning.

- **Evolutionary and Swarm Intelligence:** Sakana AI, for instance, has pioneered evolutionary algorithms for model merging. In 2025, it released TreeQuest, an open-source system enabling LLMs to collaboratively solve problems, and AB-MCTS, a bio-inspired algorithm for multi-agent coordination (Sakana AI, 2025a; 2025b).
- **Models of Biological Learning:** Research at institutions like Cornell University focuses on modeling sensory learning for adaptive robotics, while UCLA investigates neural synchronization patterns to bridge biological and artificial systems.
- **Physics-Based Approaches:** Companies like Extropic.ai apply principles from thermodynamics and reversible computing to develop energy-efficient, bio-mimetic AI, demonstrating that inspiration can come from the fundamental physics of biological systems, not just their emergent structure.
- **Hybrid Systems: Integrating Biological and Digital Computation:** The most direct form of mimicry involves integrating living neurons with digital architectures. This frontier is exemplified by the highly experimental field of **Organoid Intelligence**, which uses lab-grown brain organoids for computation. A notable example is the work from Cortical Labs with its CL1, a biological computer embedding up to 800,000 human neurons for applications in AI model exploration and pharmacology (Kasianov, 2025; New Atlas, 2025). These platforms, while still in early research stages, offer unique access to learning dynamics not currently reproducible in silicon-only systems (Smirnova et al., 2023).

2.3. A Critical Assessment: The Benefits and Limits of Mimicry

The overarching goal remains: to design systems that are useful—not necessarily human-like—in solving problems beyond our current capacity. The next section explores the implications of embodiment and animal cognition in further challenging the anthropocentric framing of intelligence.

3. The Moravec's Paradox and the Cat

A key idea in AI, **Moravec's Paradox**, notes that AI excels at tasks humans find hard, like complex math, but struggles with what we find easy, like basic movement and perception (Moravec, 1988). This isn't a failure of abstract thought, but a gap in physical competence.

Nowhere is this paradox more apparent than in what could be called the **"Cat Paradox."** A common house cat, with no formal training, can effortlessly navigate a cluttered room, hunt, and adapt to unpredictable physical conditions. These behaviors—which rely on real-time sensorimotor integration and embodied learning—remain extraordinarily challenging for even the most advanced AI and robotic systems. The cat's abilities reveal a deeper truth: intelligence is not merely a disembodied process of symbolic manipulation but is profoundly **embodied**—it emerges from the dynamic interaction between an agent, its body, and its environment.

This embodiment gap has direct implications for AI's real-world usefulness. For AI to move beyond digital screens and operate effectively in the physical world—in robotics, logistics, and scientific exploration—it must be able to perceive, act, and learn within it.

The key takeaway is not that AI should be built to replicate a cat, but that the paradox is a crucial reminder of what is missing. This conceptual shift is already underway, gaining momentum under the banner of **Physical AI**. This emerging field aims to unify large-scale models with embodied agents that learn directly through interaction with the world, often trained in rich simulations before deployment. Initiatives like NVIDIA's GR00T project for humanoid robots, Google DeepMind's work on visuomotor policies, and Meta's Habitat simulation environment are at the forefront of this effort. By learning from the principles of embodied experience, the field is beginning to bridge the gap—shifting the focus from creating systems that can *think* to creating systems that can *do*.

4. Usefulness: The Only Metric That Matters

The previous sections have argued that intelligence is multifaceted and that the pursuit of human-like AGI is a flawed goal. It follows, then, that **usefulness**, not resemblance to human cognition, should be the guiding metric for

AI development. The value of this approach is already clear. AI is uniquely positioned to process vast amounts of data to find patterns invisible to the human eye, perform complex reasoning to solve intractable scientific problems, and accelerate discovery at an unprecedented scale. We see this in biology, where DeepMind's **AlphaFold** has revolutionized protein folding, a problem that stumped scientists for decades. We see it in healthcare, where AI assists in diagnostics and drug discovery, and in materials science, where systems like Google DeepMind's GNoME have discovered hundreds of thousands of new stable materials, dramatically accelerating the design of novel technologies (Merchant et al., 2023). These are not incremental improvements; they are transformative capabilities that arise precisely because AI does not 'think' like a human.

4.1. Functional Intelligence Across Domains

The fixation on mimicking human cognitive traits—language fluency, reasoning patterns, emotional resonance—obscures a fundamental truth: AI does not need to be human-like to be valuable. In fact, its greatest contributions may arise precisely from its differences. Machines can ingest, process, and analyze terabytes of multimodal data in real time, perform massively parallel simulations, or discover correlations invisible to human intuition.

These capabilities are already extending human potential across numerous fields. In medicine, AI now assists in everything from early cancer detection and radiological triage to protein design and personalized treatment planning. In climate science, AI enables faster, more granular modeling of atmospheric dynamics. In software development, AI copilots and agentic systems generate, debug, and optimize complex code, acting as powerful force multipliers. Even highly specialized domains are being transformed; in neuroscience, AI tools are used to reverse-engineer brain circuits, while in agriculture, AI-driven systems manage autonomous farming machinery and optimize irrigation.

Perhaps one of the most powerful examples of this paradigm is using AI to understand the brain itself. Rather than trying to build a synthetic brain, researchers are using AI as a tool to model and test hypotheses about neural computation. A leading example is the work of Professor Martin Schrimpf and his colleagues, who developed **Brain-Score**, a platform that quantifies how well artificial neural networks predict neural and behavioral data from the primate brain (Schrimpf et al., 2020). This approach treats AI models not as aspiring replicas of the brain, but as scientific instruments for understanding it—a perfect embodiment of prioritizing usefulness over mimicry.

As yet another illustration, AI is becoming the engine of 'AI factories': massive data centers built specifically for training and inference, which will serve this vast array of use cases.

4.2. A Platform-Level Case Study: The NVIDIA Ecosystem

To see this principle in action at scale, it is useful to move beyond individual applications and examine an entire industrial ecosystem built on the philosophy of augmenting human capabilities. The portfolio of a single company like NVIDIA serves as a compelling case study for this utility-driven approach, with tools designed to solve specific, high-impact problems across the full stack of AI development:

- CUDA has become the de facto programming interface for high-performance GPU computing, enabling breakthroughs in deep learning, simulation, and scientific computing.
- NVIDIA DGX systems provide enterprise-grade AI supercomputing for model training and scientific research, drastically reducing time to insight.
- Triton Inference Server allows scalable, real-time model serving with support for multiple frameworks, helping operationalize AI in production environments.
- NVIDIA Clara supports AI-powered medical imaging and genomics workflows, assisting clinicians in diagnostics and treatment planning.
- NVIDIA Modulus enables physics-informed machine learning for real-time simulations in engineering and climate modeling.
- NVIDIA Omniverse integrates AI with simulation, digital twins, and robotics development—supporting complex industrial automation pipelines.
- NVIDIA cuQuantum accelerates quantum computing research by simulating quantum circuits on classical hardware, speeding up algorithm prototyping.

- GR00T and Isaac Sim offer embodied AI platforms that merge large models with physical robotics, paving the way for generalist agents in the real world.

These examples reflect a consistent focus on amplifying human capability rather than mimicking it.

4.3. Task-Alignment as a Design Principle

From an engineering perspective, usefulness can be framed as task-alignment: how well does an AI system perform relative to the goals defined by its operational context? This includes:

- Accuracy and efficiency
- Robustness to uncertainty
- Adaptability to new data
- Scalability across environments
- Interoperability with human workflows
- Consistency and Reliability
- Trustworthiness

These are practical measures tied to real-world engineering and business results. Unlike philosophical AGI benchmarks that ask 'is it intelligent?', this approach asks a simpler, more important question: 'does it work?'. Task-alignment is about engineering valuable systems.

5. Conclusion: Don't Imitate — Innovate

Artificial Intelligence is not—and should never be—about building synthetic replicas of human minds. For too long, the field has been oriented toward an anthropocentric fantasy: constructing machines that think like us, talk like us, and eventually surpass us. This narrative, centered on AGI is conceptually flawed, scientifically vague, and strategically limiting.

The argument of this paper is that the value of AI lies in useful innovation, not imitation. While defining and measuring usefulness requires significant engineering effort, it reorients the field toward what truly matters. The obsession with benchmarking machines against human capabilities has been a distraction from the real goal: building systems that extend our reach, solve urgent problems, and open new frontiers.

We have shown that useful AI is already transforming science, medicine, climate research, engineering, and creativity—not by thinking like humans, but by doing what humans cannot.

This is not about redefining intelligence—it is about liberating it from outdated constraints. The goal of AI should not be to pass as a person, but to perform as a powerful collaborator. The benchmarks we need are not Turing Tests, but measures of contribution: scientific discovery accelerated, diseases detected earlier, energy systems optimized, lives improved.

As researchers, engineers, and policymakers, we must stop asking: How close are we to replicating human cognition? and start asking: What new capabilities can we create, how can they be aligned with critical tasks, and for whom will they deliver measurable value?
Let us not chase minds. Let us build tools that change the world—through systems that are not human-like, but profoundly, undeniably useful.

6. Disclaimer

The views and opinions expressed in this paper are solely those of the author and do not necessarily reflect the official policy or position of NVIDIA or any of its affiliates. The information and analysis presented are based exclusively on publicly available data, documentation, and reports accessible on the web.

References

Fortune Business Insights. (2025). *Neuromorphic Computing Market Size | Global Report [2032]*. Retrieved from https://www.fortunebusinessinsights.com/neuromorphic-computing-market-108240

Gardner, H. (1983). *Frames of Mind: The Theory of Multiple Intelligences*. Basic Books.

Gottfredson, L. S. (1997). Mainstream science on intelligence: An editorial with 52 signatories, history, and bibliography. *Intelligence, 24*(1), 13–23.

Kasianov, R. (2025, March 4). Cortical Labs Launches $35K Biological Computer Built on Human Brain Cells. *BiopharmaTrend*. Retrieved from https://www.biopharmatrend.com/news/cortical-labs-introduces-biological-computer-built-on-human-brain-cells-1156/

Merchant, A., Batzner, S., Schoenholz, S. S., Aykol, M., Cheon, G., & Cubuk, E. D. (2023). Scaling deep learning for materials discovery. *Nature, 624*(7990), 80-85.

Moravec, H. (1988). *Mind Children: The Future of Robot and Human Intelligence*. Harvard University Press.

New Atlas. (2025, March 3). World's first "Synthetic Biological Intelligence" runs on living human cells. *New Atlas*. Retrieved from https://newatlas.com/brain/cortical-bioengineered-intelligence/

Quantum Zeitgeist. (2025, January 27). Neuromorphic Computing Advances. *Quantum Zeitgeist*. Retrieved from https://quantumzeitgeist.com/neuromorphic-computing-advances/

Russell, S. (2019). *Human Compatible: Artificial Intelligence and the Problem of Control*. Viking.

Sakana AI. (2025a, July 7). Sakana AI's TreeQuest: The Monte Carlo Revolution That's Redefining Enterprise RAG Orchestration. *RAGabout.it*. Retrieved from https://ragaboutit.com/sakana-ais-treequest-the-monte-carlo-revolution-thats-redefining-enterprise-rag-orchestration/

Sakana AI. (2025b, July 2). Multi-LLM AB-MCTS. *Sakana AI Blog*. Retrieved from https://www.sakana.ai/multi-llm-ab-mcts/

Schrimpf, M., Kubilius, J., Lee, M. J., Murty, N. A. R., Ajemian, R., & DiCarlo, J. J. (2020). Integrative Benchmarking to Advance Neurally Mechanistic Models of Human Intelligence. *Neuron, 108*(3), 413-423.

Shettleworth, S. J. (2010). *Cognition, Evolution, and Behavior*. Oxford University Press.

Smirnova, L., et al. (2023, February 27). Organoid intelligence (OI): the new frontier in biocomputing and intelligence-in-a-dish. *Frontiers in Science, 1*. Retrieved from https://www.frontiersin.org/journals/science/articles/10.3389/fsci.2023.1017235/full

Sternberg, R. J. (1985). *Beyond IQ: A Triarchic Theory of Human Intelligence*. Cambridge University Press.

Turing, A. M. (1950). Computing Machinery and Intelligence. *Mind, 59*(236), 433-460.

Wang, A., Pruksachatkun, Y., Nangia, N., Singh, A., Michael, J., Hill, F., Levy, O., & Bowman, S. R. (2019). SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems. *arXiv preprint arXiv:1905.00537*. Retrieved from https://arxiv.org/abs/1905.00537

Wang, Z., et al. (2025, July 31). Self-rectifying memristors based on epitaxial AlScN for neuromorphic computing. *Applied Physics Letters, 127*(4). Retrieved from https://pubs.aip.org/aip/apl/article/127/4/044105/3357168/Self-rectifying-memristors-based-on-epitaxial