AI & THE PHYSICAL WORLD · PERSPECTIVE

NVIDIA Says AGI Has Arrived. But Has AI Passed the Physical-World Test?

Digital intelligence has advanced at extraordinary speed. But if we want to call AI truly general, should it also be able to understand, reason and act reliably in the physical world?

Nevzat Atakli
CEO and Co-Founder, Trio Mobil

September 2026 · 6 min read

In March 2026, NVIDIA founder and CEO Jensen Huang made a provocative statement during an episode of the Lex Fridman Podcast: **"I think we've achieved AGI."**

The comment came in response to an unconventional benchmark proposed during the conversation: whether an AI could create and grow a technology business to a valuation of more than $1 billion. By that definition, Huang argued that the threshold for artificial general intelligence may already have been reached.

It is an interesting test of digital intelligence. But is it a sufficient test of general intelligence?

That may say something important about digital intelligence. But it does not demonstrate that AI has mastered general intelligence.

Digital Intelligence Has Become Remarkably General

There is little question that the capabilities of today's AI systems would have seemed extraordinary only a few years ago.

Large AI models can write software, analyze complex documents, generate images and video, reason across multiple information sources, operate digital tools and increasingly execute multi-step tasks with limited human involvement.

The boundaries between language models, reasoning systems and autonomous agents are becoming less distinct. In the digital world, intelligence is becoming increasingly transferable.

Give an advanced AI system a legal document, a software repository, a financial model or a marketing problem and it can often adapt remarkably quickly.

But there is another test of intelligence that receives much less attention: **what happens when AI leaves the screen?**

The Physical World Changes the Test

At CES 2026, NVIDIA put physical AI at the center of its robotics strategy. The company described physical AI as models capable of understanding the real world, reasoning about it and planning actions. NVIDIA also declared that the **"ChatGPT moment for robotics is here."**

That distinction between digital AI and physical AI matters.

A language model can learn enormous amounts about gravity from text, images, simulations and video. But understanding the statistical representation of gravity is different from operating reliably in an environment where gravity, momentum, occlusion, human behavior and mechanical systems interact in real time.

A generated landing page can be regenerated. A line of code can be rolled back. An AI-generated report can be corrected. **The physical world often provides no equivalent reset button.**

Industrial Environments Make This Especially Clear

My own perspective on this question comes from more than a decade of building AI and IoT systems for industrial environments. At Trio Mobil, our systems operate around forklifts, pedestrians, loading docks, production areas and warehouses. These environments expose AI to a very different set of constraints.

Consider a forklift approaching an intersection while a pedestrian walks toward the same area. A useful AI system must do more than recognize that forklifts and pedestrians exist. It may need to determine:

1. Whether the pedestrian is actually on a collision path
2. Whether another object is partially obscuring the person
3. How quickly the forklift is approaching
4. Whether the operator has already reacted
5. Whether the pedestrian is moving toward or away from danger
6. Whether the interaction represents normal traffic or developing risk
7. Whether intervention is required within fractions of a second

The problem is not simply detection. It is perception, context, prediction, decision-making and, in some applications, physical intervention. And the consequences of being wrong are fundamentally different from those of making an error in a purely digital environment.

Intelligence Is Not the Same as Reliability

This is one reason I believe the AGI discussion needs to separate capability from reliability. An AI system might demonstrate extraordinary reasoning ability while still being unsuitable for safety-critical physical applications.

In industrial environments, a system cannot simply be impressive on average. It needs to operate under changing lighting conditions, unusual human behavior, different vehicle types, blind corners, environmental noise and edge cases that may occur only rarely. And when a serious risk develops, **latency matters.**

The distinction becomes even more important when AI is expected not only to detect an event, but to trigger an alert, communicate with another system or influence the behavior of a machine. That requires more than a powerful model. It requires an architecture designed around the physical environment.

Detection Is Not Intervention

This is also why discussions about physical AI should distinguish between three different capabilities:

Three capabilities, three different levels of maturity

A system can be strong at one stage and unproven at the next

DETECTION	→	DECISION	→	INTERVENTION
Can the system identify the person, vehicle, object or developing risk?		Can it determine whether the situation requires action?		Can it reliably trigger the appropriate response quickly enough to affect the physical outcome?

Detection ≠ Decision · Decision ≠ Intervention · Each stage must be proven on its own

These capabilities can exist at very different levels of maturity. An AI camera may identify a pedestrian near a forklift. That does not automatically mean the same system can slow the vehicle safely and reliably.

A model may recognize unsafe behavior in recorded video. That does not necessarily mean it can make a safety-critical decision within the response window required by a moving machine.

As AI expands into robotics, autonomous vehicles and industrial automation, these distinctions will become increasingly important.

Physical AI Is the Next Major Test

None of this diminishes the progress being made in AI. Quite the opposite.

The emergence of physical AI may be one of the most important developments in the field because it forces intelligence to confront reality rather than representations of reality.

AI systems will increasingly need to understand cause and effect, reason about space and time, adapt to unfamiliar environments and operate alongside humans whose behavior is not always predictable.

Simulation and world models will play a major role. So will edge computing, sensors, computer vision, real-time location technologies and increasingly capable robotics models.

But the ultimate benchmark will be what happens outside the simulation.

QUESTION 01
Transfer
Can the system transfer what it has learned into unfamiliar physical situations?

QUESTION 02
Self-correction
Can it recognize when its assumptions are wrong?

QUESTION 03
Real-time reliability
Can it act within real-world time constraints — and do so reliably enough that humans are willing to depend on it?

So, Has AGI Arrived?

That ultimately depends on what we decide the word *general* means.

If AGI means extraordinarily broad digital intelligence capable of performing an expanding range of cognitive tasks, the argument that we are approaching or have crossed an important threshold is understandable.

But I would set the bar differently.

> **General intelligence should transfer across domains, reason about cause and effect, adapt to unfamiliar conditions and act reliably in the world. If those abilities break down when intelligence leaves the screen, it is not yet general.**
>
> — Nevzat Atakli, CEO and Co-Founder, Trio Mobil

An AI-generated landing page can be retried.
An encounter with a forklift or crane cannot.