



# Artificial General Intelligence Definition Explained

## Why AGI Definitions Keep Changing - The Strategic Ambiguity Behind AI's Biggest Promise

*Artificial general intelligence is often discussed as if it's a fixed destination. It isn't. The term keeps moving because it does more than describe a technical goal; it helps companies persuade different audiences to do different things.*

In 1956, a group of scientists gathered at Dartmouth University to launch a new discipline. John McCarthy, then an assistant professor, chose the name “artificial intelligence” after rejecting “automata studies” the previous year. His colleagues worried about the implications. The name tied the field to recreating human intelligence, even though no one could actually define human intelligence with scientific precision.

Nearly seventy years later, that foundational problem hasn't been solved. Psychology, biology, and neurology still offer no unified account of what human intelligence is. Historical attempts to measure and rank it have often been entangled with efforts to justify hierarchy and exclusion. Yet artificial general intelligence, or AGI, remains the industry's grand promise.

Artificial general intelligence has no stable definition because it isn't primarily a scientific target. It's a strategic communication tool that can be reshaped for the audience in front of it.



## TL;DR

AI began with an impossible benchmark: recreate human intelligence without a settled definition of what human intelligence is. That gap didn't disappear as the field matured. It became useful. Companies can describe AGI one way to regulators, another way to consumers, and still another way to investors, all without ever leaving the shelter of an undefined goal. That's why AGI can sound like a cancer cure in one room, a digital assistant in another, and a revenue threshold in a boardroom.

## The Original Problem: No Goalposts for Human Intelligence

That takes us back to the core difficulty. When McCarthy named the field, he effectively attached it to a moving target. There are no clean goalposts for artificial intelligence because there are no clean goalposts for human intelligence itself.

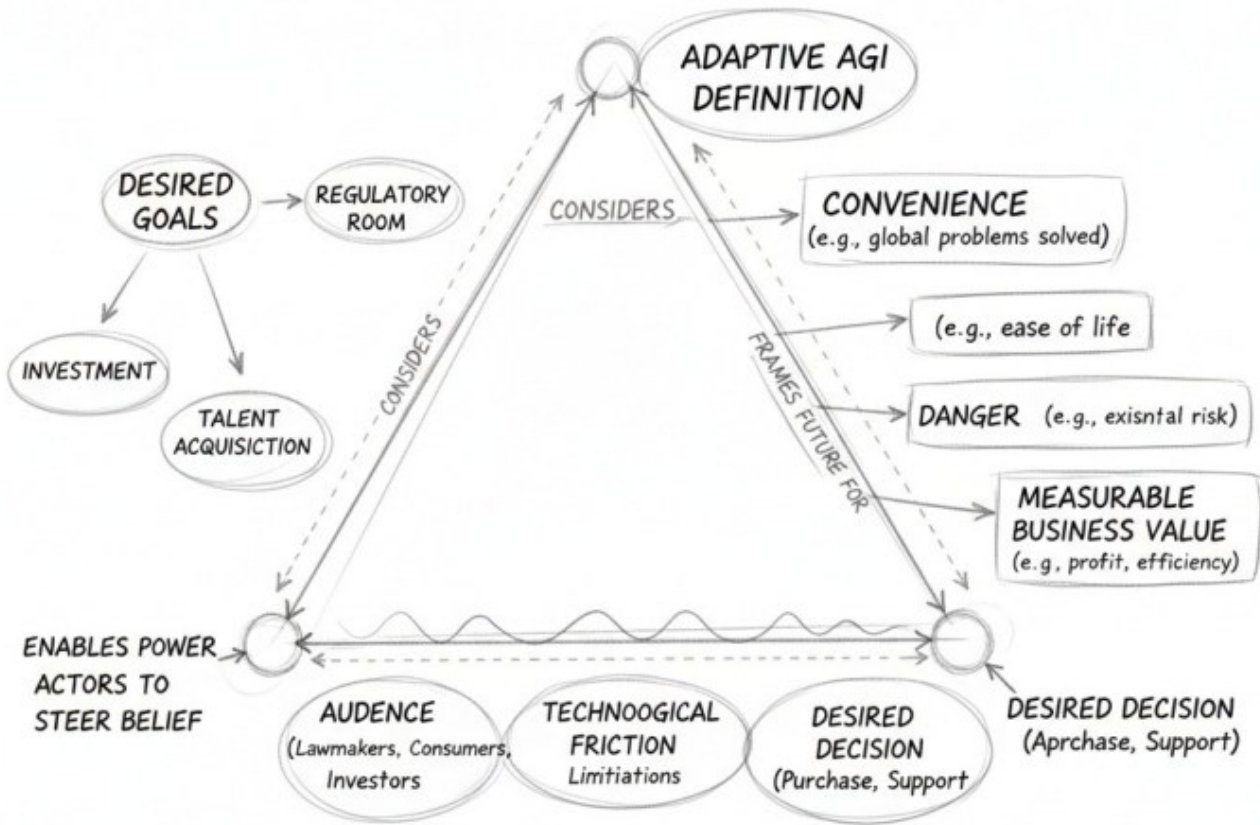
This isn't a minor technical oversight. It's the central constraint. Human intelligence includes pattern recognition, abstraction, memory, emotional judgment, improvisation, social reasoning, and creative problem-solving, among other capacities. Different cultures value different expressions of intelligence. A chess grandmaster, a jazz musician, and a forest tracker each show profound intelligence, but not in ways that reduce neatly to one scale.

Once the target is that unstable, claims about approaching it become unusually flexible. This is where the mechanism becomes clearer. Rather than treating AGI as a precise threshold, many organizations use it as a variable term that can absorb whatever meaning serves the moment. I call this the Triangulation Method: defining the same promised future differently depending on the audience, the friction, and the decision being sought.

The desire is straightforward enough. Companies want investment, regulatory room, talent, adoption, and public legitimacy. The friction is that today's systems still fall short of the sweeping promise implied by "general intelligence." So the belief has to be managed. The mechanism is adaptive definition: frame AGI as salvation, convenience, danger, or measurable business value depending on what the listener needs to hear. The decision condition is audience-specific. A lawmaker must feel caution without restriction, a consumer must feel usefulness, and an

investor must see a plausible path to return.

## THE TRIANGULATION METHOD: Strategigly Defining AGI



Same promised future → Different framing for different audiences.

## How AGI Became a Marketing Variable

Once you see that mechanism, modern AI messaging looks different. Artificial



general intelligence sounds like a technical milestone, but in practice it often behaves more like a strategic variable.

OpenAI offers a revealing example. Across public settings, the company has described AGI in at least four distinct ways. In front of Congress, AGI has been framed as a force that could cure cancer, address climate change, and reduce poverty, which positions it as too important to restrain too aggressively. For consumers, AGI becomes the ultimate digital assistant, a promise designed to make the abstract feel intimate and immediately useful. For Microsoft, AGI has been tied to a system capable of generating \$100 billion in revenue, translating an ambiguous technological future into a concrete investment logic. On its website, AGI is described as “highly autonomous systems that outperform humans in most economically valuable work, ” which sounds technical while remaining broad enough to stretch.

These aren't just different wordings of one stable concept. They're different promises. Each one is calibrated to mobilize a specific audience toward a specific decision.

When a definition changes with the room, the definition isn't just describing a technology. It's doing strategic work.

That doesn't mean the underlying technology is fake. It means the language around it can't be taken at face value as a neutral scientific description. The same term is being used to coordinate belief across multiple constituencies at once.

## **The Existential Risk Gambit**

This becomes even more visible in the way existential risk enters the conversation. In 2015, before OpenAI's official launch, Sam Altman wrote: “Development of superhuman machine intelligence is probably the greatest threat to the continued existence of humanity... AI is probably the most likely way to destroy everything.”

Placed in isolation, that reads like a straightforward statement of conviction. In strategic sequence, it looks more complicated. Altman was trying to convince Elon Musk to co-found OpenAI. Musk had already expressed concern about AI as an existential threat, so apocalyptic framing aligned with Musk's worldview and



increased the chances of securing his participation and resources.

Later, the same broader project would present advanced AI to Congress as a solution to humanity's greatest problems. That shift matters. It suggests that existential risk narratives can function not only as sincere warnings but also as positioning tools, shaped to recruit allies, attract attention, or establish moral seriousness.

A former AI researcher put the pattern plainly: “I spent months trying to reconcile different statements from the same company about what they were building. Eventually I realized they weren't contradicting themselves, they were optimizing for different outcomes with different audiences.”

That observation helps explain what otherwise feels incoherent. The contradiction isn't necessarily evidence of confusion. It may be evidence of strategy.

## **Recognizing Strategic Definitions in Practice**

Once you understand the Triangulation Method, AI claims become easier to parse. The question shifts from “Is this definition correct?” to “What is this definition trying to accomplish here?”

Start with audience. If a company describes the same system one way to regulators, another way to customers, and another way to investors, the variation is probably not accidental. Then look at timing. Claims about existential risk, timelines, or capabilities that swing sharply without a corresponding technical breakthrough often signal a change in strategic need more than a change in scientific understanding. Finally, watch for benchmarks built on elastic terms such as “human-level intelligence” or “general intelligence.” If the threshold is vague enough, progress toward it can always be redescribed after the fact.

None of this requires cynicism. A company can genuinely be building powerful systems and still use strategically flexible language to manage stakeholders. In fact, that's often the point. Technical ambition and narrative opportunism aren't opposites. They can operate together.



## What This Means for How You Read AI Claims

The practical shift is simple but important. Instead of asking only what an AI claim says, ask what it wants. Who is the intended audience? What action is this message trying to produce? What version of AGI appears here, and how does that compare with what the same organization says elsewhere?

That habit won't tell you everything, but it will help you separate technical description from audience design. It also keeps you from mistaking verbal scale for actual capability. In a field this charged, language often arrives before the evidence does.

The faint glimmer in the blackness is still there. Real advances in AI exist, and some are genuinely consequential. But if you want to see them clearly, you have to filter out the strategic fog that gathers around undefined promises. AGI keeps changing because the term is doing more than naming a destination. It's helping powerful actors steer belief while the destination itself remains unsettled.