MedCrave Step into the World of Research

Journal of Historical Archaeology & Anthropological Sciences

Short Communication

Open Access CrossMark

Why can't artificial general intelligence be achieved?

Abstract

Agnosticism argues that human beings cannot fully comprehend everything in the universe, thus making it impossible to create a general artificial intelligence whose intelligence surpasses that of humans. The achievements of artificial intelligence so far have been limited to successes in specific, simple system domains. A truly "general" intelligence should be capable of simultaneously solving many problems (e.g., 100,000 tasks). Furthermore, when it comes to complex systems, their complexity far exceeds that of games, making artificial intelligence even less capable of handling them. The computational power of AI depends on computer speed, which has a fundamental limit the speed of light. Real-world problems often involve multi-dimensional scenarios whose complexity increases exponentially without an upper bound. Thus, finite computational capability cannot satisfy infinite computational demands. Human intelligence, distributed across various functions through natural evolution, is inherently more efficient than artificial intelligence at the individual level. Therefore, AI cannot fully surpass human intelligence, but can only outperform humans in certain specialized tasks. A conceivable future scenario involves a division of labor between humans and artificial intelligence, where cooperation and complementary strengths create a mutually beneficial world.

Keywords: artificial intelligence, agnosticism, complexity, computing speed, intelligent distribution

Volume 10 Issue 2 - 2025

Hong Sheng

Unirule Institute of Economics, China

Correspondence: Hong Sheng, Unirule Institute of Economics, Beijing, China, Tel +8610-82932775

Received: July 24, 2025 | **Published:** August 20, 2025

Introduction

Agnosticism among great thinkers of the past

After artificial intelligence has made some notable breakthroughs, there has emerged an unlimited tendency toward technological optimism, believing that, at the current pace of AI development, it will soon surpass human intelligence and even replace humans as the new rulers of the Earth. One important sign of this is the realization of artificial general intelligence. However, regarding what constitutes intelligence, whether rationality can fully comprehend the universe, and the mechanisms behind intelligence and rationality itself, some great thinkers of the past have expressed less optimistic views. Kant stated that humans can never know the thing-in-itself because they must first understand where their sensibility and intellectuality come from and how they work. Concerning the ability to answer such questions, Kant was pessimistic. He said, "If we wish to make a judgment about the origins of sensibility and intellectuality, I can only see that such exploration completely exceeds the limits of human reason and is beyond our capability."[1]Hayek also believed that human understanding of the universe and especially of complex systems is "beyond reason." In the conclusion of his book The Sensory Order, he stated, "A complete explanation of even the external world as we know it would presuppose a complete explanation of the working our sense and our mind. If the latter is impossible, we shall also be unable to provide a full explanation of the phenomenal world."[2]

Some may argue that during the times of Kant and Hayek, science and technology had not developed to the extent they have today. They could not foresee the advanced development of modern artificial intelligence, so their judgments are outdated. My response is that both Kant and Hayek were extraordinary thinkers, and their judgments are not empirical but rather metaphysical and philosophical. The characteristic of metaphysical judgments is that they can provide insights in areas beyond human experience. Their judgments aim to establish the limits of human wisdom and warn humanity against misjudging its own intelligence. Some might say that artificial intelligence will surpass human beings, and that what humans cannot comprehend, AI can. My response is that since it is "artificial" intelligence, it must be designed by human rationality. If humans do not understand the reasons behind their own rationality, how can they design and create an intelligence that surpasses something they themselves do not recognize or understand?

Artificial intelligence has only succeeded in simple, isolated tasks

The significant breakthroughs in artificial intelligence today have all been achieved in isolated tasks. For example, AlphaGo defeated Lee Sedol, ChatGPT has become a widely accepted conversational agent, autonomous vehicles are officially on the road, and AI proves geometric problems to the level of an Olympiad champion, among others. The key point is that these are all isolated tasks. Machines have long surpassed humans in specific areas: cars are faster than humans, airplanes can fly while humans cannot, and computers calculate at speeds far exceeding human capability. So, do these many isolated successes combine to form an all-around champion? The crucial question is, how many tasks constitute "all-around" capability? Is it 100, 1000, or 10,000? According to human needs, even 10,000 tasks would not be enough. This is because human rationality must not only judge familiar things but also assess things they have never encountered, respond to various unforeseen events, and think about matters that sensory experience can never grasp. Any single task that artificial intelligence tackles is a problem that humans have deliberately set out to solve, with clear boundaries, defined rules, and simple variables. For example, this is true for Go. So, would 100,000 tasks be sufficient? Let's say it might be. However, when these 100,000 tasks are combined, they cannot achieve the desired outcome.

Let's first look at how breakthroughs in artificial intelligence have been achieved in isolated tasks. Taking AlphaGo as an example, on a 19×19 Go board, there are a total of 361 points. Players can choose any unoccupied point to place their stones. Each move has about 250 possible choices, and an average game requires at least 100 moves.[3]When considering the opponent's responses, this results in over 200 moves. If one were to find the best move right from the start, it would be ideal to try every possible option and identify the one with the best outcome. However, the number of possible choices is extraordinarily vast: 250 raised to the power of 100 yields approximately 6.223E+239 options; raising it to the power of 200 results in what current calculators call "infinity." The fastest computer in the world can only perform about 2E+16 calculations per second. AlphaGo significantly reduced the number of trial-and-error attempts per game by employing the Monte Carlo random method. Through millions of iterative training sessions, it ultimately identified moves close to optimal and assigned them the highest weight. This method is what we refer to as "exhaustive testing." I liken it to counting on one's fingers. The idea is that exhaustive testing, or counting on fingers, is the foundation of any form of thought; while thinking is about discovering the best choice hidden among many options using more clever and less trial-and-error than mere fingers' counting. When a computer can greatly enhance the speed of counting on fingers, it may surpass human cognitive abilities, finding answers more quickly or "intelligently." However, in form, this approach appears quite "clumsy."

If artificial intelligence were to "think" not in the pattern of human cognition but in another mode, it might have the potential to surpass humans through this high-speed exhaustive testing method. Its feasibility lies in its simplicity and technical ease of implementation. We can hypothesize that the causes of all things in the universe and the causal relationships between them are akin to a game of Go, where one selects the best choice from an extraordinarily vast array of possibilities. This "best choice" corresponds to the concept of "Dao" in traditional Chinese Confucianism and Daoism. "Dao is the essence of all things," meaning that Dao is the cause of everything. However, it is not easy to discover; it is often unclear, as expressed in the phrase, "You cannot see its beginning when you face it, nor can you see its end when you follow it," and it is described as "vague and indistinct." To seek and discover the Dao requires tremendous effort. Wolfram employs a method in his book "A New Kind of Science" that is quite similar to this approach. He first designed the simplest cellular automaton model, which is one-dimensional with three cells and two states, resulting in a total of 256 possible rules. Through exhaustive testing, he found that three of these rules were meaningful or effective, displaying orderly patterns. These three rules are akin to the "Dao" or the best moves in Go. This indicates that in extremely simple systems, one can use the exhaustive testing method.

Artificial intelligence cannot be the all-around champion in complex matters

However, when things become a bit more complex, such as when a cellular automaton transitions to one-dimensional with three cell and three-state systems, the number of possible rules increases to 7,625,597,484,987.[4] When moving to two dimensions, the number of possible rules expands even further. For instance, the number of types of limiting rules for a two-dimensional nine-cell two-state system is 4,294,967,296.[4] Humans typically face three-dimensional or four-dimensional spaces, where the number of possible choices is unimaginably large. Such a vast number of possibilities makes it far beyond human capability to exhaustively test all options, and even computers struggle with this task. As mentioned earlier, the fastest computer in the world can perform only about 2E+16 calculations per second, while the number of possible moves in Go amounts to 6.223E+239, even to be described as "infinity." Go is merely a greatly simplified game invented by humans; any real-world problem that needs to be addressed will have an even larger number of possible solutions. Therefore, relying on high-speed exhaustive testing methods to solve practical issues beyond games is likely beyond the current capabilities of computers. For example, although artificial intelligence has made significant advancements in image, sound, and text recognition, combining these AI dimensions does not just simply add their complexities together. Instead, the interaction and combination of these dimensions exponentially increase the number of possible choices by several orders of magnitude.

Moreover, to achieve artificial general intelligence means excelling in 100,000 different tasks. Each task and its corresponding solutions have unique characteristics, and most importantly, their value judgments differ. Therefore, a single model cannot be used universally; for instance, one cannot apply a value function for plants to assess the quality of animals. These tasks also exist in different dimensions; visual image recognition, sound recognition, and text recognition are vastly different, and simply bundling them together is not sufficient. Thus, "general" artificial intelligence implies having a similarly high level of understanding and responsiveness to different tasks and dimensions of information, but it cannot rely on a universal model to handle everything. Different tasks may need to be processed simultaneously, which is a common scenario humans encounter, such as facing rain and the attacks of wild animals while crossing a river. Artificial general intelligence must maintain sufficient computing power to address multiple events occurring at once. Integrating and recognizing information from different dimensions further increases the number of possible choices by several orders of magnitude, necessitating even faster exhaustive testing capabilities. However, when dealing with single tasks, the speed of computers is already constrained, and their inadequacy becomes even more pronounced when addressing multiple combined issues.

The ultimate limit of artificial intelligence computing speed is the speed of light

Some may argue that the computing speed of computers will continue to grow. According to Moore's Law, the performance of microprocessors doubles approximately every 18 months. However, there are already signs that Moore's Law is slowing down. Some have demonstrated that there are insurmountable limits to Moore's Law. On a microscopic level, transistors have become so small that they reach atomic dimensions, which is the limit of smallness; on a macroscopic level, the exponential increase dictated by Moore's Law will eventually exceed the total number of atoms in the universe, surpassing the limit of largeness. Certainly, there are phenomena that surpass Moore's Law today, known as the Nvidia phenomenon, or Huang's Law- Nvidia has increased computing speed by a thousand times over eight years. This has mainly been achieved by lowering precision without affecting results and reducing unnecessary neurons, among other measures. In fact, it involves clever methods to minimize redundant computations.[5] Additionally, the integration of numerous interconnected chips has allowed them to accumulate their computational power, forming greater computational capacity and faster speeds together. Moreover, there are quantum computers and biological computers. It seems that humanity always finds ways to overcome obstacles and continuously improve the trial-and-error speed of computers. If this trend continues, the speed of computers will only increase, and one day it will meet all the needs of artificial general intelligence.

However, this viewpoint is merely a low-order perspective, based on past experiences and failing to recognize the limits. It is akin to the experience of turkeys on a farm; they consistently receive food from their owner every day, but they do not realize that one day the owner will come not to feed them, but to slaughter them. From a godlike perspective, Moore's Law, or similar exponential growth laws, only approximately hold true under low-speed conditions. They may continue to grow, but as speeds approach their limits, they begin to fail. We know that the ultimate speed limit is the speed of light. Computers operate based on the movement of electrons, and the speed limit of electrons is also the speed of light. From a metaphysical standpoint, according to Einstein's mass-energy equivalence formula $E=mc^2$, energy equals mass multiplied by the speed of light squared. If speed were infinite, converting energy into matter would require infinite energy; space-time would become unified, allowing one to see point B from point A without the need for time, rendering past and future irrelevant, and causal relationships would be instantaneously realized. This is impossible; even if it were possible, infinite energy would be unattainable. Therefore, there must be an upper limit to speed, which is the speed of light. Physicists tell us that as speed approaches the speed of light, photons begin to encounter increasingly greater resistance. The closer one gets to the speed limit, the greater the resistance becomes, until it reaches infinity. Thus, the speed of computers must ultimately have a limit.

Now, we can estimate the upper limit of computer computing speed based on the speed of light. Suppose that for each calculation, an electron travels a distance of 1 centimeter. Dividing 300,000 kilometers by 1 centimeter gives us 30,000,000,000, or 3E+10 calculations per second. This is the speed for a single machine. Even if we employ Nvidia's methods or platform integration approaches, the increase can only be finite; even with a hundred million times increase, it would still only reach 3E+18 calculations per second. This still establishes an upper limit on the demand for computing speed. If we consider 1 centimeter too long, even if the distance for each calculation is reduced to 1 nanometer, the speed can only increase to 3E+25 calculations per second. This is still far from the previously mentioned number of possible moves in Go, which is 6.223E+239 choices, approaching "infinity." Moreover, the real problems faced by humanity are far more complex than those in Go. Perhaps humans will discover some clever and simple algorithms that save a significant number of calculations, but this requires time. Within any given timeframe, the fastest computing speed is always a finite number. Saving calculation times is equivalent to increasing speed by a finite factor, which always has an upper limit. To achieve artificial general intelligence, the demand for computation can approach infinity. On the other hand, according to Yann LeCun, the computational speed of a computer is merely one hundred-thousandth that of the human brain.[6] This is clearly just a rough estimate. If we assume that the computing speed of the brain has already reached the limit of the speed of light, then artificial intelligence will ultimately be constrained by the upper limit of computing speed, making it impossible to achieve artificial general intelligence.

Although artificial intelligence cannot infinitely increase computational speed, the demand for computational speed is limitless. To achieve artificial general intelligence, the demand for computation can approach infinity. There is a possibility that the tasks assigned to artificial intelligence may initially appear simple, with a manageable number of calculations. However, as complexity increases, the number of possible choices that need to be explored can grow exponentially. This is akin to the story of an Indian king who easily agreed to reward a grain of rice for the first square of a chessboard, doubling the amount for each subsequent square. By the time he reached the 64th square, the total was an astronomical figure. The problems assigned to artificial intelligence might seem uncomplicated, but when it is required to simultaneously recognize images, sounds, touch, smell, taste, and text- such as the feeling of spring, which is quite ordinary for humans- the potential choices become nearly infinite. Regardless of how fast artificial intelligence may become in the future, it will ultimately remain limited, and it is impossible to meet infinite demands with finite capabilities.

A hypothesis: humanity has reached the limit of intelligence in the universe

If the above judgment holds true, it leads to an important hypothesis: if artificial general intelligence cannot be realized due to its speed limit at the speed of light, meaning it cannot surpass human intelligence, then humanity has already reached or is close to this limit of intelligence in the universe. This represents the quantifiable boundary of human bounded rationality. Although artificial intelligence may surpass humans in specific tasks, both humans and AI will have limits on computing speed. When the total intelligence is distributed across various functions, the computing speed for each function will be restricted. Humans can only employ "clever" methods to save on the number of exhaustive testing. From a comprehensive perspective, the way humans distribute their intellectual resources is superior to that of artificial intelligence. When we view the intelligence models of AI and humans as two interchangeable or competing paradigms, the human model distributes intellectual resources across various functions, so that each individual function may not seem very fast. In contrast, AI concentrates its intellectual resources on specific areas, enabling it to employ high-speed exhaustive testing methods and achieve success in those particular tasks. When overall intelligence is limited by the upper limit of the speed of light, it becomes scarce and must be allocated reasonably and effectively across different areas. When one area excels too much, other areas will correspondingly fall short. Different individuals have different strengths: some are intelligent, some are graceful, some are agile, and some are profound. Additionally, some people may develop extraordinary abilities in a specific area due to their interests or professions, but at the cost of lower capabilities in other areas. This means that their superior endowment in one aspect is balanced by their lack of ability in another.

This is an adjustment in the distribution structure of intelligence caused by natural evolution. In other words, the reason why humans have evolved to have the current distribution of intelligence is a result of natural selection. They allocate more intellectual resources to functions that are most critical for survival and fewer resources to those that are less important. The fundamental principle regarding how much to allocate and what kind of structure to form is: sufficient, rather than the fastest or the best. The ultimate test is survival. Assuming that the computing speed of artificial intelligence reaches its limit, matching that of humans, the distribution of its capabilities across various functions would still be controlled by humans. It could allocate all its computing power to playing Go, but in other areas, it would be zero. This merely pushes the specialization of humanlike intelligence distribution to the extreme. When it attempts to distribute this computing power to other functions, it lacks an understanding of what kind of structure should be formed. Since artificial intelligence does not directly face natural selection, it cannot achieve a reasonable distribution of intelligence through natural selection. It relies on the specific needs of different humans for its development, and thus cannot organize intelligence as rationally as humans do; overall, it cannot use limited intelligence as effectively as humans can. This means that, not only at the limit of speed but also in structure, it cannot reach the level of intelligence that humans typically achieve.

Conclusion

Division of labor between artificial intelligence and humans

It can be said that, in an evolutionary sense, humans are the "nature" of artificial intelligence. The distribution of intelligent resources that artificial intelligence aims to achieve depends on human choices, which are merely responses to human needs. We can speculate that human intelligence and artificial intelligence are not identical in their patterns; they will have different strengths and weaknesses in various aspects. This creates space for cooperation between them. Generally, people will prioritize functions that are costlier for them and cheaper to implement with artificial intelligence, making it an economic choice based on cost-benefit calculations, rather than relying on functions that humans consider natural or instinctive, which have evolved over time (e.g., quickly grasping the main points of a scene and understanding the meanings and emotions behind it). Achieving these functions with artificial intelligence can be prohibitively expensive and not very economical. In other words, in certain areas, artificial intelligence may be more efficient, while in others, humans may excel. Therefore, from the perspective of human choices, they are unlikely to develop a "general" artificial intelligence; rather, they will create a "complementary" artificial intelligence. From the standpoint of natural selection, such artificial intelligence would not survive without human assistance.

In this discussion, we consider human intelligence and artificial intelligence as two distinct modes of intelligence. One is a specialized, exceptionally fast exhaustive testing approach, while the other is a slower exhaustive testing process that can significantly reduce the number of trials through clever methods. These "clever methods" include importance selection, classification, abstraction, simplification, intuition or insight, metaphysical generalization, and the creation of concepts such as the divine, among others. To some extent, the two can substitute for each other; however, due to their different energy efficiencies, such substitutions cannot occur arbitrarily. If we evaluate human individuals as units of intelligence in terms of energy efficiency, their relative efficiency is quite high. This is because individual humans are relatively limited, yet, as previously mentioned, the wisdom functions they produce are substantial. In contrast, current artificial intelligence often operates on a large scale with multiple-machine platforms, resulting in enormous energy consumption. The energy expended by individuals is limited to what they can personally obtain, whereas the energy consumed by artificial intelligence is supplied on a large, centralized scale. "The human brain requires only about 20 watts to operate, achieving extremely high information processing efficiency".[7] However, contemporary AI models are becoming increasingly complex, leading to greater overall energy consumption. Although the energy consumption per unit is decreasing, the energy scale for artificial intelligence, as a unit of intelligence, is far greater than that of individual humans. "The daily electricity consumption of ChatGPT is more than 17,000 times that of an average American household".[8] If this mode of intelligence in artificial intelligence is superior to that of humans, why has natural selection not evolved larger energy-consuming human individuals?

To go further, we cannot compare artificial intelligence merely as a large-scale entity with individual humans; instead, we should compare it with humanity as a whole. This is because human "computation" or cognition operates at the level of the entire human race. Today, with nearly 8 billion global inhabitants, every moment they are "testing exhaustively" in a random manner. In this age of rapid communication technology, whenever a superior choice yields good results it quickly spreads and is emulated by others. Moreover, there is a division of labor among people. Today, this division of labor is essentially an intellectual division, which saves costs in information and knowledge, significantly enhancing the energy efficiency of human wisdom and raising the overall level of intelligence far beyond that of individual humans. In fact, the training materials used by large models are the varied responses of countless individuals, generated through their interactions, and this scale of individual wisdom may be the optimal scale from a natural perspective. A centralized artificial intelligence can only be based on this form of wisdom derived from many individual humans. Without a scale or distribution similar to that of human individuals, a large-scale centralized artificial intelligence lacks sufficient input material and cannot be intelligent.

From the perspective of energy efficiency of wisdom, the relatively low-efficiency model of artificial intelligence should not and cannot significantly replace human intelligence. Yann LeCun stated, "The power consumption of the human brain is approximately 25 watts, while a GPU consumes 250 watts ten times that of the human brain." He also mentioned that "the computational speed of a computer is merely one hundred-thousandth of that of the human brain," and "the efficiency of electronic systems is one millionth of that of biological organisms".[6] If it were to replace human intelligence, it would be uneconomical from an economic standpoint and not the "fittest" from the perspective of natural selection. Some may argue that the energy consumption of artificial intelligence is external to humans and merely represents an additional enhancement of human intelligence. However, the energy consumed by artificial intelligence, both directly and indirectly, is produced by humans, and there is a relationship of "as one falls, another rises" between the two. According to the current pace of AI development, the Uptime Institute predicts that by 2025, the share of AI-related electricity consumption in global data centers will surge from 2% to 10%.[7]

Furthermore, we can foresee that as computing speeds approach the speed of light, the energy consumption per unit will sharply increase. This means that the proportion of energy consumed by artificial intelligence in the total energy produced by humans is rising rapidly, which will also reduce the proportion of other energy consumption (including direct human consumption). Thus, not only is the realization of artificial general intelligence impossible, and it cannot completely replace humans, but its supplementation of human intelligence is also limited. When measured by the cost or energy efficiency of wisdom, both humans and artificial intelligence have comparative advantages over each other, which provides space for their division of labor. Therefore, a society that fosters cooperation between the two is superior to one where either of them develops independently.

Finally, if human intelligence reaches the limits of the universe, then the failure to develop a form of intelligence similar to artificial intelligence cannot be seen as a neglect of natural evolution. If the intelligence model of artificial intelligence is more effective, humans might as well evolve into such a form. Conversely, we should realize that adopting this intelligence model would not be able to achieve the goals of human survival, so from the perspective of natural evolution, the human intelligence model is the most effective, while the intelligence model of artificial intelligence is inferior to that of humans. Even if we assume that human intelligence has reached the pinnacle of the universe, it does not mean that human intelligence will no longer progress. The path of advancement is not simply about increasing brain capacity or improving computational speed; rather, it is about progressing through the interaction of many brains and the accumulation of past knowledge. The significance of artificial intelligence lies in its ability to operate on energy external to humans, providing intelligence in a manner different from the human model. However, due to its inability to discover the rational allocation of its wisdom resources and its relatively low energy efficiency, it can only serve as a supplement to human intelligence, enhancing wisdom efficiency through the division of labor with humans. Therefore, the realization of artificial general intelligence is impossible.

Acknowledgements

None.

Conflicts of interest

The authors declare that there are no conflicts of interest.

References

1. Zeng J, Liu Y. Kant’s Wisdom [e-book]. Beijing, China: China Film Publishing House; 2007.
2. Hayek F. The Sense of Order. Chicago, IL: The University of Chicago Press; 1976.
3. Michell M. AI 3.0. Chengdu, China: Sichuan Science and Technology Press; 2021.
4. Wolfram S. A New Kind of Science. Champaign, IL: Wolfram Media Inc.; 2002.
5. Semiconductor Industry Observation. Four Elements of Nvidia’s Success. Sina.com. September 8, 2023.
6. LeCun Y. The Road to Science. Beijing, China: CITIC Publishing Group Co., Ltd.; 2021.
7. Ava Y. How to bridge the gap between artificial intelligence and human brain intelligence. Follow Up. September 22, 2024.
8. Wu Z. Amazing! Large models become ‘electric tigers’ with astonishing energy consumption. How should tech giants respond? Open the World with Martial Arts. March 12, 2024.