OCCASIONAL PAPER 200 | MARCH 2024

MIND OVER MATTER

The philosophical arguments around AI, natural intelligence and memory

John Sweller

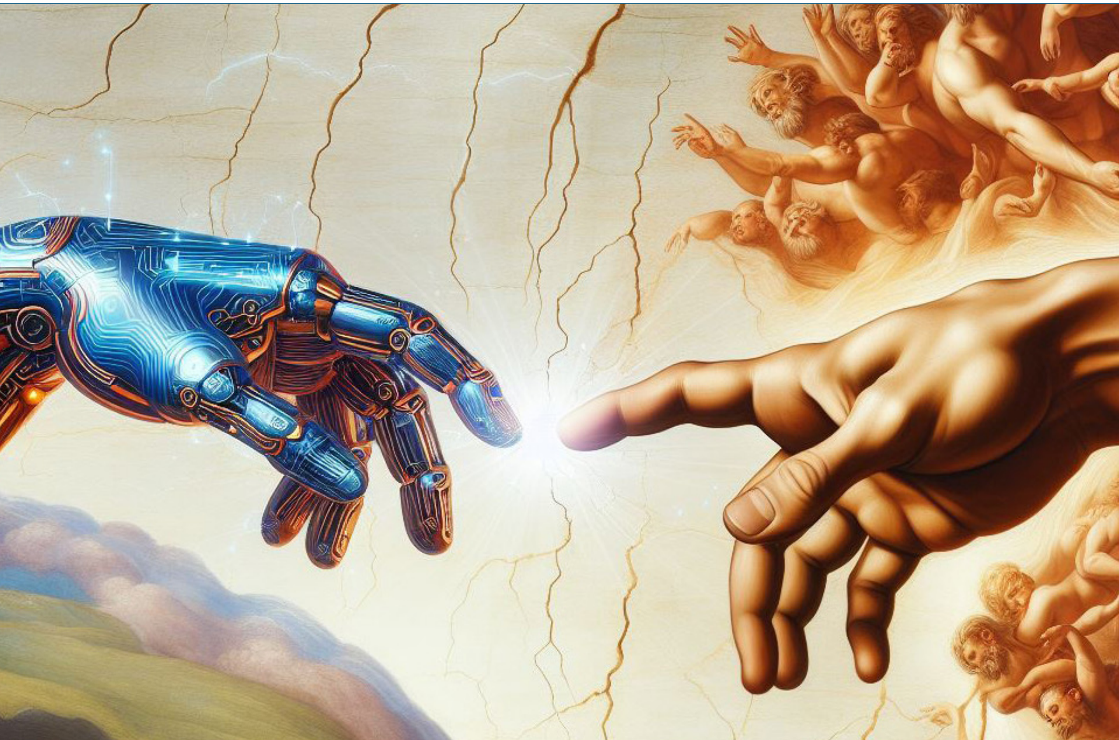

THE CENTRE FOR INDEPENDENT STUDIES

NATIONAL LIBRARY OF AUSTRALIA

A catalogue record for this work is available from the National Library of Australia

MIND OVER MATTER

The philosophical arguments around AI, natural intelligence and memory

John Sweller

Emeritus Professor of Educational Psychology
at the University of New South Wales and
Centre for Independent Studies Adjunct Research Fellow

CIS Occasional Paper 200

THE CENTRE FOR INDEPENDENT STUDIES

2024

Published March 2024
by the Centre for Independent Studies Limited
Level 1, 131 Macquarie Street, Sydney NSW 2000
Email: cis@cis.org.au
Website: cis.org.au

Foreword

Though artificial intelligence is quickly becoming a greater part of our everyday lives, the latest entry in the CIS research series provides a timely opportunity to revisit what makes any system — be it an artificial or a natural one — intelligent in the first place.

In characteristic engaging style, Professor John Sweller employs metaphor to provide a novel understanding of the foundations of intelligence. Namely, he draws on our understanding of other intelligent systems — including from evolutionary biology — for fresh insights.

As his life's work attests to, human cognition is itself an intelligent system that relies heavily on our existing knowledge. Sweller's essential work for educational contexts comes from his emphasis in underlining the critical role that knowledge plays in how humans — especially children in classrooms — navigate novel information; particularly when learning new information.

Just as the role of knowledge in human cognition has historically been underestimated by researchers — especially in education circles — it has also been downplayed in the development of artificial intelligence.

This observation helps make sense of why it is that we've seen the explosion in artificial intelligence so recently; despite the field of research now extending for more than half a century.

He again uses metaphor to make the case. When humans don't have sufficient knowledge on which to draw, we have limited useable intelligence. Similarly, without considerable information (knowledge) to access — as we now do thanks to massive storage of data — artificial intelligence applications are limited in utility.

Sweller's wider research has implications for the way artificial intelligence works — and its limitations. It's now largely accepted in education that it can be very inefficient for humans to 'discovery learn' (in effect, attempt to solve problems by practicing problem-solving rather than systematically building foundational knowledge) our way to solutions. However, the same is true for artificial intelligence, and for similar reasons. This is self-evident when considering the potential for AI 'hallucinations' — where false or implausible conclusions are generated when there is insufficient, or low quality, information available to an algorithm.

Time will tell if artificial intelligence will ultimately deliver on the ambitious promises for education systems and the wider economy and society. But, as Sweller warns us in this paper, much can be understood from reflecting on the foundations of intelligent systems as we navigate the opportunities, challenges, and limitations of the AI of today and the future.

Glenn Fahey is program director in education policy at the Centre for Independent Studies.

INTELLIGENCE is an unusual concept. We all feel we understand what is meant by the word but no one, not even psychologists who invented intelligence tests, has provided an adequate definition or even a description other than that an intelligent person obtains a high score on an intelligence test. We assumed that whatever it meant, intelligence was a characteristic solely of humans, although we conceded that some animals sometimes acted intelligently but only if their actions slightly resembled complex human actions.

The comforting feeling that however intelligence should be defined, only humans are highly intelligent, was disrupted recently by the development of artificial intelligence programs. The main but not the only skill of these programs is to write meaningful, natural language text on complex subject matter when requested to do so. They not only can write such text, but they can do so in seconds, a feat not matched by humans.

Apart from humans and artificial intelligence, I argue there is another intelligent system with abilities that exceed those of either human or artificial intelligence. That system is evolution by natural selection that created all species on earth. Of course, biological evolution is not artificial, it is natural. Not only is it natural, it functions in a manner very similar to human cognition (Sweller, 2022; Sweller & Sweller, 2006) which should not be surprising since we are ourselves a product of evolution by natural selection.

What makes evolution by natural selection and human cognition intelligent systems?

While intelligence may be difficult to define, on any useful definition, evolution by natural selection is an intelligent system. Simply based on its results, biological evolution should be classed as an intelligent system. It has managed to create processes, functions, and objects that to this point far exceed anything that humans can accomplish or even understand.

As one of a plethora of possible examples, consider the current human desire to find industrial procedures to convert plant matter into meat. While currently we can do so in a somewhat clumsy, expensive manner, this is a problem for which evolution by natural selection found a solution a very long time ago.

As one of many examples, evolution oversaw the construction of a self-replicating tiny factory several cubic centimetres in size that not only turns vegetable matter into meat but also, into a very large variety of other substances. We call that factory a mouse. We can understand a small proportion of the processes involved but the vast majority are beyond our ability to understand or replicate.

Perhaps in the future, artificial intelligence will be able to assist us in understanding the natural intelligence that produced this marvel along with the uncountable similar marvels of the biological world. Nevertheless, despite the lack of any assistance to this point from artificial intelligence, we have made considerable progress in understanding the processes that constitute the relevant natural intelligence underlying evolution by natural selection. Furthermore, those processes underly our own intelligence.

Random Generate and Test as a Generator of Moves under Uncertainty

Considered as a system, evolution by natural selection consists of a large store of genetically encoded information that is altered by random mutation. Every mutation is tested for effectiveness with effective mutations that assist the continuation of life being retained and ineffective mutations that interfere with the continuation of life being discarded. This process immediately leads to an obvious argument: *surely a system that relies on random generation of novel outcomes cannot be intelligent?*

In fact, random generation is the well-spring of all natural intelligence and creativity. Without it, humans cannot be creative, and for that matter, neither can artificial intelligence.

Analogously to random mutation, humans generate novel information by randomly generating moves during problem solving. In both cases the randomly generated information is retained if it is effective and discarded if it is ineffective. That mechanism is not only essential but also unique. There is no other way of generating novelty.

The essential role of random generation in creativity requires analysis. When faced with a novel problem, such as surviving a changing environment in the case of evolution or finding a new industrial procedure in the case of humans, previously used, known procedures will be tried first. Can previously constructed genetic information be used to meet the new environmental conditions, or can an old industrial procedure be used for a new purpose?

As indicated below, re-purposing old procedures is the simplest, most common way of handling new situations. Frequently of course, there are no old procedures available for use under the new circumstances. Under those conditions, there is no alternative but to randomly generate a new procedure and see if it works. If it does work it can be retained, if it does not work it must be discarded.

But does it not require intelligence to pick a more likely candidate move instead of a less likely one? It certainly does, but that requires knowledge of what constitutes a likely candidate and so the choice of a move is no longer purely random. If we reach a point where knowledge does not allow us to distinguish between the utility of different moves, we have no choice but to randomly choose one move and test it for effectiveness.

If it is effective, we retain it and if it is ineffective, we discard it, which is exactly what evolution does with respect to random mutation, assuming genetic structures are the evolutionary equivalent of knowledge. In the absence of knowledge in the case of humans or already available genetic structures in the case of evolution, there is no other way of generating problem solving moves. Random generate and test, called adaptive mutation in evolutionary biology, is the only procedure available.

Knowledge in the case of the human cognitive system, or current genetic structures in the case of evolution by natural selection, can circumvent this random generate and test process. Accordingly, the use of previous knowledge or the use of current genetic structures to determine moves when faced with novel conditions is the second way after random generate and test to determine moves. Furthermore, the use of knowledge or current genetic structures to determine moves provides the only conceivable determinant of intelligence.

Knowledge as a Generator of Moves under Uncertainty

Consider the spines of a porcupine. They bear minimal resemblance to fur but evolved from fur. If fur had not evolved, the spines of a porcupine with their current characteristics could not have evolved. The genetic characteristics of fur evolved primarily (but not exclusively) to allow mammals to retain heat but porcupine spines have no insulation function.

The genetic structures that allowed the evolution of fur were restructured to allow the subsequent evolution of porcupine spines. The evolution of porcupine spines, used for defence, could not have occurred without the previous evolution of fur, used for insulation. Genetic 'knowledge' of fur for insulation allowed the evolution of porcupine spines for an entirely different purpose.

There are countless instances of such evolutionary repurposing based on previous 'knowledge'. As an additional example, without the use of feathers for insulation, birds would not have feathers for flying. Without the stored genetic information, evolution by natural selection could not occur. Stored information that is adaptable and transmissible for indefinite periods of time is the essential core of an intelligent system.

If so, evolution by natural selection is an intelligent system. Random generation and test as exemplified by random mutation is an essential ingredient but the characteristics of genetically based 'knowledge' on which the random generation and test process functions determines the levels of intelligence.

Knowledge has exactly the same function in human cognition. Appropriate knowledge is an essential prerequisite for effective problem solving. If you were sufficiently fortunate to live in a society that had knowledge of the wheel either because your society had invented the wheel or obtained the requisite knowledge of wheels from another society, current early versions of modernity were open to you. Everything from transport to early clocks, to windmills could be further invented and become available to your society. Without access to the knowledge associated with wheels, you were barred from modernity.

On a more prosaic, everyday level, if your refrigerator breaks down you have a problem that requires an intelligent solution. Whom would you hire to provide that intelligent solution? Someone who is certified as having a high IQ or a certified refrigerator technician? Most societies and indeed, most individuals will place a higher reliance on the knowledge of a refrigerator technician. The refrigerator technician may well have a high IQ but is not hired because of that IQ. Knowledge of refrigerators is far more important. That technician may need to engage in random generate and test to find a solution, but knowledge and only knowledge can reduce the number of randomly generated moves to manageable levels.

The critical importance of knowledge held in long-term memory tends to be grossly underestimated. Consider the following sentences: *The person went into a restaurant. The waiter spilt soup on the person's lap. The waiter did not get a tip.* No one reading those three sentences will spend time puzzling over their meaning. The sentences and the scenario are very simple. Nevertheless, consider what needs to have been stored in long-term memory to enable us to understand the sentences.

Along with considerably more, we need to know: *a restaurant sells prepared food; waiters serve that food; soup is a liquid form of food; liquids can be poured; they are spilt if they are poured accidentally; a lap is a peculiar anatomical structure that only appears when we sit down but disappears when we stand up or lie down; waiters are given additional money if they serve their customers adequately; the function of money in a society.* If we do not hold every one of these concepts in long-term memory, the three sentences will be unintelligible.

The importance of knowledge held in long-term memory to intelligent behaviour was first recognised by De Groot (1965) who was concerned with why chess masters always defeated weekend players. The only difference between the two categories of players that he could find was that when shown a board configuration from a real game for five seconds, chess masters could reproduce the configuration with high accuracy while weekend players could remember where very few of the pieces had been placed.

That difference disappeared when the same experiment was carried out using random board configurations (Chase & Simon, 1973). Chess masters have memorised tens of thousands of board configurations along with the best move for each configuration allowing them to easily defeat weekend players who, in the absence of relevant knowledge, have no choice but to make a much heavier use of random generate and test.

Knowledge is a key and knowledge is why human cognition provides an example of an intelligent system. As is the case for evolution by natural selection, random generation and test is an essential procedure used by all intelligent systems but the knowledge base on which random generation and test operates determines useable intelligence. If you do not know how wheels, how a refrigerator, or a restaurant works, dealing with these entities will be challenging.

Artificial Intelligence

While work on artificial intelligence started more than half a century ago, most of those who worked in the area assumed, along with the societies in which they worked, that memorised information or knowledge was an unimportant factor in intelligence. Memorised information tended to be painted as 'rote learning' and unrelated to true intelligent behaviour. Accordingly, artificial intelligent systems designed to model human cognition placed minimal emphasis on a knowledge base.

The models worked after a fashion in that they could solve some types of well-defined problems such as physics or mathematics problems if they were provided with sufficient information, but they bore little resemblance to human cognition. Something was clearly missing, and it was only when it was realised that the missing ingredient was a large knowledge base and when procedures to access the knowledge bases in computers were developed, that the very recent explosion in artificial intelligence could occur.

Machine learning provided the procedure for accessing knowledge. Once computer programs could appropriately access and learn from very large data bases, those data bases provided the 'knowledge' that is essential for intelligent behaviour. Large computer data bases are the equivalent of the genetic codes required for evolution by natural selection to function or long-term memory with its enormous capacity required for human cognition. Intelligent behaviour requires a large data base. How intelligent that behaviour is depends on the sophistication of the relevant, accessible data bases.

Random generate and test is used by humans, by evolution and by artificial intelligence to determine action when required knowledge is not available. It is the only form of "problem-solving" available to any system in the absence of knowledge and so is a consistent, immutable strategy used by all intelligent systems. The other factor, stored information, is the real source of intelligent behaviour.

Consequences

Are we in danger from artificial intelligence? Unless we create intelligent machines that can reproduce themselves, the danger seems unlikely, although it remains to be seen how far artificial intelligence progresses. Without the ability of machines to reproduce independently of us, we can simply switch them off. Current computers completely rely on us to interact with the outside world. All information other than computer-based information is hidden from them. They do not have our sensory-motor systems that connect us to the outside world. That may change if we can construct computers that are physically independent of us but there is no sign of such machines in the foreseeable future.

On the other hand, evolution by natural selection, which is fully embedded in the outside world, preceded us, created us, does not rely on us, and is constantly creating new species and destroying old species. With more than 99% of the species that ever existed having become extinct (Jablonsky, 2004) we have no reason to suppose it is a benign form of intelligence. It remains to be seen if we, the offspring of evolution by natural selection, can ever fully control it.

References

Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, *4*, 55-81.
De Groot, A. (1965). *Thought and choice in chess.* Mouton. (Original work published 1946).
Jablonsky, D. (2004). The fossil record, together with modern data, can provide a deeper understanding of biological extinction and its consequences. *Nature, 427,* 589.
Sweller, J. (2022). The role of evolutionary psychology in our understanding of human cognition: Consequences for cognitive load theory and instructional procedures. *Educational Psychology Review*, *34*, 2229-2241. https://doi.org/https://doi.org/10.1007/s10648-021-09647-0
Sweller, J., & Sweller, S. (2006). Natural information processing systems. *Evolutionary Psychology*, *4*, 434-458.

MINDER OVER MATTER

The philosphical arguments around AI, natural intelligence and memory

The comforting feeling that however intelligence should be defined, only humans are highly intelligent, was disrupted recently by the development of artificial intelligence programs. The main but not the only skill of these programs is to write meaningful, natural language text on complex subject matter when requested to do so. They not only can write such text, but they can do so in seconds, a feat not matched by humans. Apart from humans and artificial intelligence, Professor Sweller argues there is another intelligent system with abilities that exceed those of either human or artificial intelligence. That system is evolution by natural selection that created all species on earth.

John Sweller is an Emeritus Professor of Educational Psychology at the University of New South Wales. His research is associated with cognitive load theory, an instructional design theory based on our knowledge of how humans learn, think and solve problems. The theory is a contributor to both research and debate on issues associated with human cognition, its links to evolution by natural selection, and the instructional consequences that follow. He is also a Centre for Independent Studies Adjunct Research Fellow.

CIS Occasional Paper 200
ISBN 978-1-922674-68-5
ISSN 0155 7386

cis.org.au

THE CENTRE FOR INDEPENDENT STUDIES