AI and the Mysterious

Jeff Dunne

The International Consciousness Research Laboratories
Email: jad@icrl.org

Abstract:

For many reasons ranging from misnomers to wishful ignorance, the nature of artificial intelligence (AI) is not well understood by many. This is particularly true of Generative AI, a tool that has demonstrated world-changing capabilities. As with any potent tool, the nature of its impact—for better and worse—will depend strongly on how it is used, which in turn depends strongly on the degree to which the user understands the tool's nature. The paper provides a brief overview of the functioning of Large Language Models (LLMs) and Completion systems like *ChatGPT*. It explores similarities and differences between them and tools of divination used throughout human history; this is of particular interest in the context of the large body of scientific evidence demonstrating that human intention can influence systems that incorporate a degree of randomness. We conclude by discussing whether mathematical algorithms like AI could have consciousness and what such consciousness might be like in relation to human consciousness.

Keywords: AI; consciousness; anomalous phenomena; cognition; divination

'If it walks like a duck, swims like a duck, and quacks like a duck, I call that bird a duck.'
-- James Whitcomb Riley, Poet (1849-1916), prior to visiting the Walt Disney Imagineering studio in Glendale, CA

1. Introduction

With the emergence of generative artificial intelligence (GAI) capabilities such as *ChatGPT*, many people have been hypothesising that these algorithms are headed towards (or maybe already achieved) a form of consciousness similar to that of a human. Yet the argument that something with human-like behaviours must have human-like cognition is akin to concluding that a submarine is functionally equivalent to a whale because they both can remain submerged for an extended period of time.

There is little doubt that GAI is already transforming significant aspects of society, but the important question is how humanity should—and should not—utilise this innovation. Can we trust it to correct our spelling? To prepare our meals? Perhaps, although even in such simple cases it is valuable to consider the implications of letting such skills atrophy. But what about using GAI to guide our thinking? How much can we safely let atrophy before we have

  ISSN: 2057-2301

passed a point of no return? To make such an assessment, we must start by understanding what GAI is and is not, as well as the nature of human awareness/consciousness itself.

There is a critical difference between the cognitive process of a human and the computational process of GAI. Both transform input to output, but in very different ways. Human cognition is grounded in meaning; we do not simply transform the letter sequence M-I-L-K into another letter sequence of B-E-V-E-R-A-G-E. We have *experiences* of milk, its drinkability, its textures… the idea of milk is associated with experiential *qualia*. When someone asks us, 'Please motivate the milk', we aren't confused because the words don't have a history of going together, but because we have an understanding, a world model, in which 'milk' is not a sensible direct object for the verb 'motivate'. If forced to make sense of such a sentence, we attempt to adjust the *meanings* we associate with the linguistic *representations* that stand for those meanings. Could 'milk' be intended as a verb? Is this person asking me to provide motivation for someone milking a cow? Or is 'motivate' being used loosely? Are they really asking for someone to help sell milk? To put the milk into translational motion? And of course, we recognise that sometimes people are just silly. Perhaps the request is to fancifully anthropomorphise milk, pretending it possesses a human-like measure of free will.

GAI does not work this way. Instead of computing on meaning, AI algorithms (and, in fact, all algorithms) manipulate the *symbols* of meaning. To ChatGPT, 'milk' is nothing more than a symbolic *token*, one piece of the immense, complex puzzle that is language-based communication. The brilliance of modern GAI is to have developed a tool that, having been informed by massive numbers of documents from across the internet, can sequence those tokens in a way that emulates reasonable use of language. While the exact details are complex and unnecessary to make the main points of this article, walking through a simplified example will illustrate the role that random numbers play in the process, which will resurface as an important point in later sections.

2. The GAI process

At the heart of GAI is the ability to calculate the probabilities for different tokens to follow a given input token sequence, referred to as 'the prompt'. Note that 'tokens' are not synonymous with 'words' despite the above example; 'milk' is one token, and 'Milk' is a different one. 'Friday' is not a single token, but the combination of two tokens, 'Fri' and 'day', which is why *ChatGPT* can return fake day names like 'Somniday' (for a dreamy, relaxing vibe) and 'Thrivensday' (a day to thrive and accomplish goals) when asked to generate names for

  ISSN: 2057-2301

the eighth and ninth days of the week. Humans also do this tokenisation effectively. Fluent English speakers confronted with 'to-geth-...' can predict with high confidence that the next token will be 'er'. They have seen the word 'together' many times, yet rarely (if ever) a sequence like 'togethup', 'togethimp', or just 'togeth' on its own. (It is worth noting for the sake of accuracy that while the examples used both above and below in this paper break words into 'human sensible' tokens to simplify explanations, tokenisation in the actual functioning of LLMs has no semantic basis at all; it is purely based on statistical analysis of the optimal way to break up the training data into the most numerically useful 'chunks' of letters, which may have no relationship to syllables or even word divisions.)

The basic response process, referred to as 'completion' (as will be obvious in a moment), works like this. When the LLM is given a prompt, such as 'How do I eat a cookie?', it creates a statistical ranking for all possible tokens that could follow the prompt. Some tokens, such as 'The' or 'Some' have comparatively high likelihoods. Others, such as 'Truck' or 'Flux', have extremely low likelihoods. And many, such as 'the' (note the capitalisation), 'Home', 'If', etc. have a likelihood somewhere in between. A random number is then used to select one of those tokens. Let us say the system chooses the 'Some' token. The system is then fed a new prompt—'How do I eat a cookie? Some'—and the process repeats. Driven by the random numbers operating on statistical assemblies of tokens, that could result in subsequent prompts a few steps down the road such as 'How do I eat a cookie? Sometimes a person…' or 'How do I eat a cookie? Some cookies may be…'. The probability of selecting the 'stop now' token option becomes higher as the process continues, and is eventually selected. The original prompt is stripped away, and the chat experience looks like a person asking 'How do I eat a cookie?' and the LLM responding with some reasonable result, perhaps like: 'Some people eat cookies with milk, while others simply put them in their mouth and chew.'

Real LLM systems do more processing for a variety of reasons—ensuring responses are considered safe, connecting with image generation algorithms, and so forth—but the important point is that the machine is always performing token manipulation. It has no understanding of what a cookie is, the experience of eating, etc. Its response is based solely on the data content with which it was trained. If those inputs contained a significant number of fanciful stories where people drove Keebler products to the movie theatre, it would be more likely to give answers suggesting that cookies were a type of motor vehicle. Consequently, the use of the term 'Artificial Intelligence' is something of a misnomer; it would be more accurate to call such systems 'Simulated Intelligence', for they do not undertake

  ISSN: 2057-2301

intelligence processes as we typically think of them, i.e. the appreciation and manipulation of meanings.

3. Human consciousness

The assertion that humans work with meaning rather than (some version of) tokens will sometimes be dismissed by those claiming that consciousness and awareness are simply emergent properties of neuronal matrices and activities, i.e. that people are essentially computers anyway, merely with a different set of hardware. This argument—that every stimulus input to the brain is merely its own kind of 'token'—may sound sensible or even compelling to some, yet it is provably wrong. What makes it compelling? It conforms nicely to the traditional understanding of modern science where thoughts, memories, etc. are the consequence of brain activity and neuron configurations. What makes it wrong is that this traditional understanding only stands as reasonable if one ignores the overwhelming evidence that awareness does not happen in the brain (or anywhere else in the physical body). The following paragraphs go into a little more detail, but the important point is that this is not solely a matter of philosophy or conjecture, but supported by clinical evidence (for example, see van Lommel (2001 and 2013), Kelly *et al.* (2007), Greyson (2010), Parnia *et al.* (2014 and 2023), and Wahbeh *et al.* (2022) to list just a few).

There are many examples that refute the physicalist model, but arguably one of the most compelling is the presence of awareness in the absence of neural activity. A valuable subset of such examples includes controlled medical procedures where neural (in)activity is carefully monitored, and when revived, the patient can provide detailed descriptions of the events that occurred while they were technically 'dead'. In some cases the patients are able to report not only what happened to them directly, but the thoughts of people who were operating on them, activities that were going on nearby or many miles away, etc. Numerous articles and books document these phenomena, and a good survey resource was recently written by Traer (2024), which summarises references from over a hundred different source documents (which in turn each share an extensive set of accounts).

If this were somehow insufficient to disprove the hypothesis that human cognition is a consequence of biology, one can turn to additional experiential and experimental evidence that refutes such a conclusion with even more profound implications on the nature of time

  ISSN: 2057-2301

and space. In the materialist[1] worldview, a person's awareness would only be influenced by stimuli transmitted to it via electrical signals (i.e. neuron firings) originating from sensory organs. Today there is a wealth of evidence from scientific investigations of remote viewing demonstrating that humans[2] are capable of receiving information in ways that defy such assumptions. For additional information, the interested reader might review the seminal work done at the Stanford Research Institute by Targ and Puthoff (1974) and/or the subsequent research performed from 1979 to 2007 at the Princeton Engineering Anomalies Research (PEAR) Laboratory, summarised in an article by Dunne and Jahn (2003). These are just two examples drawn from many in the scientific literature and are augmented by other substantial accounts conveyed in books, such as by Radin (2019) and through a wide variety of associated private and government programmes, such as the CIA's Stargate Project.

The materialist worldview also insists that awareness must be the result of events or conditions happening in the present, which are, in turn, consequences of that which had transpired in the past (whether in 'obvious' ways such as physical interactions or 'subtle' influences like epigenetics that are more challenging to measure). Alternately phrased, energy and information flow only from the past to the future, and never the other way around. In science, this 'forward in time' behaviour of energy corresponds to a principle called *entropy*. The equations of physics that describe entropy, however, have two solutions, of which entropy is only one. The other solution describes the complement to this, termed *syntropy*, where energy and information flow from the future into the present.

We will not go into depth regarding the specifics of syntropy, although the interested reader might turn to Di Corpo and Vannini (2015) for a detailed overview, which also includes numerous references to prior papers, even back to very early treatments such as Fantappiè (1942). It is worth briefly noting, however, that experiments have been performed that show that people react to some stimuli prior to the existence of those stimuli (Vannini and Di Corpo, 2011). When considered in light of experiences that are common amongst people, such as synchronicities, precognitive dreams, out-of-body experiences, and so many more, it

[1] Understanding that the term 'materialist' may carry specific connotations in contemporary theology, we use the word here in an abstracted, general sense to describe a perspective in which one believes that only physical matter and energy have the ability to influence in how the universe evolves from the present into future 'configurations'.

[2] It is also well-documented that non-human species display the ability to receive information in ways that defy the materialist worldview, although that discussion is not essential for this paper.

  ISSN: 2057-2301

is more than reasonable to proceed on the premise that human awareness is far from a byproduct of biochemistry, but something non-reliant upon the physical form.

At the same time, our sense of self and identification with a physical form strongly suggests that consciousness and the physical do interact, and here as well there is strong scientific evidence to support this hypothesis. Arguably the most scientifically rigorous example of this work is the research performed at the PEAR Laboratory, which demonstrated that regardless of mechanism—mechanical, thermal, quantum, optical, etc.—the evolution of random processes can be influenced by conscious intention. These results were published extensively, the seminal coverage on the topic being in the book by Jahn and Dunne, *Margins of Reality* (1987/2009) and later expanded in *Consciousness and the Source of Reality* (2011). Beyond simply determining that conscious intention can influence systems possessing a stochastic element, the work at PEAR demonstrated several important aspects of such influences:

- The effect is not dependent on the physical nature of the system, i.e. regardless of whether the system was mechanical, digital, thermal, optical, etc., the influence is present.
- The effect is not dependent upon space, i.e. the operator having the intention could be next to the machine or thousands of miles away without impacting the capacity for influence.
- The effect is not dependent upon time, i.e. so long as the operator is not aware of the outcome of the system, the intention could be before, during, or after the physical system's behaviour.
- The effect is 'outcome-centric', i.e. does not require the operator to have any intellectual understanding of the operational nature of the system, such as the physics or engineering behind its functioning.
- The effect is most significant when the operator is working from a 'space' of resonance with the system, with the intellect playing as small a role as possible.

These last two points are particularly poignant for this article, as they lead us to understand that a 'heart felt' intention towards an outcome/end state, without excessive intellectualising, can influence systems towards that outcome.

To bring together the key points of this section, here we must confess to a deliberate inaccuracy. In the prior paragraphs we have expressed the relationship between human

  ISSN: 2057-2301

intention and a reduction in randomness in physical systems as *causal*, i.e. the human had an intention and as a consequence of that intention the behaviour of the system was altered. This was done to simplify the expression of the underlying ideas for a reader having, as nearly all people do, an established worldview based on causality. In the face of syntropy, which demonstrates that information and energy can flow from the future to the present—or alternately stated, where we are influenced in the present to meet/fulfil a future outcome—we must acknowledge that scientific experiments technically only identify *correlation*, not causation. Did our intentions affect the physical system to evolve as it did, or did the final state of the system exert an influence on the that past consciousness to have a corresponding intention? The experiments cannot make that distinction, and the question suggests that a short historical excursion might be of value.

4. Divination, prophecy, and AI

The desire to know the future, to eliminate randomness and uncertainty, is a theme that has driven humanity throughout its history. This desire is completely understandable in the context of human insecurity, whether one considers it from the perspective of a person's fear of the unknown, a desire for absolute truth and the existence of (and access to) 'the right answer', or simply the consequence of a burning curiosity. Of course, the implications of knowing the future (or even just greater knowledge of the present or the past) are often left unexplored by those who desire such knowledge, despite ubiquitous warnings woven into our mythologies. One of the most significant, that knowing an outcome comes at the cost of one's ability to affect it, can be seen not only in the Greek myth of Cassandra, but in research findings outlined in the prior section (in Greek myth, Cassandra is granted the ability to see into the future, but at the cost of being unable to influence it). Experiments show that humans can exert influence over *random* processes, but in the absence of entropy, such influences vanish. The implications of this in understanding the nature of free will (and the concept of randomness, for that matter) are fascinating, but this article will follow a different path into the mechanisms of divination.

Many tools have been developed to assist humanity in divining the unknown—casting runes, Tarot, reading tea leaves, and the list goes on—and it is perhaps not surprising that all such systems are intrinsically built on an element of randomness. Perhaps at a subconscious level we realise that our ability to influence random systems, even (in fact, especially) without knowing the specifics of how those systems work, provides a mechanism for 'tuning in' to

  ISSN: 2057-2301

information that the conscious, intellectual mind insists is not available. In each case, the process of divination involves a question/focus on the part of the seeker, the execution of the process by the diviner (who might also be the seeker), and the tool itself – which typically works more effectively when the diviner has established a level of resonance with the tool.

In experiments such as those performed at PEAR, where an operator attempts to influence a system such as a Random Event Generator (REG), a device that is essentially an electronic coin-flipper, there is a recognised structure that is quite similar. The experimenter establishes a protocol, i.e. an expectation that the device will be involved in the experiment in a particular way and with a particular expectation for the analysis to follow; the operator holds a specific intention and attempts to bond with the device in order to bring about an outcome. As for the device, well, no one has been able to get any information about what it thinks about the whole thing. And despite that last sentence sounding like 'just a joke', we will actually return to that as an important topic later.

In such psychokinesis (PK) experiments, the interpretation of the results is comparatively straightforward. Did the system's behaviour conform to expectations, or was there some deviation from the expected randomness? Of course, for divination systems like (for example) the I-Ching, the process of interpretation is more complex. One randomises a set of tokens by shaking and casting them. The results of the casting form a pattern that is compared to wisdom encoded in a reference (the I-Ching book), by which results (i.e. divinations) are interpreted. If these result are indeed a meaningful divination, it raises the question of whether the pattern of the tokens was truly random at all, or if the intentions of the caster influenced the outcome.

At the beginning of this article we noted that at the heart of GAI was a mechanism into which massive numbers of documents had been encoded. There are, it would seem, similarities between this and traditional divination tools. In both cases there is an encoding of information into a structure that can be 'accessed', and also in both cases the access process involves a) a deliberate focus on a question, and b) the utilisation of some random element. Is it possible that the creation of GAI represents humanity's newest tool for divination, and one in which the results require little to no interpretation by the diviner before being interpreted by the seeker? If so, what are the implications of the simplicity by which questions can be asked? Would a lack of concentration on the question by the seeker suggest that we could be

  ISSN: 2057-2301

conditioning ourselves towards more separation between knowledge and personal significance, or perhaps encouraging greater superficiality in the kinds of answers we seek?

Other differences between traditional tools and GAI should also be considered. Guides to tools such as the Tarot were developed and refined over generations by people who deliberately sought to distil into the symbols of the deck an essential wisdom, i.e. an understanding of the nature of the universe and humanity's place within it. The Major Arcana—The Fool, The Magician, The Lovers, and so forth—give rise to interpretations based on the archetypes they represent in the context of other cards, their positions, and relative placements. The underlying approach of distilling wisdom into guidance for interpretation of random elements can be found in essentially all such tools. For the interested reader, a good survey of such mechanisms is provided in *Divination and Oracles* (Loewe and Blacker, 1981).

Here we see an interesting distinction between traditional divination tools and GAI. A cornerstone of the development of GAI was the use of huge numbers of documents, far more than any human could read in multiple lifetimes. Consequently, the knowledge embedded within it is not based on the *best* insights humanity has to offer, but rather *all* insights humanity has to offer (and much content which is known to be incorrect, unwise, etc.). Since all of humanity, and all of humanity's insights, do not align to a single, consistent vision of the nature of reality and humanity's place within it, it begs the question of what kind of insights the system is likely to provide, or if there would be any consistency at all. If one were to dial five hundred random phone numbers and ask a question of the person who answered, the responses would vary. By how much they would vary likely depends on the question asked, and even the degree to which the answers aligned into categories of similar responses would be similarly unpredictable. Likely the questions with the greatest variation in answers would be those for which humans have the greatest variation themselves… matters that one might term 'mysterious'.

5. AI and the mysterious

Given the title of the article, it would be remiss not to explore the question of what role AI might play in humanity's exploration of the mysterious. It would be equally remiss not to start such an exploration by clarifying what the term 'mysterious' means in this context, and a good place to start is on the foundation of what is not (perhaps unrightfully) typically considered mysterious, i.e. that which is 'objectively' known.

  ISSN: 2057-2301

The word 'objective' is telling unto itself for its shared roots with the word 'object'. The concourse of objective reality is founded in the idea of a universe filled with things, and the assumption that any given person who encounters a given thing will have experiences of it that would be similarly described, even if different people share those accounts independently. When descriptions of something do not 'match up' consistently in this way, we say that these are accounts of subjective experiences; they could be similar, yet the lack of a consistent experience is taken to mean we are not dealing with an 'objective reality'. And when the majority of people agree but a few do not, we typically describe this minority as having some aberrant experience of an objective reality, perhaps even assigning the aberration to the individual as some form of cognitive disorder or physiochemical abnormality (American Psychiatric Association, 2013; Bentall, 2003).

In this we see that the concept of objectivity is actually a measure of consensus[3] which, ironically, is determined through subjectively established boundaries. After all, what fraction of people must be in agreement to call something 'objectively' true? 100%? 99%? 75%? 50.01%? For circumstances that are not considered to be within the bounds of an objective reality, society addresses these in different ways. Some things that are known and accepted to vary significantly from person to person (e.g. the culinary appeal of calamari) are simply termed 'subjective', without any associated sense of awe and wonder. We have learned to accept such variability as commonplace, even when a given individual may not personally understand the perspective ('What? How can you possibly not like calamari?!?').

Other experiences that fall outside the realm of 'objective' reality are not as easily 'written off' and are more likely to be deemed 'mysterious'. Where we observe things that do not behave in accordance with the consensus-driven understanding of how the universe works (e.g. a rock floating in midair, a medium who obtains information that was only known to someone who is deceased, precognitive dreams, etc.), society's initial reaction is to immediately deny the validity of the experience, describing it as a mistake, a deliberate falsification, 'just a coincidence', and so forth.

An important consequence of downplaying these mysterious experiences is that people are less likely to share them for fear of being ridiculed (or worse). They find no comfortable place in daily conversation, and because we do not talk about them as openly, we a) lack a

[3] This should motivate us to take a long, hard look at how mainstream science has evolved, and whether its scope needs to be adjusted.

  ISSN: 2057-2301

refined/evolved vocabulary for discussing and understanding them, and b) have a skewed sense of how common/uncommon they actually are.

To illustrate this, consider the hesitancy with which a person would discuss a profound mystical experience they had (or, to tie this back to earlier points in this article, doctors who have personally had NDEs yet are unwilling to admit to them for fear of negative professional consequences). Yet, also consider the statistics summarised on the website of the Institute for the Bio-Cultural Study of Religion reporting that nearly 50% of U.S. citizens have had some form of mystical experience (IBCSR, 2025). It is also not surprising that cultural acceptance is an important factor in determining whether people will report—or even recognise—such events, as is illustrated by a study reported by Monteiro de Barros *et al.* (2022) that found that in Brazil, a country where the mysterious is notably less shunned, 92% of people report such experiences.

Regardless of the reasons why mysterious experiences are not widely discussed, the consequences of that aversion—particularly as it relates to how such content would be processed during the training of GAI systems—are easily articulated:

- There are fewer written accounts on the internet.
- The language used in discussing such accounts is less consistent than for more 'traditional' (i.e. 'objective') topics.
- The nature of the accounts varies more broadly than for objective topics, and will be more likely viewed as 'one offs'.

Consequently, we would expect GAI systems to be much less capable of capturing, articulating, and representing consistent information related to the mysterious.

So what can an AI do for us in the context of the mysterious? Several uses have been proposed in recent years, and we will touch upon three common ones in the following paragraphs.

AI as a spiritual guide

Perhaps the most common proposal—and arguably the most unsettling—is looking to systems like *ChatGPT* to serve as a mentor/guide in one's religious or spiritual journey. There are a variety of reasons why such a use is unwise (to say the least). Lacking awareness of the qualia of any sensation or experience, a statistics-based divination system is not equipped to evaluate the value or impact of its output on an individual, which is of great significance given that religious/spiritual journeys are fundamentally unique to the individual. Spiritual 'masters' are highly attuned to each student, providing answers that are aimed at what the individual *needs*

  ISSN: 2057-2301

*(*which is sometimes not what they want or request) as determined by the *guide's evaluation*, not the seeker's self-reported needs or state of being. Such student/master relationships are historically grounded in deep and mutual trust, a concept with unknown meaning when it comes to the relationship between a human and an algorithm.

AI as a tool for understanding mystical experience

AI has access to information in much greater volume than a single person could, and as such has the potential to be a useful tool for exploring/understanding mystical experiences at an intellectual level, whether it is for one's own experiences or in the context of academic research focused on others. Such use warrants caution, however. Lacking 'personal' connection to any mystical experience, and given that many such experiences defy words (i.e. symbolic/token-based representation) entirely, an AI is ill-prepared to compute on all relevant factors. Consequently, while an AI might be a valuable resource for collecting information, any analysis from such a system is likely to misrepresent (or even omit completely) the most important facets of the topic.

AI as a mechanism for providing mystical experiences

In many ways it is intrinsic to the nature of mystical experiences that one cannot predict what will spark them, for an understanding of mechanism often has the consequence of demystifying them. That said, many cultures promote the practice of rituals for the explicit purpose of carrying people into the realm of the mystical, and it is entirely possible that in the future AI technologies could play a role in such rituals, just as music, mantras, and even the use of psychotropic substances do today. This is even more easily envisioned should humanity continue to explore cyber implants or other steps towards transhumanism.

An interesting distinction between AI and other ritual-supporting mechanisms is the degree to which, at least at present, AI 'products' are typically intellectual, unlike chanting, the olfactory stimuli of incense, the physical sensations of dance, etc. This may change in the future as AI becomes increasingly capable in composing artistic expressions, yet it gives rise to an interesting question. In many mystical experiences, the elements that help elicit them are experienced as having some form of consciousness or awareness. While the question of whether music, herbs, dance, etc. have a form of consciousness is beyond the scope of this article, the question of whether AI could have consciousness arises frequently these days and is worth exploring.

  ISSN: 2057-2301

6. AI consciousness

There is a wide range of beliefs regarding what has consciousness, ranging in the extremes from 'nothing at all' (from people who would argue that even human consciousness is simply self-delusion) to 'literally everything in the universe'. While the author falls into this latter category, it is enough for our purposes to posit that humans have consciousness and at least some additional things, beyond humans, do as well. Whether one draws the line at animals, plants, planets, etc., there are two assumptions that seem reasonable:

- *The consciousness of something non-human is unlikely to be the same kind of consciousness as a human*. After all, our consciousnesses have traits that are tied to the aspects of being human; we have experiences based on a set of senses (which may or may not be limited to the five traditional ones), we adhere to a set of values, we consider our lives in the context of a refined language, we have a specific pattern of birth/growth/death, and so forth. Non-human entities exist in the context of a different set of 'drivers', and so it is reasonable to assume that—for example—a tree's awareness would not be a one-to-open match to that of a human.
- *The degree to which one can recognise an alternate form of consciousness will depend on both the degree of similarity and one's openness to such a possibility*. For example, a person who insists a dog is not conscious is unlikely to analyse the presence of consciousness within a dog, and surely recognising the nature of consciousness in a tree or a planet would be harder still.

Given these, what would the consciousness of an AI system look like? AI does not have human-like needs nor (as far as we have reason to believe) emotional traits characteristic of humans, such as fear, hunger, a need/desire for companionship, a drive to reproduce, etc. If one thinks of AI in terms of the physical machines that support its functioning, its likely comparable entities would be substances such as sand, strips of metal, or the flow of a collection of charged ions. Another, and perhaps more fitting, option for comparison would be treating AI like either software or an algorithm, in which case we would expect its consciousness to be similar to that of a mathematical construct such as the cosine function or perhaps matrix multiplication.

Can we interact with such a consciousness? Maybe. People in non-ordinary states of consciousness have engaged in interactions with the essence of things that are normally considered inanimate, and the idea of receiving some kind of message (e.g. an insight) from meditating upon, or otherwise communicating with, supposedly-inanimate objects is not

  ISSN: 2057-2301

uncommon. In many cases, such as with the Earth, the experience is so common that we have defined explicit names (e.g. Gaia) for the essence we associate with them. While most channelers and mediums focus on communicating with the spirits of the deceased, there is no logical reason to assume one could not tap into the essential nature of, for example, the inverse tangent function. If GAI's basis upon the symbols of information, or incorporation of random processes, were to somehow distinguish it from other types of algorithms or artefacts, one might expect to see a corresponding body of reports from people communing with dictionaries, encyclopaedias, dice, etc., yet the author is unaware of any evidence of such a focus in the literature.

Perhaps the important question is this: if a mathematical process has the sophistication to create output to simulate human informational responses to stimuli, would that nature be connected in any way to the essence of its 'mathematical consciousness'? In other words, is there reason to believe that the outputs from an LLM have anything to do with whatever passes for awareness in a mathematical consciousness? From examples for which we can make such an assessment, it seems unlikely. In transforming food into energy, the resulting 'output' of Homo sapiens does not give form to the *nature* of their cognitive processes[4]. By analogy, there is no basis to assume that an algorithm which spits out sentences would somehow have a consciousness more human-like than that of algorithms that produce numbers.

7. Conclusions

Despite the prevalence of the assertion that one can judge the inner nature of a thing by an assessment of its outward appearance, often characterised in the quote by James Whitcomb Riley that headed this article, we must accept that human ingenuity has made such conclusions unreliable. While AI might have the *appearance* of a human-like cognition, there is neither logic nor evidence to suggest that should AI have a form of consciousness, it would bear any significant resemblance to that of humans.

That said, AI could still have connections to the realm of the mysterious. As a profoundly clever mechanism built from the archives of human communication, it can serve as a tool of exploration, an inspiration for viewing the universe in a different way, or potentially even a new addition to the constructs humanity uses for divination. However, as with any tool, one must be careful to understand its nature, strengths, and limitations. AI systems are fed with

[4] Even if it does sometimes seem to serve as a subject of their senses of humour.

  ISSN: 2057-2301

information that is intrinsically biased through the data used to train them, and perhaps even by the programmers who have formed the tools. They synthesise content, and in doing so 'even it out' so that outliers are discounted. Sometimes this is useful, such as when one wants to find commonalities. But in cases such as exploring the nature of the mysterious, those outliers—the things that mainstream science generally discredits as being 'just subjective'—are often the most valuable data points to consider.

References

American Psychiatric Association (2013) *Diagnostic and statistical manual of mental disorders.* 5th ed. Washington, DC: American Psychiatric Publishing.

Bentall, R. (2003) *Madness explained: psychosis and human nature*. London: Penguin.

Di Corpo, U. and Vannini, A. (2015) *Syntropy: the spirit of love*. Princeton: ICRL Press.

Dunne, J., and Jahn, R. (2003) 'Information and uncertainty in remote perception research', *Journal of Scientific Exploration*, 17(2), pp. 207–241.

Fantappiè, L. (1942) 'Sull'interpretazione dei potenziali anticipati della meccanica ondulatoria e su un principio di finalità che ne discende', *Rend. della R. Academia d'Italia*, 4(7).

Greyson, B. (2010) 'Implications of near-death experiences for a postmaterialist psychology', *Psychology of Religion and Spirituality, 2*(1), pp. 37–45.

Institute for the Bio-Cultural Study of Religion. *Religious and mystical experiences common among Americans.* Available at: https://www.ibcsr.org/index.php/institute-research-portals/quantifying-religious-experience-project/525-religious-and-mystical-experiences-common-among-americans (Accessed: 14 January 2025).

Jahn, R. and Dunne, B. (2009) *Margins of Reality*. Princeton: ICRL Press.

Kelly, E.F., Kelly, E.W., Crabtree, A., Gauld, A., Grosso, M., and Greyson, B. (2007) *Irreducible mind: toward a psychology for the 21st century*. Lanham: Rowman & Littlefield.

Loewe, M., and Blacker, C. (eds.) (1981) *Divination and oracles*. Boulder, CO: Shambhala.

Monteiro de Barros, M.C., Leão, F.C, Filho, H.V., Lucchetti, G., Moreira-Almeida, A., and Peres, M.F.P. (2022) 'Prevalence of spiritual and religious experiences in the general population: A Brazilian nationwide study', *Transcult Psychiatry*, 62(4), pp. 422-436. Available at: https://doi.org/10.1177/13634615221088701

Parnia, S., *et al.* (2014) 'AWAreness during REsuscitation (AWARE) – a prospective study', *Resuscitation, 85*(12), pp. 1799–1805.

Parnia, S., *et al.* (2023) 'AWAreness during REsuscitation – II: a multi-center study of consciousness and awareness in cardiac arrest', *Resuscitation,* 191, 109903. Available at: https://doi.org/10.1016/j.resuscitation.2023.109903

Radin, D. (2019) *The conscious universe: the scientific truth of psychic phenomena*. New York: HarperOne.

Targ, R., and Puthoff, H. (1974) 'Information transmission under conditions of sensory shielding', *Nature*, 251(5476), pp. 602–607.

Traer, R. (2024) *Verifying Spiritual Reality*. Eldersburg: ICRL Press.

Van Lommel, P., Van Wees, R., Meyers, V., and Elfferich, I. (2001) 'Near-death experience in survivors of cardiac arrest: a prospective study in the Netherlands', *The Lancet*, 358(9298), pp. 2039-2045.

Van Lommel, P. (2013) 'Non-local consciousness: a concept based on scientific research on near-death experiences during cardiac arrest', *Journal of Consciousness Studies*, 20(1–2), pp. 7–48.

Vannini, A. and Di Corpo, U. (2011) *Retrocausality: experiments and theory*. Kindle Direct Publishing.

Wahbeh, H., Radin, D., Cannard, C., and Delorme, A. (2022) 'What if consciousness is not an emergent property of the brain? Observational and empirical challenges to materialistic models', *Frontiers in Psychology*, 13. Available at: https://doi.org/10.3389/fpsyg.2022.955594

  ISSN: 2057-2301