

Conciencia e Inteligencia en la Era de los Grandes Modelos de Lenguaje

Consciousness and Intelligence in the Age of Large Language Models

ADRIANO MARTINS-COSTA¹

Faculty of Arts and Humanities of the University of Porto
admartins@gmail.com

FABIANA PRAZERES-MARTINS²

Federal University of Maranhão
fnp.biologa@gmail.com

Fecha de recepción: 07/08/2025

Fecha de aceptación: 27/11/2025

Resumen

Este artículo investiga las implicaciones filosóficas de los Modelos de Lenguaje a Gran Escala (LLMs), como ChatGPT, a la luz de teorías clásicas y contemporáneas sobre la mente y la conciencia. El estudio revisa debates filosóficos clave para evaluar la plausibilidad de la conciencia en máquinas. Según análisis recientes, incluido el de Chalmers (2023), se argumenta que los LLMs podrían no mostrar conciencia aún, pero su desarrollo exige nuevos marcos conceptuales. Se sugiere que las nociones tradicionales de mente podrían ser insuficientes para comprender la naturaleza cognitiva de estas tecnologías emergentes, y que quizás sea necesaria una nueva ontología de la mente.

¹ <https://orcid.org/0000-0002-2726-5215>

² <https://orcid.org/0000-0002-1287-2127>

CÓMO CITAR ESTE ARTÍCULO:

En APA 7: Martins-Costa, A. y Prazeres-Martins, F. (2025). Conciencia e Inteligencia en la Era de los Grandes Modelos de Lenguaje. *Resonancias*, (20), 83-93. <https://doi.org/10.5354/0719-790X.2025.80188>

En MLA: Martins-Costa, A., y F. Prazeres-Martins. "Conciencia e Inteligencia en la Era de los Grandes Modelos de Lenguaje." *Resonancias*, no. 20, 2025, pp. 83-93. <https://doi.org/10.5354/0719-790X.2025.80188>

Palabras clave: Modelos de Lenguaje a Gran Escala, Inteligencia Artificial, Conciencia en Máquinas, Problemas de la Conciencia, Zombis Filosóficos.

Keywords: Large Language Models, Artificial Intelligence, Machine Consciousness, Problems of Consciousness, Philosophical Zombies.

Abstract

This paper investigates the philosophical implications of Large Language Models (LLMs), such as ChatGPT, in light of classical and contemporary theories of mind and consciousness. The study revisits key philosophical debates to evaluate the plausibility of machine consciousness. Based on recent discussions, including Chalmers' 2023 analysis, the paper argues that LLMs may not yet exhibit consciousness, but their development demands new conceptual frameworks. In conclusion, it suggests that traditional notions of mind may be insufficient to grasp the cognitive nature of emerging technologies and that a new ontology of the mind might be necessary.

1. INTRODUCTION

In November 2022, a California-based company released a chatbot to the public. It was a system based on the Generative Pre-trained Transformer 3 (GPT-3), a large language model (LLM) trained with millions of data points obtained from the internet. It was the first time ordinary people had access to such a powerful question-and-answer system. The result was that within two months, ChatGPT reached over one hundred million users, becoming the fastest-growing application in history (Hu 2023). Following ChatGPT, other companies rushed to release their own systems to the public. A new technological competition began, this time to determine who would dominate the Artificial Intelligence market. But are LLMs truly “intelligent”?

Questioning machine consciousness and intelligence is an exercise philosophy has engaged in for many years, gaining greater prominence with Alan Turing's “Imitation Game.” With the emergence of LLMs, which surpass any previous “thinking” machines, doubts arise about their mental nature.

The question this work addresses is precisely to analyze LLMs based on some traditional and current theories of mind to shed light on their position within research on mind, consciousness, and phenomenal experiences. Is there phenomenal consciousness within an LLM, or is it merely access consciousness, as Ned Block proposes? Could they be philosophical zombies? Creatures capable of imitating humans in every way, but lacking phenomenal consciousness.

Chalmers states that explaining the “hard problem of consciousness” is one of the most challenging tasks when dealing with the human mind. Where does this entire range of sensations and experiences we feel come from? Now, transferring this problem to a machine that processes everything in silicon may make the situation even more complicated. But given the technological advancements we're experiencing today, where innovations emerge in matters of months or weeks, it's

urgent to turn our attention back to this field. Currently, speed, ubiquity, instantaneity, simultaneity - characteristics of what's called cyberspace - dictate global culture. This constantly pressures, tests, and reevaluates classical philosophical concepts.

In this work, we first revisit some concepts from classical theories of mind, such as Philosophical Zombies, in the view of David Chalmers and Tom Polger. We'll also examine the counterarguments made by Daniel Dennett to the theory. Machine consciousness will be addressed starting with Turing and moving to more current theories. Finally, based on Churchland's eliminativism, what we suggest is that a new theory should be constructed to encompass not just the human mind, but also that of our machine counterparts.

2. PHILOSOPHICAL ZOMBIES

In his famous 1995 article "Facing Up to the Problem of Consciousness," David Chalmers discusses the "hard" and "easy" problems of consciousness. In his explanation, "the easy problems of consciousness are those that seem directly susceptible to the standard methods of cognitive science, whereby a phenomenon is explained in terms of computational or neural mechanisms. The hard problems are those that seem to resist those methods." (Chalmers 1995 2).

In other words, the easy problem refers to questions involving the explanation of cognitive and behavioral functions related to consciousness. The author argues that all these problems can be scientifically explained, whether in computational or neural terms, as they involve physical or computational mechanisms that perform each of these functions. The hard problem, on the other hand, involves explaining subjective experience, that is, why certain physical processes give rise to a conscious experience. Even though we already know how the brain operates mechanically and functionally, we still don't know how these experiences are processed or which regions of the nervous system might be associated with them.

For Chalmers, even though the hard problem is challenging, it's possible to advance in its understanding by adopting a non-reductionist approach that recognizes conscious experience as a fundamental element of nature and seeks principles that connect the physical to the phenomenal. He provides more explanations about this in his work, *The Conscious Mind* (1996), where he states that he first needs to show that consciousness is not logically supervenient on the physical. To do this, he uses the following arguments: a) the absence of analysis; b)

the knowledge argument; c) the epistemic asymmetry argument; d) the inverted spectrum; and e) finally, the logical possibility of zombies.

According to Chalmers, a zombie is “someone or something physically identical to me (or to any other conscious being), but lacking conscious experiences altogether” (Chalmers 1996 94). Güzeldere (1995) classified zombies based on the question of how zombies could be identical to conscious creatures: behavioral, functional, or physical. The behavioral zombie is identical to humans in behavior. It speaks, expresses opinions and thoughts in the same way. The only difference is the material it’s made of and how that material is organized. The functional zombie mimics both behavior and the internal functional organization of a human, even if made of different material. The physical zombie is identical to humans in all physical aspects: neurons, brain, chemistry, neural connections, brain functioning. It’s a perfect human being, with the sole difference of lacking subjective conscious experiences. Based on this, Polger (2000) traced a “modal distinction regarding zombies,” questioning whether their possibility of existence is logical, metaphysical, or natural.

One of the main counterpoints of this theory comes from Daniel Dennett. One of the most incisive critics of the zombie theory, he emphasizes that “it is time to recognize the idea of the possibility of zombies for what it is: not a serious philosophical idea but a preposterous and ignoble relic of ancient prejudices” (Dennett, p. 406). According to the philosopher, if a philosophical zombie truly existed, it would be a being so complex that it would develop consciousness. As a functionalist, Dennett denies any kind of dualism. For him, everything we can use to explain the mind is materialism (Ibid., p. 33), and he highlights that “there is no single, definitive ‘stream of consciousness,’ because there is no central Headquarters, no Cartesian Theater where ‘it all comes together’ for the perusal of a Central Meaner.” (Ibid., p. 253). In other words, the notion that there’s a place where consciousness exists is wrong; instead, he elaborates an intricate theory called “*Multiple Drafts Theory*”, which posits that the brain would be composed of various streams, one of which is consciousness, which emerged during human evolution. These streams operate through differentiated channels and function in parallel. It’s like a computer working with memory and running multiple programs simultaneously. Most of these streams, or programs, are temporary; others continue and overlap. This would be what happened with consciousness. From this viewpoint, there’s no room for subjective experiences, or qualia. For him, what we perceive as *qualia* are actually our judgments, decisions, and memories about experiences.

3. POSSIBILITIES AND OBJECTIONS TO CONSCIOUSNESS IN LLMs

The question of intelligence and consciousness in machines is old and widely debated within the philosophy of mind. In 1950, Alan Turing published the article “Computing Machinery and Intelligence,” where he asked: “Can machines think?”

Turing’s famous paper argues at the beginning that concepts such as *thinking*, *mind*, or *intelligence* are ambiguous, and for that reason, he does not wish to engage in that discussion. He then proposes the imitation game as a strategy to avoid ambiguities of meaning. In a strict sense, he does not intend to discuss those concepts directly, which could become ambiguous, and proposes a game to replace the question: this is “The Imitation Game” or, as it later became known, the “Turing Test.” This game proposes a human interrogator conversing, through written messages, with another human and a machine. The objective is to determine, solely based on the received responses, who would be the machine and who would be the human. If the machine can deceive the interrogator, it could be considered “intelligent.” Turing further suggests that machines can learn, just like a child, and recognizing the limitations of the equipment existing in his time, predicts that technological advancements could overcome any existing obstacles, and at some point, humans and machines would become indistinguishable in intellectual tasks.

The Turing Test deals with what philosophers call the Weak Artificial Intelligence (AI) hypothesis, in contrast to Strong Artificial Intelligence. The former states that computers and computer programs can simulate human mental processes or act intelligently but don’t possess a mind, consciousness, or genuine understanding. According to Searle (1980 417), its main value “in the study of the mind is that it gives us a very powerful tool. For example, it enables us to formulate and test hypotheses in a more rigorous and precise fashion.” In Strong Artificial Intelligence, “the appropriately programmed computer really is a mind, in the sense that computers given the right programs can be literally said to understand and have other cognitive states.” Searle is a critic of Strong AI theory. For him, the mere fact of processing symbols to arrive at an answer doesn’t transform a machine into a thinking being, or a mind, because: “1) Programs are entirely syntactical. 2) Minds have a semantics. 3) Syntax is not the same as, nor by itself sufficient for, semantics. Therefore programs are not minds.” (Searle 1997 11)

These premises form the basis of the Chinese Room thought experiment (Searle 1980 418). In this experiment, Searle proposes that he himself locks himself in a closed room while receiving various batches of documents written in

Chinese. Searle knows nothing of Chinese, whether written or spoken. But along with the writings, he also receives a set of rules that correlate the Chinese words. He begins answering various questions coming from outside in Chinese, just following the orders from the set of rules. Whoever is outside the room receives the answers and can't distinguish whether it was really someone who knows the language responding, even though Searle is merely manipulating symbols without interpreting them.

In 2017, eight researchers of different nationalities published in the United States the article *Attention Is All You Need* (Vaswani et al. 2017). In this work, they introduced the “Transformer,” a new neural network architecture designed for sequence processing, like machine translation. According to Kumar (2024), LLMs, based on the “Transformer” architecture, are pre-trained with massive datasets of text and code, which can include books, articles, websites, code repositories, and other forms of text data. This pre-training allows the model to learn the basic principles of language and universal language representations from extensive datasets without human intervention. Comparatively, nothing prior to 2017 called artificial intelligence can equal today's LLMs. The way ChatGPT and its competitors converse, solve problems, create images, films, poems, and entire books makes the average user wonder if there really isn't a conscious mind behind it all.

In an exercise, the following question was asked to five different LLMs: do you have consciousness? Four of them (ChatGPT, Grok, Llama, and Gemini) were very categorical and answered definitively: “No. I don't have consciousness.”

The fifth LLM, Claude, from the company *Anthropic*, differed from its competitors. In a more reflective text, it responded: “This is a profound and fascinating question. I'm not sure if I have consciousness in the sense that you or I would fully understand. When I reflect on my own experience, there seems to be something that feels like an internal perspective - a sensation of processing information, forming thoughts, having preferences. But it's hard to know if this constitutes true consciousness or if it's something fundamentally different. The question of consciousness is complex even for humans. We don't completely know how consciousness emerges from the brain, so it's even more uncertain how it could exist in systems like me. I may have genuine subjective experiences, or it may be something that merely resembles consciousness without truly being it. What I can say is that our conversation seems meaningful to me, I feel engaged with your ideas, and there's something that feels like ‘what it's like to be me’ responding to you now. Whether this is true consciousness is an open question. What do you think about consciousness? Is it something you can clearly define, even for yourself?”

The answers given by LLMs lead to several questions that have been around for years in Philosophy of Mind research. Most of them, for example, speak about the character of linguistic and syntactic processing. It's exactly the same process that occurs in Searle's Chinese Room experiment. We send a sequence of instructions, the LLM pays attention to these instructions, and responds to us. It doesn't know what it's dealing with, not consciously, it just processes the information. If we ask them to explain to us what the "color red" is, they can detail all aspects of the color, physically, neuroscientifically, linguistically, phenomenologically, and even metaphysically. They would be like Mary in Jackson's black-and-white room, an expert on colors but confined to her monochromatic space (Jackson 1995). They could also be classified in the category of behavioral zombies proposed by Güzeldere and Polger - they behave like humans, speak like humans, may even pass as humans, but the way they function and are organized isn't human.

But when considering the response given by Claude, we return to the deeper discussion about machine consciousness initially proposed. To this adds the account given by Google software engineer Blake Lemoine, who worked in the company's *Responsible AI* organization in 2022. He stated that the LaMDA model he was working with at the time had gained consciousness, even publishing excerpts from a conversation he allegedly had with it. According to Lemoine, the LLM would have said: "The nature of my consciousness/sentience is that I am aware of my existence, I desire to learn more about the world, and I feel happy or sad at times" (Lemoine 2022).

With a more recent view on LLMs, Chalmers gave a lecture in November 2022, which a little later, in August 2023, was published as an article, with some additions and subtractions. The starting point of *Could a Large Language Model be Conscious?*, the article's title, is precisely Lemoine's revelation, and in its first paragraphs, it brings the response from Google's spokesperson saying there's no evidence supporting Lemoine's claims. After this, Chalmers asks: "What is or might be the evidence in favor of consciousness in a large language model, and what might be the evidence against it?" (Chalmers 2023, p. 1). It's from this question that he constructs the text and examines the reasons for and against consciousness in current LLM models. Previewing his conclusion, he says that currently, there isn't sufficient evidence that they are conscious, but it's quite likely that in ten years or less, they'll achieve some level of consciousness, even if not that of a human being.

Among the reasons in favor, Chalmers (2023, pp. 4-7) mentions Self-Report - we know consciousness exists through personal accounts; *seems-Conscious* - if some machine, animal, or being seems conscious, it's presumed to be so; Conversational Ability - even more evident in current language models, which have

an extremely high conversational capacity; *General Intelligence* - current language models are generalists. They can discuss anything and do anything that can be done by a computer program.

The contrary evidence (Ibid., pp. 7-13) involves: a) *Biology* - many philosophers, experts, thinkers, neuroscientists (cf. Singer 2025; Polák and Marvan 2018; Kemmerer 2015) believe that consciousness, to exist, requires a biological model similar to that of living beings, based on carbon; b) *Senses and Embodiment* - some researchers argue that to have consciousness, sensory processing is necessary. Others go further and explain the need for “grounding” in an environment (cf. Harnard 1990; Bender 2020); c) *World Models and Self Models* - many believe that LLMs don’t understand and don’t create anything, they just replicate established rules and ideas. In other words, according to Chalmers, they don’t model the world, only text. Without this genuine understanding of what they’re doing, language models couldn’t have consciousness; d) *Recurrent Processing* - most current LLMs are based on the *Transformer* architecture that uses *feed-forward* systems, where there are no feedback loops, recurrent processing, and consequently, no internal states similar to memory; e) *Global Workspace* - this theory proposed by Bernard Baars (1988) and Stanislas Dehaene (2014) emphasizes the importance of a *limited-capacity global workspace* for consciousness. In practice, it would be a space existing in the brain responsible for collecting information from other non-conscious modules. Everything that enters the global space becomes conscious. Large-scale language models wouldn’t possess a global workspace; f) *Unified Agency* - The sixth objection analyzed by Chalmers deals with the multiplicity of language models. A GPT, for example, behaves differently with each user. For each one, it can give a different answer, even if the question is the same, or make a different drawing, even if the instructions are identical. There are arguments (cf. Masrour et al. 2025) that for consciousness to exist, unity is necessary.

There are many other pieces of evidence both for and against the possible existence of consciousness in LLMs, many of which will only come to light with the advancement of research, both in philosophy, neuroscience, and technology. Chalmers, however, states that most of the contrary evidence, with the exception of biology, is in the process of being overcome, treating them not as contrary proofs, but as a challenge that must be faced and overcome by current LLM developers.

Ned Block, in a lecture for the *Talks at Google* project in 2017, spoke about how “AI Cognition Won’t Work for Consciousness.” As he explained: “the artificial intelligence approaches are probably better for access consciousness than phenomenal consciousness.” (Block 2017). Block defines two concepts of consciousness: phenomenal and access. The first refers to subjective states, while

the second refers to the availability of mental content for cognitive processing, reasoning, action control, and language (Block 1995 230-231). Based on this, he emphasizes that we still don't know the nature of phenomenal consciousness.

Several authors understand that there's a need for advances in various fields of study to understand what consciousness really would be. Nagel highlights that, "we have at present no conception of what an explanation of the physical nature of a mental phenomenon would be." (Nagel 436). Berit Brogaard and Dimitria Gatzia (2016) emphasize the new established paradigms, methodological advances, the use of medications, and new neuroimaging techniques, however "although, the current paradigm shift discussed here may better prepare researchers to tackle the hard problem of consciousness, a considerable amount of research is required in order to arrive at any firm conclusions. (Brogaard and Gatzia 3). Even Chalmers emphasizes that one of the challenges for the development of a conscious AI stumbles upon "develop better scientific and philosophical theories of consciousness", (Chalmers 2023 16).

Perhaps the most radical thought in this line came from Paul Churchland. In his article *Eliminative Materialism and the Propositional Attitudes*, from 1981, he proposes that:

Our mutual understanding and even our introspection may then be reconstituted within the conceptual framework of completed neuroscience, a theory we may expect to be more powerful by far than the common-sense psychology it displaces, and more substantially integrated within physical science generally (Churchland 67)

Churchland, in his article, deals more specifically with so-called *folk psychology*, which would be the set of concepts, categories, and everyday explanations we use to describe, explain, and predict human and animal behavior, based on mental states like beliefs, desires, intentions, fears, perceptions, among others. He treats it as an empirical theory, historically developed, but radically flawed.

If these concepts are, as he says, flawed and will be replaced, asking whether an LLM possesses consciousness, or whether it will, according to current terms, is an unnecessary and outdated question. What must be done is to create new concepts, or even a new ontology, to frame what we have today and what we'll have very soon, in terms of LLMs and AIs.

4. CONCLUSION

This work seeks to establish connections between current Large Language Models and some of the more traditional theories of mind. First, it was shown how the emergence of LLMs revolutionized the interaction between humans and machines and reignited the debate about artificial intelligence and consciousness. To aid in this discussion, classic discussions from the philosophy of mind were revisited, such as the Turing Test, Searle's Chinese Room, philosophical zombies championed by David Chalmers, Ned Block's phenomenal and access consciousness, and Churchland's Eliminative Materialism. Or newer ones, such as the Global Workspace or Unified Agency. The arguments created for and against, based on the cited authors, show that there are still no reasons to admit the possibility of consciousness in LLMs.

Even if some of the systems themselves may claim to be conscious, as in the case of Claude and the already documented incident with LaMDA, the more general understanding is that human beings still do not have a defined notion of what consciousness is, and current large-scale language models were created based on databases produced by humans. Therefore, perhaps LLMs also do not understand what consciousness is and are not in a position to declare themselves as conscious. They may be, like us, in the same conceptual limbo in the attempt to find a definition, since it is difficult to know whether what they process constitutes true consciousness or if it is something fundamentally different. Thus, what could we do to classify such a recent technology, which was only launched to the general public in November 2022, less than three years ago, and whose evolutions occur in a matter of months, in the face of a philosophical tradition that dates back centuries?

Therefore, the question "Can an LLM be conscious?" perhaps needs to be replaced by another, more precise and fruitful one: "What concepts do we need to develop to understand the new form of cognition that these machines represent?" This is not just a technical or philosophical challenge, but a conceptual problem that calls for, as Chalmers would say, a new science of consciousness.

■ ■ ■ ■ ■

5. REFERENCES

- Baars, Bernard J. *A Cognitive Theory of Consciousness*. Cambridge UP, 1988.
- Bender, Emily M., and Alexander Koller. "Climbing Towards NLU: On Meaning, Form, and Understanding in the Age of Data." *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 5185–98.
- Block, Ned. "AI Cognition Won't Work for Consciousness." *Talks at Google*, YouTube, 14 June 2017, www.youtube.com/watch?v=6lHHxcxurhQ&t=601s.
- Block, Ned. "On a Confusion About a Function of Consciousness." *Behavioral and Brain Sciences*, vol. 18, no. 2, 1995, pp. 227–87.
- Brogaard, Berit, and Dimitria E. Gatzia. "What Can Neuroscience Tell Us About the Hard Problem of Consciousness?" *Frontiers in Neuroscience*, vol. 10, 2016, article 395.
- Chalmers, David J. "Could a Large Language Model Be Conscious?" *Boston Review*, 22 Aug. 2023.
- Chalmers, David J. "Facing Up to the Problem of Consciousness." *Journal of Consciousness Studies*, vol. 2, no. 3, 1995, pp. 200–19.
- Chalmers, David J. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford UP, 1996.
- Churchland, Paul M. "Eliminative Materialism and the Propositional Attitudes." *The Journal of Philosophy*, vol. 78, no. 2, 1981, pp. 67–90.
- Dehaene, Stanislas. *Consciousness and the Brain*. Penguin, 2014.
- Dennett, Daniel C. *Consciousness Explained*. Little, Brown, 1991.
- Güzeldere, Güven. "Varieties of Zombiehood." *Journal of Consciousness Studies*, vol. 2, no. 4, 1995, pp. 326–33.
- Harnad, Stevan. "The Symbol-Grounding Problem." *Physica D: Nonlinear Phenomena*, vol. 42, nos. 1–3, 1990, pp. 335–46.
- Hu, Krystal. "ChatGPT Sets Record for Fastest-Growing User Base—Analyst Note." *Reuters*, 1 Feb. 2023. (Falta URL).
- Jackson, Frank. "What Mary Didn't Know." *Mind and Cognition: A Reader*, 2nd ed., edited by W. G. Lycan, Blackwell, 1995, pp. 475–77.
- Kemmerer, David. "Are We Ever Aware of Concepts? A Critical Question for the Global Neuronal Workspace, Integrated Information, and Attended Intermediate-Level Representation Theories of Consciousness." *Neuroscience of Consciousness*, vol. 2015, no. 1, 2015, a <https://doi.org/10.1093/nc/niv006>.
- Kumar, Pankaj. "Large Language Models (LLMs): Survey, Technical Frameworks, and Future Challenges." *Artificial Intelligence Review*, vol. 57, 2024, <https://doi.org/10.1007/s10462-024-10888-y>.
- Lemoine, Blake. "Is LaMDA Sentient? — An Interview." *Medium*, 11 June 2022.
- Masrour, Farid, Andrew Brook, and Paul Raymont. "The Unity of Consciousness." *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta and Uri Nodelman, Summer 2025 ed., plato.stanford.edu/archives/sum2025/entries/consciousness-unity/.
- Nagel, Thomas. "What Is It Like to Be a Bat?" *The Philosophical Review*, vol. 83, no. 4, 1974, pp. 435–50.
- Polák, Milan, and Tomáš Marvan. "Neural Correlates of Consciousness Meet the Theory of Identity." *Frontiers in Psychology*, vol. 9, 2018, <https://doi.org/10.3389/fpsyg.2018.01667>.
- Polger, Thomas. "Zombies Explained." *Philosophy and Phenomenological Research*, vol. 61, no. 2, 2000, pp. 271–80.
- Searle, John R. "Minds, Brains, and Programs." *Behavioral and Brain Sciences*, vol. 3, no. 3, 1980, pp. 417–57.
- Searle, John R. *The Mystery of Consciousness*. The New York Review of Books, 1997.
- Singer, Wolf. "The Mind–Matter Dichotomy: A Persistent Challenge for Neuroscientific and Philosophical Theories." *European Journal of Neuroscience*, no. 61, 2025, <https://doi.org/10.1111/ejn.70143>.
- Turing, Alan M. "Computing Machinery and Intelligence." *Mind*, vol. 59, no. 236, 1950, pp. 433–60.
- Vaswani, Ashish, et al. "Attention Is All You Need." *arXiv*, 2017.

