Thinking Machines, Speaking Minds: Language, Philosophy, and Artificial Intelligence – A Case Study

Hiba Harb

DOI: 10.18355/XL.2025.18.03.08

Abstract
The article discusses the intersection of philosophy, language, and artificial intelligence (AI), focusing especially on AI-generated language and its role in education. It discusses the ever-changing capabilities of AI, especially in terms of its ability to create meaningful conversation, and how these systems (ChatGPT for example) imitate human thinking without really thinking (intentionality). It considers some of the key philosophical conversations like the difference between strong and weak AI, and some of the ethical and epistemological issues connected to AI-generated language, when thinking about authorship, responsibility, and trust. An important research gap is also addressed, by considering how students will interpret and take in language generated via AI, especially when it pertains to educational contexts. The article presents findings from a survey of students and their perception of credibility, ethics, and AI's role in knowledge generation. The article discusses how while students may find the utility of AI, they remain aware of the ethical issues tied to the use of AI, especially issues connected to authorship. The article concludes with potential directions for future research, such as the development of a new philosophy of machine language, in order to better consider AI, cognition, and education. The article adds to the debate on the role of AI in education, and how to use it ethically and productively in academic contexts.
Key words: Artificial Intelligence, Language, Thought, Machine, Philosophy

1. Introduction

Since Ancient civilizations, the question of thought, language, and their articulation has always been highly considered in philosophical debates. Greek thinkers had already questioned the nature of the soul, and the human ability to reason and express themselves. Descartes defined thought as proof of human existence in his famous Cogito, ergo sum (Wittgenstein, 1953). Nowadays, these questions have intensified with the progression of the language sciences and cognitive psychology and, more recently, artificial intelligence (AI). At present, people are experiencing a unique integration of different fields: the philosophy of mind, formal linguistics, epistemology, computer science, and educational sciences are coming together all around the same object of study: the thinking machine. This revolution is not simply a technical transformation, but a paradigm shift: modern technologies and AI-supported tools, and in particular large language models like ChatGPT, have surpassed significant barriers (Winograd & Flores, 1986). In addition to the ability of these AI tools to calculate or automate tasks, they can also generate language, support conversations, produce arguments, and even narratives. It appears through the output of these tools that they give an impression of autonomous thought, a kind of intelligence in action. This rupture is raising deep questions: Can machines think? Can they speak meaningfully? And what does it really mean to "think" or to "speak" in this case? In addition to these challenges, there is an observable challenge in educational and social contexts: the use of AI in learning is changing how we deal with knowledge and intellectual authority, and how we produce content (Rushkoff, 2019). Teachers, researchers, and students are now confronted with texts produced by non-human authors that challenge notions of originality, understanding, and accountability,

and they are also comparing the human ability to reflect, and the algorithmic ability to predict.

The objective of this article is to analyze this change in philosophical and pedagogical points of view, and to highlight not only how these text-generating AI work, but also how their human users (especially students) come to interpret and evaluate these texts. The educational lens allows for critical observation of how the representations of language, intelligence, as well as communication from/for younger generations continue to evolve in the digital landscape (Penrose, 1989). The research gap this article targets to fill lies in the literature concerning AI, not only in the adaptive and responsive capabilities of the technology, but also its implications in pedagogical and philosophical perspectives from the experience of learners. Our methodological approach is based on a comprehensive conceptual analysis, linked with theoretical concepts from the philosophy of language (Wittgenstein, Searle, Austin), linguistics (Chomsky, Saussure), and Humanities and Technology (Dreyfus, Heidegger, Dennett), and enhanced with a field study with university students in order to catch real experiences of AI discourse in practice (Searle, 1999). This work aims to contribute to critical reflection on the status of language in the age of strong AI by bringing together a discussion among philosophy, the humanities and technology in light of the demonstrated, and growing implications of the matter. It is also worth asking how we might rethink our perception of meaning, consciousness, and truth, if words are no longer created by humans only?

2. The Philosophical Foundations of Language and Thought

The link between language and thought has been an ongoing topic in the history of philosophy, raising varied and often contradictory notions as to how words, ideas, and consciousness are alternating with one another. Understanding the contexts from which artificial intelligence has begun generating discourses requires the recognition of the underlying philosophies that set the definitions for language and thought (Turing, 1950).

2.1 Presentation of Key Thinkers

Rationalist René Descartes puts forth the notion thinking is the core of a human being ("cogito ergo sum") and, therefore, only human beings possess true language. For Descartes, a Mechanism may mimic (and perhaps even behave like) mechanical behavior, but it cannot produce language that is meaningful as it does not have mind or consciousness. Ludwig Wittgenstein, in his Tractatus Logico-Philosophicus and later in Philosophical Investigations, advocates a more pragmatic view of language, positioned alongside Cartesian dualism (Kauffman, 2009). Wittgenstein argues meaning is not an inherent essence, but a function of use: "the meaning of a word is its use in the language (Wittgenstein, 1953)." In this way, speaking is more than communicating information but rather seen as engaging others in a language game, and a social context. Noam Chomsky supports a more formal linguistic tradition; he uses the term universal generative grammar and asserts the capacity for language is innate to all humans as derivative of what he calls mental structures (Piatelli-Palmarini, 1989). According to Chomsky, language serves as a window into the mind, and discovering its workings requires accessing the underlying mental configurations. However, Michel Foucault offers a critical and contextual approach to languages and to understanding knowledge (Varela et al., 1991). In The Order of Things, he demonstrates that language operates across fields of power, structuring discourses and truth regimes of a particular period of time. Furthermore, according to Foucault's view, language is never neutral: it contributes to the construction of subjectivities and organizes reality (Minsky, 1985).

2.2 The Role of Language in the Construction of Reality and Consciousness

Language is more than a mean of communication: it shapes our relation to the world, it outlines our consciousness and defines our thoughts. As Wittgenstein states: "the limits of my language mean the limits of my world.". Furthermore, language plays a role in the construction of reality, by producing conceptual divides and solidifying collective representations (Johnson-Laird, 1983). In terms of consciousness, language helps to integrate thought, creates internal dialogues, and allows for self-reflection. It is difficult to imagine a fully developed self-aware consciousness without language. This question is particularly timely, as it emerges in the context of discussions about artificial intelligence: can a machine that lacks subjectivity truly "understand" what is produced, or simply reorder signs based on statistical likelihood?

2.3. The Link Between Mind, Meaning, and Intentionality

The term of “Meaning” is at the center of philosophical discussions about language. For humans, it is often associated with intention: the ability to 'aim' at something based on thought or language. The philosopher Franz Brentano and John Searle articulated that intention as a necessary condition of consciousness (Searle, 1980). For example, to speak about something meaningfully implies an intention to direct that intention to something. Yet this 'intention' does not seem to be included in contemporary artificial intelligence systems. Language models do not aim at anything: it orders words without 'understanding' what it makes or produces. We can then distinguish between the lived-intentioned meaning (i.e., represented meaning; subjective; conscious), or a determined meaning (produced; statistical; and mechanical). This distinction is precisely what leads to questions about machine discourse and its value (Schmidt, 2018).

3. Artificial Intelligence and the Simulation of Thought

The concern of whether a machine can think or understand creates profound dilemmas between technical performance and its philosophical conceptualization. In its historical development, artificial intelligence has turned the corner from strict logical systems to models that learn to either produce language or solve complex problems (Langer, 1942). But is that enough to grant it a trace of true intelligence, or even consciousness?

3.1 Evolution of AI: From Symbolism to Machine Learning

The early work on artificial intelligence between the 1950s and the early 1970s was grounded in the symbolic approach, a view that intelligence was the manipulation of symbols based on explicit logical rules. Expert systems were a typical example of this perspective and rested on so-called decision trees and "if... then" symbolic rules. While this approach provided significant and appropriate influence, it was also relatively quick to show its limitations when confronted with the complexity and ambiguity of natural language or human thinking (Gray, 2020). From the 1990s onward, and the emergence of machine learning, then deep learning, in the 2010s onward, AI faced a significant change. Instead of programming rules, there are now training statistical models with large volumes of data. Artificial learning models have enabled the emergence of systems capable of recognizing images, converting text into other languages or formats, or generating text in a fluent and coherent way such as ChatGPT does. Despite this turn toward high-performance models, it raises an interesting distinction between simulation of intelligence and the truth of thought (Heidegger, 1977).

3.2 Philosophical Debates: Strong AI vs. Weak AI

The philosophical discussions regarding AI reflect two opposing visions of AI: weak AI vs. strong AI. Weak AI supports that machines can mimic certain intelligent behaviors but don't necessarily have any conscious or understanding. However, Strong AI holds that a machine can, in principle, develop a form of true intelligence similar to human intelligence, or some form of consciousness (Brynjolfsson & McAfee, 2014). This difference is significant in the discussions surrounding AIs in the use of a language: is it an imitation of intelligence or a form of artificial cognition?

3.3 Can machines "understand"?

John Searle highlighted this question based on what is now famously called the Chinese Room Argument (1980). Searle imagines a "man who speaks no Chinese" is locked inside a room with a manual for responding correctly to messages in Chinese. His responses would look coherent from outside the room, but in fact, the man is just following the manual and manipulating symbols according to specific rules (Searle, 1999). According to Searle, computers do the same thing by following a set of rules that simulate understanding. This argument is also linked to the Turing's test suggested by Alan Turing in 1950. If a machine can chat with a human, and that human cannot tell whether they are dealing with a human or a computer, then we can say the machine is doing something "intelligent". Although this test relies on behavior, it does not in any way guarantee that the machine understands. The AI may pass the test without a conscious thought or meaning behind their responses (Dennett, 1996). It highlights a paradox: is the observable performance of a system sufficient to discuss intelligence or thought? Or do we assume an inner mind wishing subjectivity for "understanding" is at play? The distinction between syntax (manipulation of symbols) and semantics (understanding meaning) is significant. But again, existing AI tools, as impressive as they may be, seem limited in terms of syntax while the question of understanding remains open.

4. Language, Meaning, and Machine Discourse

As generative AI systems, such as ChatGPT, become more common across educational, professional and social contexts, an important philosophical question comes back into focus: does the language generated by machines have meaning? If it has meaning, what kind? It’s important to review how generative AI systems work, before considering their conceptual implications against the theories of language.

4.1 How language-generating AI systems work

Models like ChatGPT are transformer-type models, which means they do not possess consciousness or deep understanding of the human sense of the “meaning”. They are programmed on massive text corpuses and operate on a principle of prediction, predicting the next word in a sequence of language based on co-occurrence statistics. Their performance is based on their ability to model linguistic, discursive, and pragmatic regularities, even though they never "know" what they are talking about. In other words, ChatGPT has no way to "understand" the content it generates, it approximates it through models of human language, where its effectiveness ultimately hinges on its ability to iterate language as a model of human language (Dreyfus, 1972). However, it does not "understand" the meaning of the texts it produces, even though its production operates linguistically and pragmatically similar to human language.

4.2 Philosophical Implications: Is Generated Language Meaningful?

The ability of an AI to produce coherent and well-structured relevant discourse creates a philosophical question: can there be meaning without understanding? In a Wittgensteinian perspective, meaning is produced through use and context: if a human responds to an AI-generated text in an educational context, can meaning be attributed

in some sense to the AI-generated text? Some contemporary philosophers believe this meaning can be interpreted by humans (Wittgensteinian, 1953). One may still argue that there is no intention or meaning embedded internally within the machine in this process. However, meaning in language does not carry meaning until it is intentionally produced by a subject with consciousness, originality, and related to the world. This is where the fundamental limit of machine language appears; it produces a simulation of meaning, rather than actual meaning (Binns, 2018).

4.3 Differences between human intentionality and statistical prediction

Human intentionality—as defined in phenomenology (Husserl) or in philosophy of mind (Searle)—is a mode of consciousness "is directed toward." We can think about an object, and we can signify an idea (McCarthy et al., 1955). Cognitive orientation toward the world is what generates a link between thought and relationship to language. In addition, language models like ChatGPT possess no intentionality. A machine has no desires, beliefs, or representation of the world. They aim at nothing but statistically predict the most likely outcome given the initial prompt. Regardless of the automatic nature of machine language, it is not manifested from conscious experience.

This distinction is an essential undertaking: human language is an expression of lived-experience, machine language is its likely formal composition. This raises larger epistemological and ethical implications for a machine language: can we trust it as language? Can machine language be used as a source of reasoning or in education?

5. Ethical and Epistemological Dilemmas

The rise of generative AIs like ChatGPT reveals important ethical and epistemological issues. One of the most important challenges goes back to how systems that can create complex contexts will most likely be understood and used in a professional or academic context. In this section, we will address the credibility of the authorship of texts created by the machine, the challenges, and the cognitive place of these machine-generated discourses (Kauffman, 2009).

5.1 Responsibility and Authorship: Who Owns the Text Produced by an AI?

The issues of accountability and authorship of texts made by machines present an essential question: who is accountable for the produced content? Is it the user who formatted the query, the developer of the algorithm, or the AI itself? There is no clear attribution of authorship in the classical sense of the term, beyond the machine, as there is no intent and consciousness attributed to the machine. To illustrate, when a student submits an academic assignment of an answer to a query and generates it through ChatGPT, the student is not simply acting as the mechanism for submitting the task. The student also has accountability, although indirectly, for qualities, accuracy, and ethical construction of the generated content (Floridi, 2020). Certainly, in an academic situation, the human user has assumed the accountability and quality of the information generated by the AI. For example, if the student submits an AI-generated assignment that is incorrect or contains plagiarized content, the student has accountability, even if it is a machine-generated text. Indeed, the student is still responsible for verifying information and using it ethically. This brings up another complication in scientific research (Gergen, 1994). Can a scientist publish their findings by citing a machine if they use AI to generate hypotheses or analyses? This brings into question scientific integrity or transparency, as humans must always be responsible for data validation and verification.

5.2 Risks of linguistic manipulation and disinformation

Another major complication relates to the risks of linguistic manipulation. Since AI writes text from gathered data using statistical models, it can generate texts that are sometimes biased, unintentionally biased, or even modified. AI writing could even be produced using neutral, objective language, yet biased thinking can be hidden through the data set from which it was produced (Hao, 2021). For example, if AI relies primarily on American texts or Western texts, it may repeat cultural, political, or ideological biases in its responses, especially if the question asked relates to international politics or social values. This occurred in 2018 with the AI created by Amazon to sort resumes that discreetly took down female applicants based on the AI dataset. This demonstrates that AI can produce false or misleading results (He & Liu, 2019). Similarly, generative artificial intelligence can disseminate falsehoods at much greater speeds. Malefactors can use these tools to create false or manipulative and misleading narratives. For example, think of disinformation campaigns using an AI system for producing incorrect articles or social media content that make it extremely difficult for individuals to differentiate between the truth and rumors. Thus, regulating the use of generative AI systems becomes a vital step in hindering exploitative or negative behaviors (He & Liu, 2022).

5.3 Can AI be a reliable source of knowledge? Epistemic status of machine-generated discourse

In terms of epistemic status about Machine-generated discourse, another important question arises: Can machines offer valid knowledge? When we think of classic criteria for knowledge like authenticity, justification, and coherence, generative AIs cannot be a source of knowledge like a human. Indeed, AIs are not capable of critical thinking or self-reflection. AIs are only recomposing a pre-existing bodily of information without any measure of validation. For example, an AIs generated answers to a scientific question is based solely on data it has read, and it cannot contextually interpret nor utilize rigorous methodologies for validation (Hirsh, 2016). For example, if an AI were questioned about complicated medical cases and were able to produce an answer that seemed semi-accurate, it might have omitted critical nuances or important information around a medical condition. Human Physicians, while having read the same information, are still capable of analyzing and contextual the medical information based on a patient's medical presentation. This illustrates the basic limitation of AIs: absence of judgment or experience for understanding this world in which they interact. Therefore, AI can exist as a supportive action but is not an independent source of knowledge (Anderson & Dill, 2000). AI can design concepts, generate research summaries, or propose general information, but human-thought is ultimately still required to validate reliability and relevance to the situation.

6. Case Study / Application (optional)

Study Objective

The main goal of this research is to examine students' opinions on language-developing artificial intelligence (AI) as it relates specifically to academia such as ChatGPT. We intend to investigate students' trust in using these AI tools, their interpretation of the meaning and relevance of the language produced by AI tools, as well as their ethics on using AI in the production of Ai-generated material. Additionally, we expect that this project will also expose students' philosophical assumptions about responsibility in machine-generated content (Zhang & Zhang, 2023).

Methodology

The survey includes five questions, and the results will be analyzed with statistical data as well as in-depth analysis about the philosophical implications of each response.

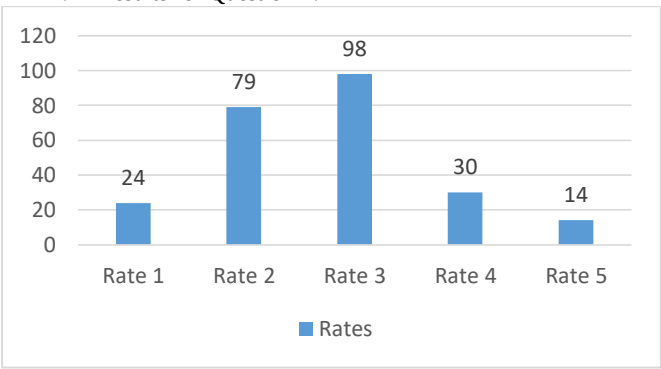
The data will be used to construct students' perception of AI and their approach towards machine-generated language.
The survey was distributed to 245 students over the course of four semesters. All participants were enrolled in the English Language Teaching program at Lusail University in Qatar. They were taking both “Foundations of Education” and “Technology in Education” courses at the time of the survey.

Survey Questions

1. Question 1: "On a scale of 1 to 5, how would you rate the authenticity of information produced by an AI such as ChatGPT in your studies?"
The objective of this question is to measure the trust that students have in AI as a source of academic content.
2. Question 2: "Do you think that AI-generated language can have as much meaning and relevance as that of a human?"
The objective of this question is to investigate the understanding that students have about meaning and relevance of AI-generated language compared to that of a human.
3. Question 3: "Do you have any ethical concerns regarding AI in the creation of academic content (e.g. essays, summaries, etc.)?”
The objective of this question is to understand students' ethical concerns in using AI in an academic context.
4.Question 4: "Do you think that AI understands what it generates?”
The objective of this question is to assess students' understanding of AI, specifically with regards to topics in the philosophy of mind and the nature of intentionality in machines.
Question 5: "Would you be more willing to use AI tools to help you with your academic writing if they acknowledge academic integrity? Why or why not?"
The objective of this question is to explore student acceptance and reflection about using AI for academic tasks related to ethical and liability issues.

Analysis of the results

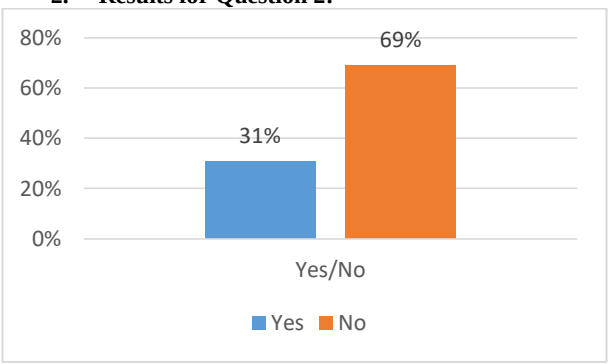
1. Results for Question 1:



Most students rate the credibility of AI-generated information generally low (79 students rated 2, and 98 students rated 3). This suggests their trust level in AI is low and tends to align with the heavy reliance on human sources that have been verified. A few students tended to rate it higher (4/5 - 5/5), perhaps particularly in empirical disciplines where their research suggests the AI responses are more "accurate" and "objective."
These results provide insight into general doubt of AI outputs' reliability, particularly the use of AI in disciplines where human decisions and nuances might matter more. It

also reflects a tendency towards viewing AI as statistical rather than a strong understanding of the content it generates. Students may also have orientation towards how AI builds information and their human verification of that information.

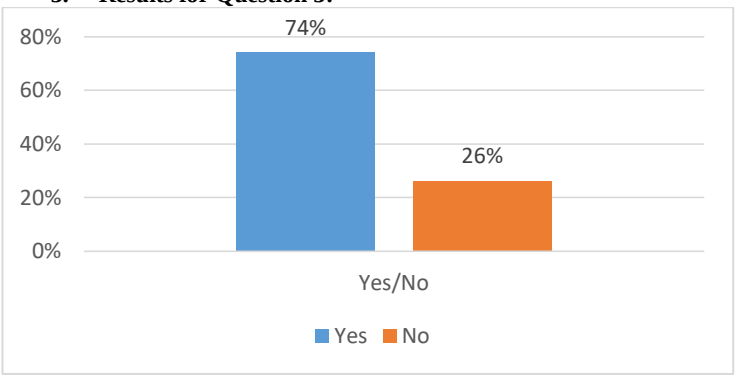
2. Results for Question 2:



The majority of students find that AI does not generate human meaningful language. They argue that while AI can provide correct or rational answers, those answers often lack the nuance, originality, or cultural interpretations that indicate the use of human language. Some students are more receptive to the idea of AI producing language that is completely relevant to the task it is performing, especially in fields that rely on exact science.

This question reveals perceptions of the difference between machine-generated language and human-generated language. The students' responses may highlight philosophical distinctions of language as a tool for communication (reasonably used by AI) and language as a reflection of human communication goals. Philosophers like Wittgenstein, believes that the meaning of language is related to its common use; students may simply feel, even if AI uses human language effectively, that AI is not a part of this socially embedded process of cognitive interpretation.

3. Results for Question 3:

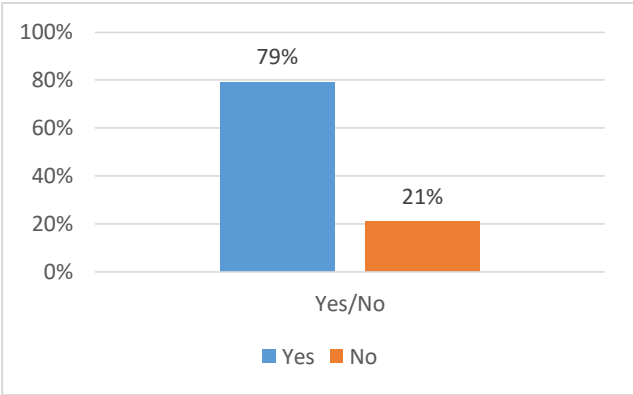


Many students indicate ethical concerns, particularly in regard to plagiarism, authenticity of information generated, or the risks of AI manipulating information. Others worry about the impact of AI on academic creativity, believing that excessive use of these tools may hinder critical analysis and reflection. Furthermore, some students may be more positive and confident that AI can be ethically used, particularly if information is shared clearly and considered as supportive to humans.

The findings demonstrate the gap between the pragmatic use of AI for repetitive tasks and the ethical concerns brought up by risks such as plagiarism or inequity of access to technology. These dilemmas are worth considering in philosophical discussions about authenticity of creation and the author's obligation.

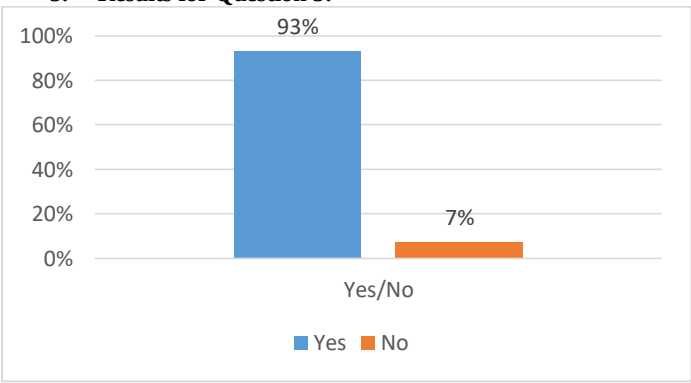
4. Results for Question 4:

A considerable portion of students responded that AI understands the content it creates since it can explain the responses provided and make corrections. However, 21% responded "no" based on the premise that AI has no consciousness, since it's generating language based on rules of an algorithm, not a particular understanding of meaning. However, others may have a less decisive view and argue that AI does not have consciousness like humans but is able to 'understand' it functionally, it means that it is able to respond correctly based on statistical patterns.



Whether students think that AI has or does not have 'understanding' in the human sense goes directly to the heart of philosophical discussions regarding the nature of intelligence and understanding. If students do think that AI does not have an 'understanding' in the human sense, this raises questions related to consciousness and intentionality: to be authentically 'intelligent', does a machine need to have subjective experience that is consciousness? The findings may line up with Searle's Chinese Room argument, indicating that even if an AI were able to simulate intelligent answers, the AI, however, does not authentically 'understand'.

5. Results for Question 5:



Most students support the idea of using AI to write assignments but only with well-defined guidelines such as academic integrity, no plagiarism, and human monitoring. They may further suggest that a human would need to check any content before submitting academic work. Furthermore, some students may oppose the use of AI because of their concerns around knowledge creation or originality to the work.

This question measures the acceptance of the use of AI in academic contexts. If students think AI is a useful tool but needs to be controlled by strict ethical guidelines, this might suggest that students recognize the potential usefulness of AI but still want to maintain principles of intellectual responsibility. Likewise, the complete rejection of its use would be interesting to analyze based on fears involving automation and technological drift.

7. Conclusion

Artificial Intelligence (AI), specifically its language-generation systems such as ChatGPT, invites intriguing philosophical questions that extend far beyond technological applications. In this paper, we have highlighted the notion that AI represents a mimicry, or at least an imitation of human cognitive processes as they relate to language. However, our conclusion must be that AI generates texts that are coherent and fluent but do not demonstrate comprehension nor does it demonstrate linguistics intentions (Chomsky 1957). Indeed, AI accomplishes all of its language-generating tasks through algorithms that process and become so-called 'smart' tools that manipulate data, but without consciousness or subjectivity, which the very basis of what differentiates human from AI cognitive processes. In sum, AI does not actually create meaning, it utilizes its learned models from very large data sets and then produces information accordingly. In general, one of the strongest claims of this analysis is to establish that while AI's linguistic capabilities might produce overwhelming linguistic achievements, AI does not yield understanding, intention, or subjectivity that define human language (Kukulska-Hulme & Shield, 2021). Because this type of philosophical debate (in particular, on the strong vs. weak AI claim) is an important consideration in respect to what AI is, human language and the absence of consciousness is still the main distinction, and artificial intelligence as it exists today does not "understand" language in a human sense, it just mimics statistical models (Kant, 1781). However, this practice does raise fundamental implications of what meaning is and is not, and what AI outputs should be credited as meaning or meaningful. There are some philosophical, ethical, and epistemological considerations related to AI generated textual language production. For example, a text generated by an AI in relation to authorship puts the understanding of intellectual discourse into a point of consideration. AI is a tool to create automatic text generation. To put this into context, the matter of authorship - especially when considering the perceived authorship between a human and a machine- raises concerns to individuals as an intellectual capacity of creating by human executors. In addition, AI's capacity to subtly manipulate language raises concerns about misinformation and manipulation, especially in academic and social realms (Kolb, 1984). Nevertheless, despite these advancements, AI still has recognizable limitations when it comes to imitating human thought and language. The absence of consciousness and subjectivity in current AI systems, precludes the capacity for true "understanding." It imitates linguistic behaviors but does not create new concepts or ideas in a genuinely original manner (Liang & Zhang, 2017). Furthermore, the lack of intentionality or contextual understanding creates challenges for determining the reliability and validity of the actual generated content. Moving forward, future research will need to engage with several avenues. The development of strong AI that would be able to simulate not just linguistic behaviors but also cognitively deeper processes will be notable. A novel philosophy of machine language might emerge, oriented around rethinking how

machines compose and interpret language (Muller & Bostrom, 2021). This philosophy may conceptualize not only the structure of words and sentences but also the underlying intentions, social contexts, and goals that guide machine-generated discourses. Additionally, it would be of value to better understand the implications of ethical issues that arise when utilizing AI in academic contexts. Educational policies will need to evolve to diplomatically integrate AI into instructional practices of teachers while addressing linguistic manipulation and liability issues (Norvig & Russel, 2016). Finally, research could similarly look at the dialectical relationship between humans and AI to understand creativity and problem-solving, and thereby define new genres of collaboration and interaction. In summary, while AI has experienced remarkable advances in generating language, it is not yet capable of simulating human thought or grasp the smaller nuances of language. Future research shall strive to engage in developing intelligent AI systems that are able to consider comprehension and intentionality, as well as the ethical implications, in order to strengthen human-machine collaboration.

Bibliographic references

Anderson, C. A., & Dill, K. E. (2000). Video games and aggressive thoughts, feelings, and behavior in the laboratory and in life. Journal of Personality and Social Psychology, 78(4), 772-790. https://doi.org/10.1037/0022-3514.78.4.772

Binns, P. (2018). Artificial intelligence: A modern approach to education. Journal of Educational Technology, 35(2), 13-19.

Brynjolfsson, E., & McAfee, A. (2014). The second machine age: Work, progress, and prosperity in a time of brilliant technologies. W. W. Norton & Company.

Chomsky, N. (1957). Syntactic structures. Mouton.

Dennett, D. C. (1996). Consciousness explained. Little, Brown and Co.

Dreyfus, H. L. (1972). What computers can't do: A critique of artificial reason. MIT Press.

Floridi, L. (2020). The ethics of artificial intelligence: A critical introduction. Cambridge University Press.

Gergen, K. J. (1994). Realities and relationships: Soundings in social construction. Harvard University Press.

Gray, C. (2020). How the world became a computer: Computational technologies and the transformation of education. Routledge.

Heidegger, M. (1977). The question concerning technology and other essays. Harper & Row.

Hao, K. (2021). The battle for the soul of artificial intelligence. MIT Technology Review.

He, H., & Liu, Z. (2019). AI in education: A systematic review. Education and Information Technologies, 24(5), 2787-2811. https://doi.org/10.1007/s10639-019-09856-4

He, H., & Liu, Z. (2022). AI in education: A systematic review. Education and Information Technologies, 27(3), 2131-2148. https://doi.org/10.1007/s10639-022-10684-x

Hirsh, A. (2016). Understanding AI and its impact on education. Springer.

Johnson-Laird, P. N. (1983). Mental models: Towards a cognitive science of language and thought. Harvard University Press.

Kant, I. (1781). Critique of pure reason (P. Guyer & A. Wood, Trans.). Cambridge University Press, 1998.

Kauffman, L. (2009). Philosophy of artificial intelligence. Oxford University Press.

Kukulska-Hulme, A., & Shield, L. (2021). Artificial intelligence and education: Implications for language learning. Journal of Computer Assisted Learning, 37(5), 1511-1524. https://doi.org/10.1111/jcal.12487

Kolb, D. A. (1984). Experiential learning: Experience as the source of learning and development. Prentice Hall.
Langer, S. K. (1942). Philosophy in a new key: A study in the symbolism of reason, rite, and art. Harvard University Press.
Liang, H., & Zhang, H. (2017). Ethical issues in AI-powered education. Journal of Educational Technology Systems, 45(3), 341-355. https://doi.org/10.1177/0047239517735045
McCarthy, J., Minsky, M. L., Shannon, C. E., & Turing, A. M. (1955). A proposal for the Dartmouth summer research project on artificial intelligence. AI Magazine, 26(4), 12-17.
Minsky, M. (1985). The society of mind. Simon & Schuster.
Müller, V. C., & Bostrom, N. (2021). Future of AI and education: The challenges of developing intelligent education systems. Oxford University Press.
Norvig, P., & Russell, S. (2016). Artificial intelligence: A modern approach (3rd ed.). Pearson.
Penrose, R. (1989). The emperor's new mind: Concerning computers, minds, and the laws of physics. Oxford University Press.
Piatelli-Palmarini, M. (1989). The interactive mind: A dialogue with Chomsky. Cambridge University Press.
Rushkoff, D. (2019). Team human. W. W. Norton & Company.
Schmidt, R. (2018). Artificial intelligence in education: The current state of affairs. International Journal of Educational Technology, 39(4), 45-58.
Searle, J. R. (1980). Minds, brains, and programs. Behavioral and Brain Sciences, 3(3), 417-457. https://doi.org/10.1017/S0140525X00005756
Searle, J. R. (1999). Mind, language, and society: Philosophy in the real world. Basic Books.
Turing, A. M. (1950). Computing machinery and intelligence. Mind, 59(236), 433-460. https://doi.org/10.1093/mind/LIX.236.433
Varela, F. J., Thompson, E., & Rosch, E. (1991). The embodied mind: Cognitive science and human experience. MIT Press.
Wittgenstein, L. (1953). Philosophical investigations. Blackwell.
Winograd, T., & Flores, F. (1986). Understanding computers and cognition: A new foundation for design. Ablex Publishing.
Zhang, T., & Zhang, D. (2023). Ethical implications of AI in education: A survey of challenges and solutions. International Journal of Educational Technology in Higher Education, 20(1), 45-59. https://doi.org/10.1186/s41239-023-00323-w

Words: 5679
Characters: 38 279 (21 standard pages)

Assis. Prof. Hiba Harb, PhD
College of Education and Arts
Lusail University, Jabal Thuaileb PO Box 9717, Lusail
Qatar
hharb@lu.edu.qa
ORCID: 0000-0002-3308-8008