WHAT LLMS CAN'T DO: INTENTIONALITY, AND HENCE CREATIVITY

Bolette Hansen
Lund University
bolettehhansen@hotmail.com

Jordan Zlatev
Lund University
jordan.zlatev@semiotik.lu.se

Alexandra Mouratidou
Maria Curie-Skłodowska University
alexandra.mouratidou@mail.umcs.pl

Abstract

Systems based on large language models (LLMs), such as ChatGPT, produce texts that often seem indistinguishable from human writing, and this has fuelled claims that they are creative. This article argues that they are not.

Digital Age in Semiotics & Communication, Vol. IX, 2026, 71–98
Under CC BY 4.0
https://doi.org/10.33919/dasc.26.9.4

The argument proceeds in three steps: creativity in language presupposes understanding what one produces; understanding requires intentionality; and LLMs lack intentionality. After a historical review of research on "language machines", from *Eliza* and Searle's Chinese Room to connectionist models and present-day LLMs, the article draws on cognitive semiotics and the phenomenology of Merleau-Ponty to explicate intentionality. Using the theory of the Semiotic Hierarchy and the Motivation & Sedimentation Model, it identifies three preconditions for meaningful language use: embodiment, consciousness and situatedness. LLM-based systems lack all three, as they recombine sedimented linguistic forms without bodily experience, reflective consciousness or participation in actual social situations. Their output may thus mimic human language impressively, but it carries no meaning for the system and therefore cannot be genuinely creative. LLMs may nevertheless extend human creativity as tools, provided that human meaning-making is not underestimated.

Keywords: cognitive semiotics, embodiment, large language models, phenomenology, situatedness

1. Introduction

In recent years, systems based on large language models (LLMs) have become hugely popular due to their versatility in supporting a wide range of tasks, ranging from routine information processing to endeavours commonly regarded as creative, such as producing texts in English (and other languages) that are seemingly indistinguishable from those written by people (Wu et al. 2023). By making LLM-based chatbot systems widely accessible and dramatically lowering the threshold for engaging with them, such "Generative Artificial Intelligence" has become a pervasive feature of our current discourse. However, there has also been widespread criticism of this trend. As is well known, such systems regularly produce so-called "hallucinations", i.e., fabrications that have nothing to do with reality but may sound plausible (Wu et al. 2023). Even more seriously, there is widespread anthropomorphising of such systems in both professional and popular literature (Inie, Zukerman, and Bender 2026). This is misleading, given that LLM-based systems lack anything human-like, most importantly any *understanding* of the information that they are processing. Arguably, therefore, research focusing on whether LLMs can *appear* human-like constitutes a "misdirected research effort, specifically around the application of LMs to tasks

intended to test for natural language understanding" (Bender et al. 2021: 615).

Here we build further on such critiques, with the help of *cognitive semiotics* (e.g., Sonesson 2026) and the phenomenological approach to language that it implies (Zlatev 2023), and link them to the debate on whether LLMs can be considered truly creative. The argument is as follows: (1) If a system is to be creative (in its use of language), it must understand what it produces. (2) LLMs lack understanding. (3) Therefore, LLMs cannot be truly creative either.

Originating from the Latin word *creare* ('to make, bring forth, or produce something novel'), creativity remains an extremely complex concept, or rather a family of related concepts (Sternberg 1999). Instead of attempting to tackle it directly, for the purposes of our argument, we posit premise (1) based on the following shared intuition: we would not call a text "creative" unless it was produced by an agent capable of understanding what it means (even if they may not know exactly how they have produced it). As for (2), we will explicate "understanding" through the concept of *intentionality*. We devote Section 2 to a historical review of "language machines" and, in the process, rehearse discussions on *what they can't do* (Dreyfus 1992). In Section 3, we introduce relevant concepts from cognitive semiotics and the phenomenology of sign use and language. In Section 4, we will defend premise (2) and thereby reach the conclusion in (3): LLMs lack genuine creativity. What is not excluded, however, is the possibility that these systems can be used to extend human creativity through *extended agency* (Mendoza-Collazos & Sonesson 2021), understood as the enhancement of human agency with the help of artefacts. We summarise our argument in Section 5.

2. On the history of language machines

LLMs are algorithms that rely on "neural network architectures" and "deep learning". Kernighan (2021: 234, our emphasis) defines natural language processing (NLP) as:

> a sub-area of machine learning that deals with how to make computers process human language – how to *understand* what some text means, summarize it, translate it into another language, convert it to speech (or speech into text), or even generate meaningful text that appears to have come from a human.

But it is exactly this jump from "processing" to "understanding" that is questionable. LLMs are the outcome of some 80 years of technological development in the field of artificial intelligence (AI), where the ability to use

language, in both comprehension and production, has been a central goal. In this section, we review this technological development, schematically dividing it into three stages. As we show, while their versatility has gradually extended with computing power, the essential limitations of "language machines" remain constant.

2.1. Early stage

AI emerged in the mid-20th century, as researchers took an interest in how computers might perform tasks so far restricted to people (Kernighan 2021). In the famous paper "Computing Machinery and Intelligence", Turing (1950) proposed that assessing whether a machine is intelligent can be based on whether it is mistaken for a human being, depending on its ability to engage in fluent interaction in a natural language. This criterion later became known as the "Turing Test".

In 1966, Joseph Weizenbaum developed one of the first chatbots, *Eliza*. Even though it operated by simplistic rules, it proved "surprisingly effective, at least when first encountered" (Kernighan 2021: 234–235). However, Weizenbaum himself noted that this finding reflected the gullibility of people rather than any intelligence in the machine itself. Another appealing idea was that of "machine translation", but attempts to develop rule-based machine translation systems proved unsuccessful, as it became obvious that real human languages and their users do not function like formal languages: linguistic rules are full of exceptions, and language use often does not conform to them (Clark 1996).

But if a more powerful system could pass the Turing Test, would this imply that it was truly intelligent? A strong negative answer to this question was offered by Searle (1980) through his famous thought experiment known as "The Chinese Room". In short, Searle imagines himself in a closed room, manipulating Chinese characters according to a set of rules in English (a language he knows) despite knowing no Chinese. Although his responses may convince outsiders that he understands the language, he merely follows formal rules without any understanding of the meaning of the texts. Although this scenario may appear simplistic, it was received as a strong argument against the claim that NLP systems can possess any kind of understanding, since what Searle in the thought experiment is in fact doing is *implementing* an expert system, consisting of "input", "rules" and "output". And according to the principles of "Turing functionalism", it should not matter whether the implementation is in silicon, neurons, or whatever. All that matters are the algorithms: the rules according to which the system operates. As Searle (1980: 422) concludes: "the main point of the present argument is

that no purely formal model will ever be sufficient by itself for intentionality because the formal properties are not by themselves constitutive of intentionality". Ziemke (2025) presents a recent discussion of these issues, also updated with respect to modern computational systems and robotics.

In relation to this argument, and the disappointing results of early AI systems, Searle (1980) insisted on an ontological distinction between *weak* and *strong* AI.

> According to weak AI, the principal value of the computer 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. But according to strong AI, the computer is not merely a tool in the study of the mind; rather, 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 1980: 417)

In other words, "weak AI" is the use of computer models as simulations of mental processes, which can serve as *tools* to help us understand them better, as in mathematical modelling. However, they are far from being *identical* with these processes. It is "strong AI" that makes the latter leap, by assuming that if the models are not distinguishable from human performance on the behavioural level, then they are *de facto* the same, or at least not qualitatively different.

2.2. Intermediate stage

While early AI systems achieved success in focused domains like chess, they were limited in terms of scalability: success in one domain did not translate to others. The rigidity of such systems posed insurmountable limitations, as the ambiguity of the world proved not to be reducible to "artificial reason" (Dreyfus 1992). In response to criticism of their lack of scalability and flexibility, "artificial neural network" architectures, or *connectionist models*, gained popularity towards the end of the 20th century. Inspired by the functioning of biological neurons, such architectures consist of "nodes", ordered in layers, and connections between them, which, like the synapses of real neurons, can "learn" by changing their strength based on various input-output patterns, interfacing with the "world" (Krohn, Beyleveld, and Bassens 2020; Kernighan 2021; Russell & Norvig 2021). These nodes repeatedly receive input, process the sum, send the output forward, and change the "weights" between the nodes based on feedback. However, such models are highly simplified abstractions of actual neurons, as they operate with the

help of strings of code and calculations rather than actual biology and chemistry (Krohn, Beyleveld, and Bassens 2020; Russell & Norvig 2021). This new technology seemed to breathe new life into AI, as Turing (1950) had already suggested that developing human-level intelligence would be more efficient by "teaching" machines to "learn" from their successes and mistakes, rather than programming intelligence by hand.

As for *meaning*, such systems modelled this in terms of "distributed representations" in the "vector space" defined by nodes of their various layers (Russell & Norvig 2021). These were used to predict the probability of a word given its context or the next word in a sequence in the following manner:

> When we're creating word vectors, the overarching concept is that we'd like to assign each word within a corpus to a particular, meaningful location within a multidimensional space called the *vector space*. Initially, each word is assigned to a random location within the vector space. By considering the words that tend to be used around a given word within the natural language of your corpus, however, the location of the words within the vector space can be gradually shifted into locations that represent the meaning of the words. (Krohn, Beyleveld, and Bassens 2020: 27)

Relationships between words are captured in terms of proximity: the closer two words are within the vector space, the closer their apparent meaning is assumed to be due to the similarity of their contexts. Thus, synonyms, and even misspellings of words, are expected to have near-identical "meanings", since they share identical distributions and occupy nearly identical locations in this space. While new words may not initially be very accurately positioned, over time, with more instances of the newly added word in the dataset, the accuracy of the vector location would improve (Krohn, Beyleveld, and Bassens 2020). That said, there is no guarantee that word vector algorithms capture particular semantic relationships, since any "meaning" is reduced to generalisations over the distribution of words in large datasets (Russell & Norvig 2021).

2.3. Present stage

The creation of *very* large datasets during the last decade was made possible by advances in computing power, as well as by the expansion of the internet, with its vast amount of textual, visual and auditory data (Russell & Norvig 2021). The technologies and algorithms developed earlier are currently applied to corpora of billions of words (Krohn, Beyleveld, and Bassens

2020). As of 2021, "deep learning" models are considered better at generalisations about grammatical and semantic information than traditional machine learning systems (Russell & Norvig 2021). However, one downside of current models is their *opacity*: one is blind to what "features" have been identified and why, and as a result there is no explicit explanation or understanding of system outputs (Kernighan 2021). We return to these issues below and in Section 3, after first considering some of the claimed achievements of this new incarnation of language machines.

Today, NLP is used for various applications such as machine translation (e.g., *Google Translate*), writing assistance (e.g., *Grammarly*), voice assistants (e.g., *Siri* and *Alexa*), text-to-speech and vice versa (e.g., text messages) and voice cloning (e.g., *Speechify* and *Murf AI*). NLP is also at the core of LLMs such as LaMDA, Orca and the famous GPT, which will be used to exemplify our argument throughout the rest of the paper.

When dealing with LLMs, it is important to distinguish between the LLM itself (e.g., GPT-5) and the "virtual agent" (i.e., the chatbot) it inhabits (e.g., ChatGPT and Copilot)[1], though the two are often conflated (Shanahan 2024). The user interacts with the agent, which takes on the role of an "interlocutor". The training material for ChatGPT's language model is a large mixture of text sources. On this basis, a language model with hundreds of billions of parameters across multiple layers is trained (Wu et al. 2023).[2] ChatGPT utilises the previously described deep learning technology to process and generate text in a conversational manner. The user provides a prompt in natural language, in response to which the algorithm processes and generates a statistically probable text based on language patterns, as shown in Figure 1.

The method of *in-context learning* enhances the "conversation-likeness" of exchanges by utilising contextual information from the ongoing interaction. OpenAI introduced this method with GPT-3, enabling the model to learn and complete tasks by "imitation" of previous examples (Wu et al. 2023: 1125). For instance, as demonstrated in Figure 2, a paragraph from a letter draft can be shortened using ChatGPT by providing the task description together with the original paragraph as a prompt. The second prompt

---

[1] LLMs often serve as the foundation for several virtual agents. For example, the chatbots ChatGPT and Copilot are both built upon GPT (developed by OpenAI), while Snapchat and Bing also utilise it for their AI chat functions (Lardinois 2023; Malik 2023).

[2] GPT-3 is built upon 175 billion parameters across 96 layers (Wu et al. 2023: 1123). GPT-4 allegedly has 1.8 trillion parameters across 120 layers (Schreiner 2023). Unlike earlier models, specifications about GPT-5 are withheld.

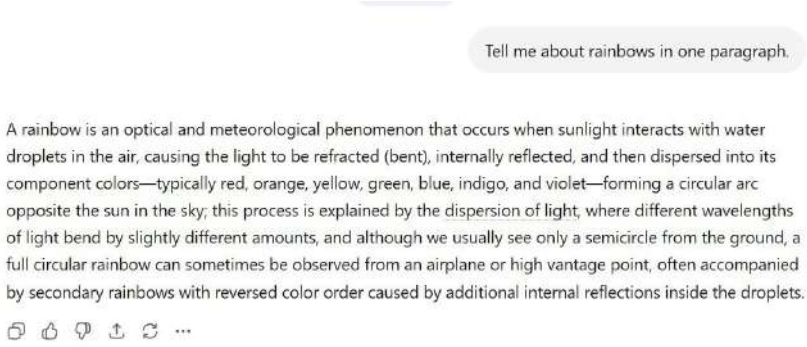

Figure 1: Example of a prompt and the response generated by ChatGPT. Generated on 1 May 2026.[3]

includes only the subsequent paragraph, without a task description; yet the model continues to apply the previous task description to the new input.

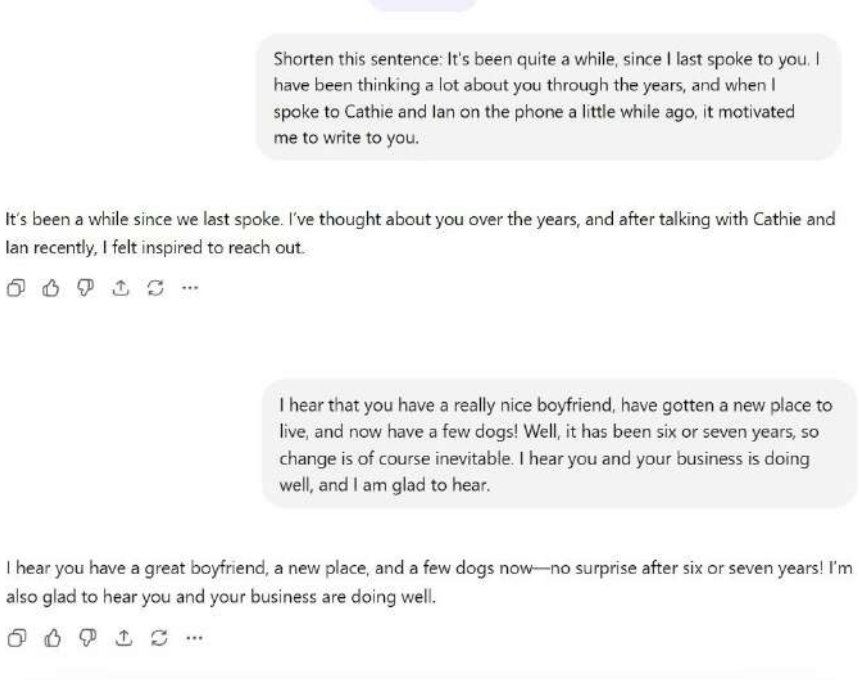

Figure 2: Example of in-context learning in ChatGPT: a task is completed by imitating previous instructions within the conversation thread. Generated on 1 May 2026.

Despite such achievements, this new technology has also been met with much criticism. As stated above, LLMs rely on huge datasets, algorithms and subsequent fine-tuning. At each point, they remain vulnerable to acquiring representational bias and reproducing negative sentiments towards specific groups, whether inadvertently or deliberately (e.g., Bender et al. 2021; Dhingra et al. 2023). Uneven patterns of internet participation are the initial contributing factor to biased training data. As internet access is not evenly distributed, this can result in overrepresentation of certain user groups (Bender et al. 2021). Furthermore, the data is only scraped from a sub-section of the internet. For instance, GPT-2 was trained on data from outbound links from Reddit, a website where users in the US in 2016 were found to consist predominantly of men (67%) and younger individuals (64% aged 18–29). Additionally, structural factors like moderation practices also make user-content sites like Reddit, X and Wikipedia less welcoming towards marginalised populations, despite presenting themselves as open to all. This means that "a limited set of subpopulations can continue to easily add data, sharing their thoughts and developing platforms that are inclusive of their worldviews", reinforcing systematic representational imbalance (Bender et al. 2021: 613). To mitigate such risks, additional filtering (e.g., for appropriateness) has been introduced. However, an "algorithmic definition of what constitutes inappropriately sexual communication will inherently be concordant with some views and discordant with others", implying that LLMs inevitably favour some perspectives above others (Bender et al. 2021: 615). And adding guardrails to the models introduces another layer of bias.

To improve adaptability (by providing up-to-date information) and accuracy (by reducing "hallucinations"), newer models increasingly incorporate retrieval-augmented generation (RAG) architecture. This enables chatbots to access and integrate real-time data external to the base LLM. However, the LLM does not in itself have access to real-time data but is "fed" external sources (Belcic & Stryker 2026). Only a limited number of external sources are retrieved, and they are not necessarily the most reliable or relevant. At the same time, source attribution significantly enhances user trust and perceived credibility (Ravi & Sindhgatta 2025), which may in turn foster false confidence and over-reliance on system outputs.

All these issues, however, deal with how to improve the technologies based on LLMs as a form of *weak AI*: useful tools that may resemble human performance and could further serve to understand, and even enhance it. The main problem is a different one: the misplaced attribution of genuine meaning and understanding to such systems and thus confusing them with *strong AI*. As pointed out above, even if a language machine enhanced with

the latest LLMs were to pass the Turing Test, that would not imply any kind of understanding, since LLMs and their chatbots are stuck in their "Chinese rooms" (Searle 1980). Though their machine-produced text can *appear* meaningful, it is produced by manipulating linguistic *expressions*, void of any true understanding of what they actually mean. This was pointed out by Searle (1980) and Dreyfus (1992) for earlier NLP systems and holds no less for modern LLMs, which are essentially based on "a system for haphazardly stitching together sequences of linguistic forms it has observed in its vast training data, according to probabilistic information about how they combine, but without any reference to meaning: a stochastic parrot" (Bender et al. 2021: 617). Like parrots mimicking human speech, LLMs mimic human language with a degree of randomness and uncertainty, stemming from their nature as data-driven, "black box" deep learning models (Wu et al. 2023). But nothing in this implies any kind of understanding.

At the same time, the performance of LLMs is increasingly being benchmarked on tasks that require sensitivity to meaning and understanding, such as evaluations of ChatGPT on the US Bar Examination (Katz et al. 2024) and Theory of Mind tasks (Kosinski 2023). As stated in the introduction, such efforts are both misplaced and alarming, as they can only enhance the tendency towards anthropomorphising these language machines: attributing to them human features that they fundamentally lack (Inie, Zukerman, and Bender 2026). And the key feature, central for our argument, is that of *intentionality*. As stated above, Searle (1980) appealed to this to help specify why NLP machines necessarily lack understanding of their linguistic inputs and outputs. Unfortunately, he left this key concept insufficiently explained, insisting that "it is a product of the causal features of the brain" (1980: 417), without specifying these in any detail. That is why in the next section, we introduce some key concepts from cognitive semiotics to help fill this gap, before returning in Section 4 to show why "computers *still* can't do" intentionality (Dreyfus 1992).

3. The nature of meaning

3.1. Cognitive semiotics and phenomenology

Cognitive semiotics combines concepts, theories and methods from semiotics, linguistics, cognitive science and phenomenology in order to provide a more integrative approach to the concept of *meaning* (e.g., Sonesson

2012, 2026; Zlatev 2015; Konderak 2018; Oakley & Zlatev 2024). In this section, we will focus on how several of its theories help explain what intentionality is, and why it is a precondition for the meaningful use of language, and indeed for any sign use.

Cognitive semiotics is strongly influenced by *phenomenology*: "the study of human experience and of the ways things present themselves to us in and through such experience" (Sokolowski 2000: 2). One of the key insights is that meaning presupposes embodied and socio-culturally embedded being-in-the-world (Merleau-Ponty [1945] 1962; Sokolowski 2000; Zahavi 2019). *Intentionality* is one of the foundational concepts of phenomenology, where it is understood as the directional nature of consciousness, providing our "openness to the world" (Merleau-Ponty [1964] 1968: 37). It is not to be confused with the psychological notion of goal-directedness, which is only one form of intentionality. Rather, it is an essential feature where "[e]very act of consciousness, every experience, is correlated with an object" (Sokolowski 2000: 8). When we perceive, remember, anticipate, or imagine any intentional object – a cup, a horse, a friend, a mathematical formula – our consciousness directs itself towards this object, allowing it to appear, under one *kind of* intentionality or another: as perceived, remembered, imagined or signified. As described in the following sub-sections, phenomenology is at the core of cognitive semiotics, including its central concept of *semiosis*, understood as meaning making in general, and not as sign use, which is only one particular manifestation of it.

3.2. Levels of semiosis: from animation to language

The model of the *Semiotic Hierarchy* is particularly useful for our argument, since it aims to provide a unified framework for the analysis of different types of meaning and their interrelations (Zlatev 2009, 2018; Zlatev & Konderak 2023). One of the main ideas is to understand intentionality and semiosis as reciprocal, as shown in Figure 3: while the former is what connects us to the world, the latter is its inevitable flip side, the significance of this experience for an embodied, embedded subject.

A recent version of the model (Zlatev & Konderak 2023) distinguishes between five levels of this relationship, as shown in Figure 4. Higher levels stand to lower levels in a relation of *Fundierung*: they are built upon and presuppose lower ones, but at the same time are not reducible to them, as each level introduces some qualitatively new semiotic property. Further, there is a process-structure duality to each level, described by Merleau-Ponty (1962) as a *dialectics of spontaneity and sedimentation* (Zlatev 2018).

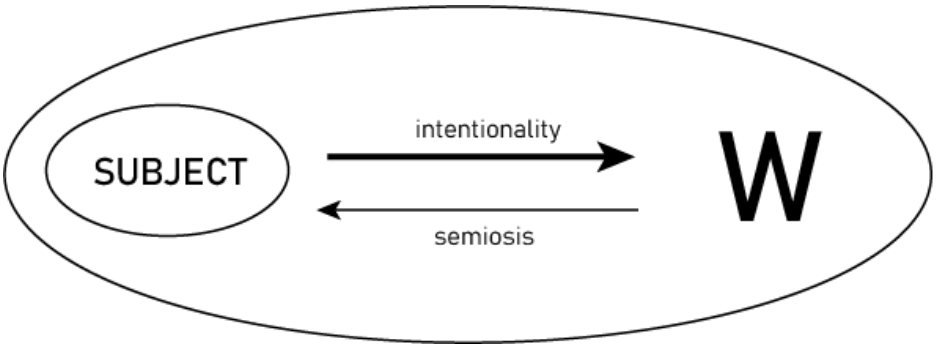

Figure 3: The reciprocal, though asymmetrical, relation between the subject and the world, and the dual nature of this relation as consciousness/intentionality and semiosis; adapted from Zlatev and Konderak (2023).

Layer and label	*Intentionality*	*Structures of semiosis*
5. Language	Linguistic intentionality	Predication, syntax
4. Sign use	Signitive intentionality	Signs
3. Intersubjectivity	Shared intentionality	Mimetic schemas, conventions, rituals
2. Subjectivity	Perceptual intentionality	Noemata, Leib/Körper duality
1. Animation	Operative intentionality, drive intentionality	Schemas, goals, emotions

Figure 4: The five levels of interlinked intentionality/semiosis in the Semiotic Hierarchy, adapted from Zlatev and Konderak (2023).

At the lowest level lies *drive* or *operative intentionality,* also referred to as "animation", since it is closely linked with the experience of movement and life itself, serving as "a non-representational background of subjectivity" (Zlatev & Konderak 2023: 179). Meaning on this level already involves consciousness, albeit not focal but marginal with respect to its object (*noema*) and pre-reflective as process (*noesis*) (Mouratidou 2025). This includes the constant presence of affective tone, whereby all acts of movement and perception receive inherent value (Thompson 2007). In the cognitive science literature, meaning on this level is commonly described in terms of *body schemas* (Gallagher 2005).

The next level is termed "subjectivity", even if this is somewhat misleading, as subjective experience is already present on the lowest level. The point is that what dominates here is *noematic* (object-directed) intentionality,

which "is the kind of intentionality that is clearly directed to an intentional object, in perception, memory, or imagination" (Zlatev & Konderak 2023: 180). There is an inherently *categorial* dimension to such intentionality, as opposed to operative intentionality, as noemata have parts/attributes and wholes, allowing for a basic form of pre-verbal thinking: "We now register the whole as containing the part. A relation between whole and part is articulated and registered" (Sokolowski 2000: 90). Another essential structure of intentionality/semiosis on this level is that of the dual experience of our bodies as sensing and experiencing subjects (*Leib*), on the one hand, and as possible objects of perception (*Körper*) on the other. With my body as *Leib,* I experience and act in the world, but at the same time, I can recognise my body as an entity in the world, an intentional object for someone else, *Körper*. I may experience joy or sadness, which my body expresses in a certain way. This duality serves to ground intersubjectivity through empathy, as we are able to experience the lived bodies of others in the same way as we do our own: as combinations of *Leib* and *Körper*. We may not feel exactly the intensity of what the other feels, but that is precisely what distinguishes them as "others" (Zahavi 2019).

This serves as a foundation for the next level of *intersubjectivity*, understood here as the sharing of perception-action schemas, goals, and even conventions, but not yet signs (see below). Donald (2001) and Zlatev (2008) have argued that in childhood and human evolution intersubjectivity develops through *bodily mimesis*: the volitional use of the body to help understand the world, others and oneself. Some acts of mimesis can also be rehearsed in self-learning and imitated in pedagogy. When such acts concern practical actions, mimesis functions as a process of *typification,* forming socially agreed-upon categories (Sonesson 2026).

Bodily mimesis is also capable of bridging the gap to the next level: sign use (Sonesson 2007; Zlatev 2013) and its particular kind of consciousness: *signitive intentionality,* and its very particular form of semiosis:

> The sign as such is a whole made up of two parts, expression and content, and there is a *double asymmetric relationship* between them. First, from the point of view of immediacy, expression is more accessible to consciousness than content. In the second place, content is more in focus (more prominent, more important) than expression. (Sonesson 2012: 85)

The crucial point is that in the case of signitive intentionality, the intentional object is intended indirectly through the mediation of signs, which

can be of several kinds: *indicational* (e.g., guiding posts in the forest to indicate pathways), *pictorial* (e.g., a sculpture) and *symbolic* (e.g., a logical argument). These categories correspond to the semiotic grounds of *indexicality* (contiguity), *iconicity* (resemblance) and *symbolicity* (conventionality), respectively (Sonesson 2012; Zlatev 2018). Figure 5 illustrates this integrative conception of the sign.

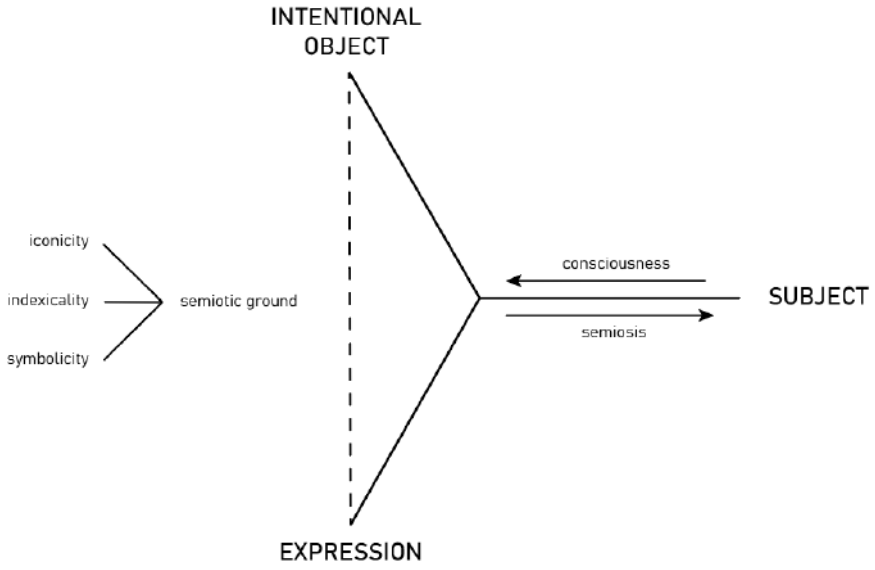

Figure 5: A representation of the sign concept, with its dependence on the consciousness/intentionality of the subject.

We now move into the fifth and final level of *language* and *linguistic intentionality.* Importantly, this is not to be defined by the purely conventional relation between expression and content, nor by its "arbitrariness" (Saussure [1916] 1960), as both words and composite linguistic constructions are strongly motivated by pre-linguistic experience (see Section 3.4). What distinguishes this level from pre-linguistic sign use is that linguistic symbols can be combined in finely articulated structures, forming constructions such as predication, modification, conjunction and subordination. These structures can in turn "be used to represent or construe complex situations, and express corresponding thoughts" (Zlatev & Konderak 2023: 185). As a semiotic system used for communication and thinking, language emerges from spontaneous acts of speaking, in the dialectic of spontaneity and sedimentation.

An important implication of the Semiotic Hierarchy is that linguistic semiosis presupposes and builds upon all preceding levels of meaning. While it is distinct from meaning making on the lower levels, it would be impossible without them. This is one of the reasons why language use is meaningful for us, but not for machines.

3.3. Language and expression

As pointed out above, language rests on pre-verbal semiotic capacities: "The reason we can communicate […] is because we have the ability to go from perception to categorial thinking. It is not the case that we can think, because we have language; rather, we have language because we can think, because we have the ability to achieve categorial intentions" (Sokolowski 2000: 91).

At the same time, language is not a mere externalisation of pre-verbal thought, since it is through language that a specific type of thinking emerges: "speech […] does not translate ready-made thought, but accomplishes it" (Merleau-Ponty 1962: 207). When we use language to refer to everyday objects, feelings and actions, these are in part constituted by language and the "available meanings" that inhabit the words. Such meanings of sedimented language arise through shared and negotiated use in our social interactions. Thus, words are not experienced as labels associated with stimuli, but as meanings readily available in a pre-reflective manner: "The word and speech must somehow cease to be a way of designating things or thoughts, and become the presence of that thought" (Merleau-Ponty 1962: 211).

In his phenomenology of language, Merleau-Ponty discerns and defines "language" as an interplay between *langage parlé* and *langage parlant*. The former is "already spoken language" in the sense of "language after the fact, language as an institution" (Merleau-Ponty [1969] 1973: 10), comprising the sedimented structures that outlive individuals. By contrast, *langage parlant* is the living, expressive activity of speaking, from which the sedimented structures initially emerge. Through processes of sedimentation, *langage parlé* gradually sediments into an articulated symbolic system that belongs to an "anonymous collectivity" (Zlatev 2023: 42). This system is thus characterised by formal and impersonal structures that possess precision and systematicity without implying that it is static: *langage parlé* emerges from "former acts of expression" and remains in a dynamic "constant state of becoming and change" (Zlatev 2023: 52).

Langage parlant resembles the modern conception of *languaging* (e.g., Kendon 2014) and is characterised by inherently meaningful and spontaneous usage. This is the phenomenon of *expression*, using the term here not

in contrast to content or intentional object, as in the previous sub-section, but as a spontaneous act in which meaning is being realised. Although expression is not restricted to language, it is defining for *langage parlant*. As Merleau-Ponty (1962) argues, whether expression takes the form of music, painting or language, its essence lies in bringing forth new ideas. This transformation of previous experience, without erasing it, constitutes the "fresh usage" that *langage parlant* entails. However, without the sedimented structures of *langage parlé*, there would be no resources to draw from in order to say something new.

The spontaneous expression of language, thus, arises from bodily experience in combination with language as a system. To summarise this dialectical relationship: *langage parlé* provides the resources for saying something original, while *langage parlant* is the living and expressive use of those resources. Over time, this sediments into grammatical structures and available meanings. Rather than viewing one as taking primacy over the other, phenomenology shows them to be simultaneous and intertwined, both drawing on and feeding into each other. These insights are reflected in the model described in the following sub-section.

3.4. The Motivation & Sedimentation Model (MSM)

Similar to the Semiotic Hierarchy, the Motivation & Sedimentation Model (MSM) involves the interaction between different levels of meaning. But here, these are understood even more dynamically, since the model focuses on the interplay between bodily experience, fluid social interaction and sedimented cultural formations. It was inspired not only by the phenomenology of language, reviewed in the previous sub-section, but also by the *integral linguistics* of Eugenio Coseriu (1985), which likewise operates with three analogous levels, but limited to language. Much of the empirical work on MSM has focused on metaphors (Stampoulidis, Bolognesi, and Zlatev 2019; Moskaluk, Zlatev, and van de Weijer 2022). However, it has also been extended to the analysis of other cultural phenomena, including the origins of money (Oakley & Zlatev 2024) and poetic creation (Mouratidou & Blomberg under review).

In a way, MSM (see Figure 6) can be seen as an elaboration of the dialectic of spontaneity and sedimentation emphasised by Merleau-Ponty (1962; see also Zlatev 2023), via the concepts of *motivation* and its correlate, *sedimentation*.

The most fundamental is the *Embodied* level, corresponding to the intentionalities at the lowest three levels of the Semiotic Hierarchy: "The Embodied level of meaning is both pre-linguistic and pre-signitive (*i.e.*, not based

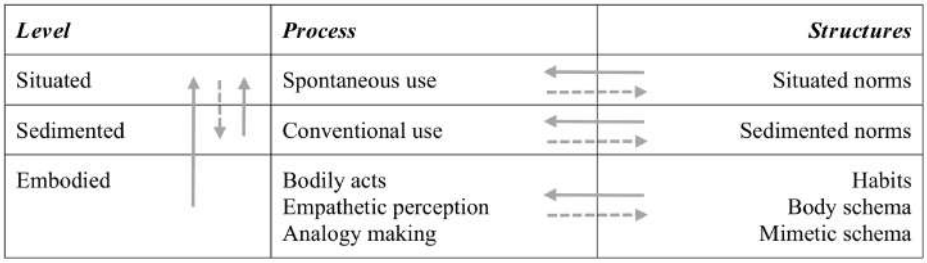

Figure 6: The Motivation & Sedimentation Model: the motivation relation is represented by the solid lines, and the sedimentation relation by dotted lines, on short-term (horizontal) and historical (vertical) time scales. Adapted from Zlatev (2023).

on sign use), engulfed in the richness of 'the visible': processes of Gestalt perception and motility, bodily interactions with things, empathetic direct perception of other persons" (Zlatev 2023: 55). Thus, the Embodied level excludes sign use as well as language but includes capacities such as enactive perception, bodily movement and interaction with the environment, as well as empathetic dealings with others. Horizontally, it sediments into bodily structures like habits and body schemas, while at the same time being shaped by them. Processes such as categorial intentionality (see Section 3.2) also belong to this level.

The *Sedimented* level corresponds to Merleau-Ponty's notion of *langage parlé*: language as a system, distilled from former acts of use. However, it can be generalised also to all kinds of semiotic conventions, including phenomena that have been theorised in semiotics as "lexica" and "encyclopedia": shared cultural knowledge (Eco 1976). The common feature is that the intentionality of all the structures on this level is latent, or "frozen": they have meaning potential, but this potential can only become actual when adopted by human beings in actual social circumstances.

The *Situated* level, by contrast, includes the living and spontaneous use of language and other forms of sign use in social contexts, thus corresponding to *langage parlant*. This use both draws from and sediments into structures of *situated norms*: "locally emerging patterns of coordination and (implicit) agreement that emerge in particular social interactions" (2023: 57). Such norms include, for instance, different manners of speaking with a friend or an acquaintance, as well as sensitivity to social appropriateness and situational congruence. But once again, the intentionality here is only potential and becomes realised only when "unlocked" in actual, and unique, situations. Such *situated intentionality* is what allows words and other signs to

have fully fledged "pragmatic" meaning and reference, hooking up to referents, both existing and imagined.

Crucially, the three levels are linked by relations of motivation and sedimentation. The key concept of *motivation* implies a non-deterministic type of influence, a resource used by an embodied subject to constrain its activity. This takes place on each of the three levels (the "Process" column), but most fundamentally on that of the Situated level. Crucially, the Embodied level does not directly motivate the Sedimented level but the situated use of language, which is also motivated by the Sedimented level. The acts of expression on the Situated level sediment into conventionalised structures on the Sedimented level, thus being indirectly grounded in the Embodied level. This is how the model can both help explicate the phenomenological notion of *expression* presented in the previous sub-section and substantiate Searle's (1980) claim that the intentionality of language use presupposes embodied intentionality, and consciousness in general (Section 2.1). In the following section, we use these insights from phenomenology and cognitive semiotics to show why LLMs lack the capacity for meaning making.

4. Why LLMs lack semiosis

Building on what was presented in the previous two sections, we focus on three preconditions for language use to possess meaning/intentionality: embodiment, consciousness and situatedness. In each case, we show that LLMs lack this prerequisite, and thus, why their manipulations of linguistic tokens are inherently devoid of meaning.

4.1. Embodiment

As made explicit in the Semiotic Hierarchy (Section 3.2), all forms of semiosis (meaning making) are ultimately grounded in our bodily engagement with the world, with the lowest level of operative intentionality providing the foundations. Thus, the experiencing and expressive body constitutes the basis upon which signitive and linguistic intentionality rest. This insight is concisely expressed by John O'Neill in his introduction to Merleau-Ponty's *The Prose of the World*: "It is through our body that we can speak of the world, because the world in turn speaks to us through the body" (1973: xxxiii).

In contrast, the "bodies" of the language machines reviewed in Section 2 are completely lacking, as software "agents" like ChatGPT and their LLMs can only generate text based on probabilistic computations over their many nodes and layers, trained on trillion-word corpora of linguistic tokens, without any access to the world at all. In genuine language use, linguistic signs

are combined by embodied subjects (speakers, writers or signers) into articulated structures, with the help of grammatical elements, to form constructions such as predication, modification, conjunction and subordination.

Although LLM-based systems appear to exhibit these characteristics by producing grammatically coherent outputs, their "languaging" is not grounded in any kind of bodily experience, including its categorial structures. As explained in Section 2.3, LLMs rely on artificial neural networks that operate as "stochastic parrots", in the words of Bender et al. (2021). The kind of "meaning" attributed to LLMs is based on "distributed representations", where semantic information is claimed to be encoded through probabilistic proximity within "vector spaces". However, such "representations" do not in fact represent anything outside of themselves. They form a closed system that is void of any real reference to the world and its intentional objects and thus lack any kind of intentionality.

Some may argue that machines that do have "bodies", in particular robots, may exhibit a form of "animation" insofar as they can move, interact and manipulate objects, and "perceive" with the help of machine vision algorithms. However, just as with their use of language, these capacities arise solely through programming and "machine learning", without any kind of meaningfulness (value) for the system itself (Zlatev 2003). Thus, in contrast to living organisms, machines lack any genuine connection to the life world (Ziemke 2025).

4.2. Consciousness

The use of any kind of signs, including those of language, requires consciousness: first as the precondition for having any kind of experience, and second as *reflective* consciousness, which enables us to denote these experiences. We do not rely on "vector spaces" or mathematical sequence prediction to produce or understand language. Instead, language use depends on the relation between expressions and the intentional objects, mediated by signitive intentionality (see Figure 5). In contrast, LLM-based systems do not *know* anything, either about language or about the world that they are supposed to be about. Human-created training data forms the foundation for generating their productions, and the system is further fine-tuned through "reinforcement learning" to align with human preferences.

A key form of human (and animal) consciousness involves *empathy*, based on (partially) shared *Leibkörper* experiences (see Section 3.2). As shown in the following example (Figure 7), ChatGPT may seem to display

empathetic language use. But to take this as an expression of genuine empathy is not different from the gullibility that many displayed when interacting with *Eliza* in the 1960s (Section 2.1).

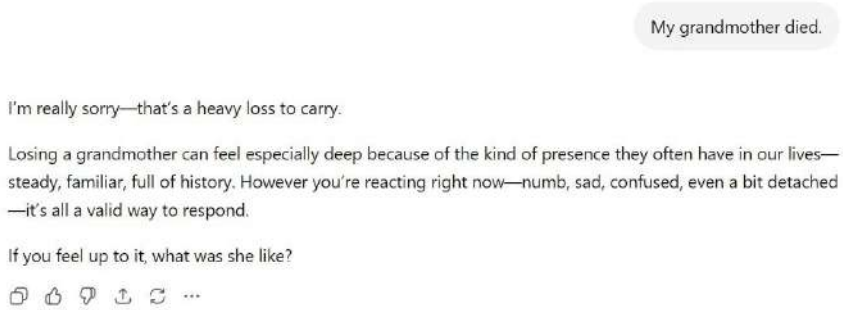

Figure 7: ChatGPT mimicking "empathy". Generated on 2 May 2026.

A potential objection could be that we also learn how to use language appropriately through social interaction by imitating others and responding to feedback. However, while human learning involves bodily mimesis (Section 3.2), LLMs operate through sequence prediction and adhere to feedback in a binary (i.e., positive or negative) manner. Admittedly, reflective consciousness and empathy are complex phenomena. However, both "occur in social contexts where joint action is possible, and social practices predominate" (Gallagher 2021: 336).

In most cases, our language use is deliberate and thus based on reflective consciousness. But even in cases where our speech diverges from our intentions, we typically apologise, acknowledging in this way that we are the originators of the act. This interconnection is reflected in Pacherie's (2007) threefold distinction of intentions (in the psychological sense of the term as goal-directedness): (a) those formed prior to actions (i.e., distal), such as planning to write up a list in order to go to the grocery store; (b) present or proximate intentions (equivalent to Searle's (1983) "intention-in-action"), which accompany actions; and (c) motor intentions, which carry out our actions by transforming the contents of previous intentions into bodily acts, such as picking up the pen to write the grocery list or getting into the car to get to the store. Even if we are not conscious of the underlying neural or muscular processes that facilitate (c), we are always more or less aware of what we are doing in terms of (a) and (b). By contrast, LLMs lack the basis for having any of the above.

More specifically related to language, LLM-based systems may draw on "former acts of expression" (i.e., training data) as resources (*langage parlé*)

for generating text, producing utterances such as those in Figure 7. But this would only correspond to the Sedimented level of MSM, and not to the Situated level of genuine expression (*langage parlant*) (see Figure 6). The productions of LLMs do not convey speakers' intentions, nor do they appropriate meaning as their own. Instead, they recombine sedimented expressions in response to external prompts, without any bodily experience that needs to underlie genuine acts of expression, in the sense of "fresh usage" described in Section 3.4. In sum, LLMs lack (a) any kind of subjective experience, (b) reflective consciousness to be able to have signs differentiated from their intentional objects and (c) their own goal-directedness: features corresponding to the Embodied and the Situated levels of MSM.

What makes LLM-based systems appear impressive to users is that they have access to trillion-word texts produced by millions of people, whereas human beings rely on their individual experiences and memories in order to express themselves. LLMs thus exceed the amount of information that is available to any individual person. Though the generated text may convey ideas that are new to the user who engages with the system, it will not be new compared to the vast material accessible to the LLM. Thus, such systems cannot bring forth new ideas that go beyond the training material, as there is nothing that would motivate this (e.g., embodied experiences/sensations or imagination). LLMs cannot give rise to new meanings and subsequently contribute to genuine acts of expression. They can only combine bits and pieces of previous expression in accordance with the algorithm. In essence, without consciousness, LLMs rely solely on stochastic manipulations of "symbols" – but only in scare quotes, as there is no meaning of any kind behind these.

4.3. Situatedness

Intentionality, both in the philosophical and psychological senses, implies a public, intersubjective context: a *situation*. Language is essential in how we assume ownership of and responsibility for who we are, both for ourselves and others, and ultimately, for the world. We do not only state facts; rather, we express our judgements and claims, positioning ourselves in relation to others and becoming what Sokolowski calls "agents of truth", since "[t]he grammatical articulation of words in the sentence corresponds to an articulated display of something in the world" (Sokolowski 2008: 32). In contrast, the language generated by LLM-based systems cannot be situated in such a sense, since it lacks the previous two features discussed: embodiment and consciousness.

Truly situated intentionality requires being an actual participant in a social situation. Human communication, as inherently both *polysemiotic* and *multimodal* (Zlatev et al. 2023), is an embodied, dynamic, and intersubjective activity in which meaning is "negotiated" between speakers and listeners, rather than transmitted as a self-contained entity.[4] The boundaries between speaker and listener are fluid and continuously negotiated. Meaning does not reside in linguistic signs alone but emerges through embodied participation in a common experiential world, where the acts of speaking and listening are aspects of a shared meaning-making process:

> Through language I discover myself and others, in talking, listening, reading, and writing. […] Our thoughts and purposes are embodied in bodily gestures which in the act of expression structure themselves toward habit and spontaneity, and thus we make our world (O'Neill 1973: xxxiv).

One might argue that some form of situatedness occurs in LLMs due to "in-context learning" (Section 2.3). This allows for a level of context sensitivity, as demonstrated by ChatGPT's ability to interact by mimicking previous examples. It can, for instance, imitate writing styles or align its responses with perspectives shared by the user. However, this kind of "context awareness" is confined to the current input and does not come from any organic participation in a social encounter. Rather, it only "parrots" the previous expressions (and views) of the user in a rather superficial manner. Thus, it remains fundamentally un-situated.

As pointed out above, while LLM-based NLP systems may reproduce sedimented linguistic norms, they ultimately lack access to *situated norms* (see Figure 6), which require negotiations between different, in principle equally valid, perspectives. But LLMs lack precisely any kind of *perspective*, as they lack embodiment and consciousness, as established above. Whatever "knowledge" of context they may have is restricted to patterns of language extracted from the training data. By contrast, situated norms, such as social appropriateness and situational congruence, influenced by interpersonal relations, depend on intersubjective engagement with the world, including

---

[4] *Semiotic systems* are defined as signs and signals of a particular type and their interrelations. The three major and universal human semiotic systems are language, gesture and depiction. *Polysemiotic* communication implies multiple semiotic systems working together, as when speaking and gesturing at the same time. In a cognitive semiotic context, the term *multimodality* is used for the combination of multiple perceptual *modalities*, e.g., visual, auditory, tactile, etc. Polysemiotic communication can be multimodal, or unimodal, as when pictures and written text are combined in a cartoon (Zlatev et al. 2023).

sensitivity to factors such as relationships with their dynamics, the occurrence of actual events and bodily cues (e.g., gesture, prosody, and gaze), all of which are inaccessible to LLMs.

5. Summary and conclusions

Today's "language machines" have come a long way since the 1950s, as reviewed in Section 2. Their contributions to *weak AI*, technology that resembles how human beings perform on tasks that we spontaneously judge to be intelligent, are undeniable. However, there is no good reason to suppose that they do, or can, instantiate *strong AI:* understanding the texts that are being presented to them, or what they themselves produce. And the fundamental issues persist, since they lack *intentionality*. Searle's (1980) thought experiment about the "Chinese Room" remains relevant for current LLMs, because they continue to operate only on empty "symbols" and not on meaning.

In Section 3, we introduced relevant concepts and models from cognitive semiotics, which jointly helped explicate three features that are necessary for intentionality: embodiment, consciousness and situatedness. In Section 4, we argued that LLMs demonstrably lack all three of these. First, they lack experiential bodies, *Leib*, that place them in contact with the sensory and intersubjective world. Second, they lack reflective consciousness, necessary to grasp the representational relation between expressions and the intentional objects that they denote, as well as empathy and intentions. And third, they lack the ability to be situated in a genuine interpersonal context, as opposed to imitations of one. In terms of the Motivation & Sedimentation Model, since LLMs strictly operate on the Sedimented level, with no access to either the Embodied or Situated levels of perceptual and social experience, it is impossible for them to have either embodied or situated intentionality. Even though their machine-produced texts may mimic human language use impressively well, there is absolutely no meaning behind this.

Why then do some people think that there is? As Bender et al. (2021: 611) argue, the reason is that we tend to "impute meaning where there is none". This tendency stems from our own deeply ingrained intersubjectivity and signitive intentionality. These compel us to participate in, and project meaning onto, what resembles linguistic interaction, even if this is essentially an echo chamber, and there is no one else there but ourselves and a good enough approximation of language, based on a distillation from trillions of "words" out there in cyberspace.

In sum, this paper conceptually examined the language capabilities of the currently influential LLM-based systems, illustrated through their most famous representative, ChatGPT. Our argument has led from the premises (1) creativity (in language, but also more generally) presupposes understanding, and hence intentionality, and (2) LLMs lack intentionality, to the conclusion (3) LLMs cannot be truly creative. While we have given some indirect support for (1), with observations that LLM-based systems engage only with sedimented language, and therefore there is no reason to expect that they can be creative, we have for the most part assumed this premise, as arguing for it would require us to first define creativity, and this is beyond the scope of the present paper.

Having said this, we do not deny that LLMs can be used as *extensions* of human intentionality, and thereby of human creativity. But this is not qualitatively different from how various kinds of cognitive artefacts can "extend the mind" (Clark & Chalmers 1998). However, as argued by Sonesson (2026), any such extensions need to start from the "naked", unextended mind, with its embodiment, situatedness and reflective consciousness. While LLM-based systems like ChatGPT may be a valuable tool, they also pose societal, and even existential, risks insofar as they contribute to the devaluation of human (and animal) meaning and intentionality. In this paper, we focused on the latter, since we are concerned that if overestimated, AI will lead to the underestimation of genuine human meaning making and creativity. For, despite everything, ChatGPT is nothing but a much fancier version of *Eliza*: only an imitation of the real thing.

References

Belcic, I., C. Stryker. 2026. What Is Agentic RAG? *IBM Think*. Available at: https://www.ibm.com/think/topics/agentic-rag (accessed 1 May 2026).

Bender, E. M., T. Gebru, A. McMillan-Major and S. Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. New York: Association for Computing Machinery, 610–623. Available at: https://doi.org/10.1145/3442188.3445922.

Clark, A., D. Chalmers. 1998. "The Extended Mind". *Analysis*, Vol. 58, No. 1, 7–19. Available at: https://doi.org/10.1093/analys/58.1.7.

Clark, H. H. 1996. *Using Language*. Cambridge: Cambridge University Press.

Coseriu, E. 1985. "Linguistic Competence: What Is It Really?". *The Modern Language Review*, Vol. 80, No. 4, xxv–xxxv.

Dhingra, H., P. Jayashanker, S. Moghe and E. Strubell. 2023. "Queer People Are People First: Deconstructing Sexual Identity Stereotypes in Large Language Models". *ArXiv*, arXiv:2307.00101. Available at: https://doi.org/10.48550/arXiv.2307.00101.

Donald, M. 2001. *A Mind So Rare: The Evolution of Human Consciousness*. New York: W. W. Norton.

Dreyfus, H. L. 1992. *What Computers Still Can't Do: A Critique of Artificial Reason*. Cambridge, MA: MIT Press.

Eco, U. 1976. *A Theory of Semiotics*. Bloomington: Indiana University Press.

Gallagher, S. 2005. *How the Body Shapes the Mind*. Oxford: Oxford University Press.

Gallagher, S. 2021. Phenomenology of Agency and the Cognitive Sciences. In Erhard, C., T. Keiling (eds.). *The Routledge Handbook of Phenomenology of Agency*. London and New York: Routledge, 334–348.

Inie, N., P. Zukerman and E. M. Bender. 2026. "De-anthropomorphizing 'AI': From Wishful Mnemonics to Accurate Nomenclature". *First Monday*, Vol. 31, No. 2. Available at: https://doi.org/10.5210/fm.v31i2.14366.

Katz, D. M., M. J. Bommarito, S. Gao and P. Arredondo. 2024. "GPT-4 Passes the Bar Exam". *Philosophical Transactions of the Royal Society A*, Vol. 382, No. 2270, Article 20230254. Available at: https://doi.org/10.1098/rsta.2023.0254.

Kendon, A. 2014. "Semiotic Diversity in Utterance Production and the Concept of 'Language'". *Philosophical Transactions of the Royal Society B: Biological Sciences*, Vol. 369, No. 1651, Article 20130293. Available at: https://doi.org/10.1098/rstb.2013.0293.

Kernighan, B. W. 2021. *Understanding the Digital World: What You Need to Know about Computers, the Internet, Privacy, and Security*. (2nd ed.). Princeton: Princeton University Press.

Konderak, P. 2018. *Mind, Cognition, Semiosis: Ways to Cognitive Semiotics*. Lublin: Maria Curie-Skłodowska University Press.

Kosinski, M. 2023. "Evaluating Large Language Models in Theory of Mind Tasks". *ArXiv*, arXiv:2302.02083 (version 6). Available at: https://doi.org/10.48550/arXiv.2302.02083.

Krohn, J., G. Beyleveld and A. Bassens. 2020. *Deep Learning Illustrated: A Visual, Interactive Guide to Artificial Intelligence*. Boston: Pearson Education.

Lardinois, F. 2023. Microsoft Launches the New Bing, with ChatGPT Built In. *TechCrunch*, 7 February 2023. Available at: https://tcrn.ch/3I7GZI4.

Malik, A. 2023. Snapchat Launches an AI Chatbot Powered by OpenAI's GPT Technology. *TechCrunch*, 27 February 2023. Available at: https://tcrn.ch/3J72wBz.

Mendoza-Collazos, J. C., G. Sonesson. 2021. "Revisiting the Life of Things: A Cognitive Semiotic Study of the Agency of Artefacts in Amazonia". *Public Journal of Semiotics*, Vol. 9, No. 2, 30–52.

Merleau-Ponty, M. 1962 [1945]. *Phenomenology of Perception*. (transl. Colin Smith). London: Routledge.

Merleau-Ponty, M. 1968 [1964]. *The Visible and the Invisible*. (transl. Alphonso Lingis). Evanston: Northwestern University Press.

Merleau-Ponty, M. 1973 [1969]. *The Prose of the World*. (transl. John O'Neill). Evanston: Northwestern University Press.

Moskaluk, K., J. Zlatev and J. van de Weijer. 2022. "'Dizziness of Freedom': Anxiety Disorders and Metaphorical Meaning-making". *Metaphor and Symbol*, Vol. 37, No. 4, 303–322. Available at: https://doi.org/10.1080/10926488.2021.2006045.

Mouratidou, A. 2025. *Mind the Body: Verbal and Non-verbal Signs of Choice Awareness*. PhD thesis. Lund: Lund University.

Mouratidou, A., J. Blomberg. Under review. "The Voice of Mute Experience: A Phenomenological Account of Poetic Creation". *Phenomenology and the Cognitive Sciences*.

Oakley, T., J. Zlatev. 2024. "Origins of Money: A Motivation & Sedimentation Model (MSM) Analysis". *Semiotica*, Vol. 2024, No. 257, 1–27. Available at: https://doi.org/10.1515/sem-2023-0031.

O'Neill, J. (1974). Translator's introduction: Language and the voice of philosophy. In M. Merleau-Ponty, *The prose of the world* (J. O'Neill, Trans., pp. ix–xxxv). Northwestern University Press.

OpenAI. 2026. *GPT-5.3 and GPT-5.5 in ChatGPT*. OpenAI Help Center. Available at: https://help.openai.com/en/articles/11909943-gpt-53-and-gpt-55-in-chatgpt (accessed 1 May 2026).

Pacherie, E. 2007. "The Sense of Control and the Sense of Agency". *Psyche*, Vol. 13, No. 1, 1–30.

Ravi, D., R. Sindhgatta. 2025. Exploring Trust and Transparency in Retrieval-Augmented Generation for Domain Experts. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. New York: Association for Computing Machinery, 1–7.

Russell, S., P. Norvig. 2021. *Artificial Intelligence: A Modern Approach*. (4th ed.). Hoboken: Pearson.

Saussure, F. de. 1960 [1916]. *Course in General Linguistics*. (transl. Wade Baskin). London: Peter Owen.

Schreiner, M. 2023. GPT-4 Architecture, Datasets, Costs and More Leaked. *The Decoder*, 11 July 2023. Available at: https://the-decoder.com/gpt-4-architecture-datasets-costs-and-more-leaked/.

Searle, J. R. 1980. "Minds, Brains, and Programs". *Behavioral and Brain Sciences*, Vol. 3, No. 3, 417–424. Available at: https://doi.org/10.1017/S0140525X00005756.

Searle, J. R. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.

Shanahan, M. 2024. "Talking about Large Language Models". *Communications of the ACM*, Vol. 67, No. 2, 68–79. Available at: https://doi.org/10.1145/3624724.

Sokolowski, R. 2000. *Introduction to Phenomenology*. Cambridge: Cambridge University Press.

Sokolowski, R. 2008. *Phenomenology of the Human Person*. Cambridge: Cambridge University Press.

Sonesson, G. 2007. From the Meaning of Embodiment to the Embodiment of Meaning: A Study in Phenomenological Semiotics. In Ziemke, T., J. Zlatev and R. M. Frank (eds.). *Body, Language and Mind. Vol. 1: Embodiment*. Berlin: Mouton de Gruyter, 85–128.

Sonesson, G. 2012. Semiosis beyond Signs: On Two or Three Missing Links on the Way to Human Beings. In Schilhab, T., F. Stjernfelt and T. Deacon (eds.). *The Symbolic Species Evolved*. Dordrecht: Springer, 81–96.

Sonesson, G. 2026. *Semiotic Investigations: From Meaning to Pictures and Communication*. Berlin: De Gruyter Brill.

Stampoulidis, G., M. Bolognesi and J. Zlatev. 2019. "A Cognitive Semiotic Exploration of Metaphors in Greek Street Art". *Cognitive Semiotics*, Vol. 12, No. 1, Article 20192008. Available at: https://doi.org/10.1515/cogsem-2019-2008.

Sternberg, R. J. (ed.). 1999. *Handbook of Creativity*. Cambridge: Cambridge University Press.

Thompson, E. 2007. *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Cambridge, MA: Harvard University Press.

Turing, A. M. 1950. "Computing Machinery and Intelligence". *Mind*, Vol. 59, No. 236, 433–460.

Wu, T., S. He, J. Liu, S. Sun, K. Liu, Q.-L. Han and Y. Tang. 2023. "A Brief Overview of ChatGPT: The History, Status Quo and Potential Future

Development". *IEEE/CAA Journal of Automatica Sinica*, Vol. 10, No. 5, 1122–1136.

Zahavi, D. 2019. *Phenomenology: The Basics*. London: Routledge.

Ziemke, T. 2025. The Role of Intentionality in Human–Robot Interaction. In Johnson, E. (ed.). *How That Robot Made Me Feel*. Cambridge, MA: MIT Press, 37–62.

Zlatev, J. 2003. "Meaning = Life (+ Culture): An Outline of a Unified Biocultural Theory of Meaning". *Evolution of Communication*, Vol. 4, No. 2, 253–296.

Zlatev, J. 2008. The Co-evolution of Intersubjectivity and Bodily Mimesis. In Zlatev, J., T. P. Racine, C. Sinha and E. Itkonen (eds.). *The Shared Mind: Perspectives on Intersubjectivity*. Amsterdam: John Benjamins, 215–244.

Zlatev, J. 2009. "The Semiotic Hierarchy: Life, Consciousness, Signs and Language". *Cognitive Semiotics*, Vol. 4, 169–200.

Zlatev, J. 2013. "The Mimesis Hierarchy of Semiotic Development: Five Stages of Intersubjectivity in Children". *Public Journal of Semiotics*, Vol. 4, No. 2, 47–70.

Zlatev, J. 2015. Cognitive Semiotics. In Trifonas, P. P. (ed.). *International Handbook of Semiotics*. Dordrecht: Springer, 1043–1067.

Zlatev, J. 2018. "Meaning Making from Life to Language: The Semiotic Hierarchy and Phenomenology". *Cognitive Semiotics*, Vol. 11, No. 1, Article 20180001.

Zlatev, J. 2023. "The Intertwining of Bodily Experience and Language: The Continued Relevance of Merleau-Ponty". *Histoire Épistémologie Langage*, Vol. 45, No. 1, 41–63.

Zlatev, J., P. Konderak. 2023. Consciousness and Semiosis. In Pelkey, J. (ed.). *Bloomsbury Semiotics, Volume 1: History and Semiosis*. London: Bloomsbury Academic, 169–192.

Zlatev, J., S. Devylder, R. Defina, K. Moskaluk and L. B. Andersen. 2023. "Analyzing Polysemiosis: Language, Gesture, and Depiction in Two Cultural Practices with Sand Drawing". *Semiotica*, Vol. 2023, No. 253, 81–116.