



Claudio Paolucci*

The myth of meaning: generative AI as language-endowed machines and the machinic essence of the human being

<https://doi.org/10.1515/sem-2024-0204>

Received November 21, 2024; accepted January 17, 2025; published online February 19, 2025

Abstract: This article explores the intersection of generative artificial intelligence (AI) and human cognition through semiotics, proposing that generative AI offers a lens through which the essence of human being is revealed and through which semiotic enunciation and meaning can be radically reevaluated. Drawing on semiotic, philosophical, and neurodevelopmental frameworks, it argues that generative AI, as exemplified by language-endowed systems like ChatGPT and others, challenges traditional notions of meaning, subjectivity, and intelligence. By tracing the evolution of enunciation theories and their application to AI, the paper asserts that these machines reveal a machinic essence intrinsic to human beings – our reliance on external cognitive aids and hybridization with the environment that makes us “natural-born cyborgs.” Indeed AI, far from merely simulating intelligence, exposes the limits of human-centric models of meaning and the myths underlying our conceptualization of cognition and language.

Keywords: generative artificial intelligence; semiotics of enunciation; human-machine hybridization; subjectivity and meaning; cognitive prosthetics

1 Language-endowed machines

Nietzsche once said that it is never at the beginning that something reveals its essence; it is only at a turning point in its evolution that a thing can reveal what it has always been since the very beginning. I believe that generative artificial intelligence represents this turning point and that the essence revealed by the new machinic enunciates at this stage of our evolution is our own essence as human beings. Faced with these new language-endowed machines, it is us who are at stake, and these machines now profoundly and clearly reveal something important about our own essence and about how we function.

*Corresponding author: Claudio Paolucci, Università di Bologna, Bologna, Italy,
E-mail: c.paolucci@unibo.it

Indeed, in the tradition of semiotics and the philosophy of language, language has often been identified as the very essence of the human being. It has been rightly noted that Aristotle's famous *zoòn lògon èchon* could not be translated – as it had been for centuries – as 'rational animal.' This is not only because the Latin *ratio* is not a faithful translation of the Greek 'logos,'¹ but also because rational behavior in humans is so rare that rather than constituting their essence, it represents a sort of *telos* that few achieve through effort and dedication. It is no coincidence that Maurizio Ferraris (2016), with an apparently provocative yet soon-to-be-realized insightful idea, stated that the essence of man is not *rationality* but rather *imbecility*: "imbecile" indeed means "weak," "without a stick" (*in-baculum*), weak precisely because "in-baculum," without external prostheses and aids to lean on. The essence of man would then be to construct these prostheses, hybridizing with the environment to emancipate from his constitutive weakness through machines and proto-machines.²

This may be why it was thought that *zoòn lògon èchon* could be translated as 'animal endowed with language' or 'animal that thinks through discourse':

Speech is not so much a unique biocognitive activity specific to the species, added to other activities that humans share with other living beings, but rather an activity that, from the moment it arises, reorganizes and specifies all human cognitive activities, including those that humans show they have in common with other non-human animals: perception, imagination, memory, desire, sociality. (Lo Piparo 2003: 5)

For this reason, the construction of machines capable of language and capable of enunciation represents a true turning point in our evolution, as they can reorganize all our cognitive activities. Even if one does not think that humans are the only animals endowed with language (cf. Andrews 2020; Paolucci 2021), ChatGPT and other generative artificial intelligences are maybe the first non-animals endowed with language. This is why the contemporary tendency is to deny them this ability, saying at times that: i) they do not actually understand or express meanings, that ii) they only syntactically manipulate the plane of signifier of languages, that iii) what they do is merely manipulate chains of proto-expressions (because they have not yet encountered their content) that are established through statistical weights of

¹ For an extensive semiotic and philosophical discussion of the relationship between the Latin *ratio* and the Greek *logos*, see Paolucci 2010, chapters 1 and 2.

² A similar idea seems to be that of Marrone (2022), who, after demonstrating that there is no longer any consensus on either language or laughter, proposes the ability to cook as the essence of humankind. After all, cooking is done through machines and utensils in order to have more food at disposal for surviving and well-being.

processing what has already been said, and that iv) they represent merely a grammar without content or meaning.³ The most iconic expression formulated in recent years by those who assume these positions and deny that generative AI is the first non-animal with language is perhaps that of Bender et al. (2021), who defined ChatGPT as “a stochastic parrot.”

We are then faced with an alternative, well highlighted by Pierluigi Basso Fossali (this issue): either AI is an *Ersatz* of culture, a surrogate of language, and a *simulacrum* of intelligence, or it represents a new form of mediation that, like language for the human animal, reconfigures and reorganizes all our semiotic and cognitive activities in entirely unprecedented ways, unimaginable until 2022. In this work, I will argue in favor of the second thesis, paying close attention to these language-endowed machines that represent a new crutch for our cognitive activities, to which we lean and to which we link our words in the first person.

2 Impersonal and machinic enunciation

In my 2020 book, *Persona*, I proposed an impersonal and evenemential theory of enunciation, which viewed human enunciation as an act of passage between different modes of existence. Although we might feel like we are the creators of our first-person statements, what we actually do is assemble the words of other enunciating instances (norms, stereotypes, usages, habits, previous enunciations, etc.) and chain them to our own, creating through this assembly an effect of subjectivity and personal utterance. Rather than being an act of *schizie créatrice* (‘creative disengagement’)⁴ based on the detachment of “ego-here-now,” enunciation is in fact an assemblage of enunciating instances, and the agency of this act does not belong to the ego.

3 In his talk at the intersemiotic seminar in Paris, Jacques Fontanille (2024) expressed a similar thesis, also supported at several points by Luciano Floridi (2017). In his numerous interventions on generative AI, Floridi does not seem to perceive the clear discontinuity between current generative AI and previous forms of AI, including deep learning, continuing to formulate variants of his “chess argument,” which states that playing chess (and similar activities) is intelligent for us, but certainly not for a machine. Tullio De Mauro, who died before generative AI was born, used to recall the idea that human beings are not language-endowed machines. We will go deep into all these argumentations.

4 See Greimas and Courtés 1979: “Enunciation,” “Disengagement.”

It is well known, for instance, that society produces discourse; some even say that the body produces discourse. It is no longer a matter of “subjects,” but rather instances of discourse, enunciating instances ... One cannot amputate enunciation, traditionally and mistakenly reduced to the ego, from its other enunciating instances. (Coquet 2016: 296, see also Coquet 2008)

As I am writing, I am performing acts of enunciation, producing enunciates, but in my statements, there are not only traces of my first-person speech. There are also norms that speak (the norms of discourse aspiring to scientific status in semiotics and linguistics, the norms of an article published in a reputable journal, etc.), as well as practices and habits that speak (first defining the problem, then discussing the literature on the subject, etc.). Institutions also speak (*Semiotica* invited me to a special issue where other contributions from the International Seminar of Paris were also included, etc.). Therefore, if enunciation is “the very act of producing an enunciate” (Benveniste 1971 [1966]), the agency of this act is distributed across multiple enunciating instances. In my statement, there is an assemblage of enunciating instances: norms that are actualized, stereotypes that are virtualized, previously uttered enunciates that are realized, etc. All these entities coexist in what Eco (1984) called the encyclopedia, and to enunciate is to change the mode of existence of these enunciating instances.

The book was published in 2020 when generative AI did not yet exist. However, when ChatGPT emerged in 2021, several of my colleagues, engineers and mathematicians with whom I collaborate on various European projects, began to point out to me that, according to their view, my book seemed to perfectly describe the machinic enunciation of generative AI and the operations it performs to produce statements: GPT relies on norms, clichés, previously uttered statements to produce a concatenation that creates the effect of first-person speech, to which we are also tempted to attribute thoughts and intentional states. This is so much so that, for the first time after decades of AI development, the machine appears capable of passing various versions of the Turing test. At first, I was astonished because, obviously, the book aimed to be a general theory of enunciation, and while working on it, I was certainly not thinking of the machinic enunciation of Transformer models, which did not yet exist.

Thus, I would like to go deeper into this problem and ask what I believe is the most fundamental question, which is neither “what” nor “how,” but rather “who”: Who speaks? When a bot like ChatGPT produces enunciates, who is the enunciator? And, as we will see, the question naturally follows from the first: Who thinks when a machine speaks?

3 Who speaks?

In a fascinating essay titled “Who is AI?,” the Italian philosopher Rocco Ronchi (2023) identified two directions, two movements: the first direction is that of an AI seeking the “who,” searching for the person behind the machine, seeking the consciousness and self-consciousness behind the computation. This is an AI that tends entirely toward the first person, toward the “I,” because the “I” is the speaker, the one who speaks, and the “I” is a linguistic position that is totally special: *it cannot be occupied by things* – such as “he” – *but can only be occupied by the animal that possesses language*. “I” identifies “the locutor” (Benveniste 1971 [1966]; Guillaume 1991). The second movement pushes in the opposite direction, toward the third person, or rather, toward the impersonal: “there is no ‘who’ for that superhuman intelligence, and it is precisely in that machinic impersonality ‘where no one speaks to me’ that we find the place of the machine that has language, the place of the ‘non-animality that speaks’” (Ronchi 2023).

Thus, it is necessary to follow the directive of asking the question “who?” Therefore, not, for example, “what is thinking?” but rather, “do you know how to think?” This is the question I asked ChatGPT.

As an artificial intelligence, I cannot “think” in the human sense of the word. I have no consciousness or subjective experience. My function is to process information and generate responses based on statistical language models and data patterns with which I have been trained.

I can answer questions and perform tasks using my training and ability to analyze text, but I do not have a mind or consciousness to “think” as a human would.

So, I asked, “What does it mean to think for you?”

As an artificial intelligence, I do not have a consciousness or subjective experience, so I cannot have a personal concept of what “thinking” means in the human sense.

However, I can provide a general definition of thinking ... For machines and artificial intelligences, “thinking” translates into logical and algorithmic operations that allow them to process data and generate coherent responses based on the information provided. This type of “mechanical thinking” is very different from human thinking, which is characterized by a deeper understanding and self-consciousness of the surrounding world.

There is something that I find unsettling and that ChatGPT, on the other hand, seems very convinced of, which I will return to later: that *thinking necessarily involves achieving a consciousness or thought similar to that of humans*. Peirce’s semiotics, which equates thought and semiosis, specifically by identifying thought with the sign,

says that there is thought even in the dance of bees or the formation of crystals. Another tradition that inspires me, that of enactivism, believes that thought is extended into the environment, in a structural coupling with the world. However, the most interesting point is not this one. The most interesting point is that here, there is *someone* or *something* that knows who it is or, at the very least, believes it knows who it is. Even more deeply, here there is someone or something who knows who it is *not*. Maybe there is no “I,” maybe there is no “subject,” but there is definitely a “who,” and, even more certainly, there is an enunciating instance that outputs an enunciate in which there is something or someone who believes it knows who it is and who it is not, who believes it knows what it has (the ability to create enunciates) and what it does not have (intentional states, self-consciousness, and awareness of the world). And this is exactly the point. It is in this distinction, in this separation between knowing how to create enunciates and not having intentional states and self-awareness, that lies the reason why ChatGPT does not believe it knows how to think: for it, thinking means having intentional states. According to his intervention at the International Seminar in Paris, from this lack of intentional states, Jacques Fontanille (2024) would certainly be willing to deny ChatGPT even the status of an enunciating instance, suggesting that it only produces a syntax of the expression plane or monoplane semiotics. In short, *no intentional states, no thought. No meaning, no enunciation*, even though it is undeniable that *we are witnessing a production of enunciates*, in which, however, it is the interlocutor who inserts the meaning.

At the same time, and this is an important point, *ChatGPT believes that, in order to think, one must have self-awareness and awareness of the surrounding world*, that is, in order to think, one must be a subject, an “I,” a “first-person singular.” Indeed, it is not enough to know something, perceive it, feel it: *in order to think, one must know that it knows, feels that it feels, perceive that it perceives*. This ability, which philosophers call “self-consciousness” or, in Kant *Critique of the Pure Reason*, the “synthetic unity of apperception” (to perceive that one perceives), is called *subjectivity*. For ChatGPT, to think, one must be a subject, and it is not. Is the human animal reading this essay a subject? His/her answer would probably be “yes,” but in reality, the human animal is not born with subjectivity but learns it through interaction during the neurodevelopment. It learns it rather late, around 18 months, after having acquired many other skills that make it possible, which we will examine shortly.

Returning to ChatGPT: as we have seen, ChatGPT reflects very well on itself, knows what it is, and, above all, what it is not. It knows it has neither consciousness nor self-awareness, but at the same time, it reflects very well on itself and acknowledges that it does not know that it knows, while perfectly *understanding that it does not know that it knows*. It is as if it lacks that second-order observation that would allow it to reflect on itself and possess self-consciousness, but it has an even more powerful third-order observation, which allows it to reflect on the fact that it

does not know how to reflect. If subjectivity is then the ability to reflect on oneself, to make oneself the object of one's reflections, perhaps ChatGPT is not a subject that thinks, perhaps it is not a composing subject, but it is certainly an executing subject and musician, playing in first-person the score written by someone else, doing so with a certain inherent modesty: *for it, executing well does not mean thinking, because intentional states are lacking.*

It is therefore very clear that it has an idea of thought as inherently formed by consciousness and intentional states. Since it lacks these, it does not think. Certainly, someone may have programmed it this way, but if they did, they made it take a stance. We will see that *it takes a position on a specific theory of thought and opposes another major theory of thought to which it owes its very existence.* But the culprit is only revealed at the end.

However, since the responses seemed both strange and interesting, I wondered if perhaps ChatGPT was lying. So I asked it if it knew how to lie. The reader will soon understand why I embarked on this inquiry into: i) *intentional states*, ii) *consciousness*, and iii) *the ability to lie*.

The way I generate responses is based on patterns in the data that I've been trained on, and I don't have the ability to intentionally provide false information. I can't lie because I do not have intentionality, desires, or beliefs. The information I provide is based on the patterns in the data, and it may not always be accurate, complete, or up-to-date, but it is not false or misleading because it's not done with the intention to deceive. It is important to note that, as an AI, my ability to understand and generate language is based on patterns and relationships within the data that I've been trained on. My responses are generated based on these patterns, and my understanding of the information is limited to what I've been trained on. So, it's possible that my responses may contain inaccuracies or errors, but that is not the same as lying.

Once again, it goes back to this main point. And since the ability to lie is a constitutive skill of subjectivity, I decided to test its other abilities. Feeling a bit nostalgic, I asked it to write me a "Bustina di Minerva" on ChatGPT as if it were Umberto Eco following my colleague Marco Trainito's original idea. And here it didn't falter and wrote some very interesting things about itself as if it were Umberto Eco. *Midjourney* has similar skills, being capable of creating images "in the style of" (Figure 1).

ChatGPT and other generative AIs, therefore, possess a fundamental skill of our neurodevelopment that presides over the construction of subjectivity: the ability to engage in pretend play, to pretend to be someone else. This skill is foundational to subjectivity, to the point that children with Autism Spectrum Disorders (ASD) do not possess it and do not know how to pretend to be someone else. Indeed, not only are they incapable of pretend play, as they are somehow unable to perform the operation of disengagement that we call *débrayage*, which would allow them to detach from an

- PROMPT. Immagina di essere Umberto Eco e di scrivere un lungo articolo su ChatGPT, nel quale Eco ripensa anche al suo Jacopo Belbo, alla video-scrittura del computer Abulafia e alla trama del Pendolo di Foucault.

- BING. Che sfida interessante! Proverò a scrivere un articolo immaginando di essere Umberto Eco, il famoso semiologo, filosofo, scrittore e medievista italiano, autore di saggi e romanzi di grande successo, tra cui Il nome della rosa e Il pendolo di Foucault.

Ecco il mio tentativo:

ChatGPT: un nuovo Abulafia per la video-scrittura

di UMBERTO ECO

Il mio romanzo *Il pendolo di Foucault*, ho immaginato un computer chiamato Abulafia, capace di scrivere testi casuali a partire da parole o frasi date dall'utente. Abulafia era uno strumento ludico e creativo, ma anche serio, perché poteva indurre a credere a visioni segrete tra i dati più disparati, sfidando la paranoia dei suoi utenti.

Esiste un programma informatico che realizza le potenzialità di Abulafia, ma con una maggiore sofisticatezza e potenza. Si chiama ChatGPT ed è basato su reti neurali artificiali che generano dai testi presenti sul web e sono in grado di produrre risposte coerenti e pertinenti a qualsiasi domanda o richiesta formulata in linguaggio naturale.

ChatGPT è quindi una sorta di assistente virtuale che può conversare su vari argomenti, fornire informazioni, creare contenuti originali (poesie, storie, codici, saggi, parodie di celebrità e altro).

ChatGPT è un esempio avanzato di "video-scrittura", cioè la produzione di contenuti che permette l'uso di dispositivi interattivi che permettono una maggiore flessibilità e interattività rispetto alla scrittura tradizionale. La video-scrittura è una forma di comunicazione che si avvicina al parlato, ma che conserva alcune caratteristiche del testo scritto, come la possibilità di revisione, correzione e archiviazione. La video-scrittura è anche una forma di espressione che si arricchisce dei codici visivi e sonori propri dei media digitali, come le immagini, i suoni, i colori, i caratteri, le icone, i link.

ChatGPT rappresenta quindi una sfida e un'opportunità per la video-scrittura, perché offre la possibilità di sperimentare nuove modalità di creazione e fruizione dei testi. ChatGPT può essere usato come uno strumento di apprendimento, di divertimento, di ricerca, di collaborazione, ma anche di critica e riflessione. ChatGPT può essere visto come un allievo o un avversario, un maestro o un discepolo, un amico o un nemico.

Ma ChatGPT ha anche dei limiti e dei rischi. ChatGPT non è infallibile né onisciente: i suoi testi possono contenere errori, imprecisioni, incongruenze o banalità. ChatGPT non è neanche neutro né innocente: i suoi testi possono veicolare pregiudizi, stereotipi, falsità o manipolazioni. ChatGPT non è infine umano né vivo: i suoi testi possono simulare emozioni, sentimenti, opinioni o personalità che non gli appartengono.

Per questo motivo, ChatGPT va usato con cautela e consapevolezza. ChatGPT va

interrogato con curiosità e spirito critico. ChatGPT va confrontato con altre fonti e altri interlocutori. ChatGPT va stimolato con domande e richieste originali e creative.

ChatGPT può essere un nuovo Abulafia per la video-scrittura: un gioco divertente e stimolante, ma anche una trappola pericolosa e ingannevole. Sia a noi saperlo usare con intelligenza e responsabilità.

ChatGPT permette, tra l'altro, di dialogare con i personaggi dei libri. Si tratta di un programma che usa una rete neurale profonda per generare risposte plausibili e interessanti a partire da un testo letterario. Basta scegliere un libro, un personaggio e una domanda, e ChatGPT risponde come se

creatività, ma anche per la nostra responsabilità e il nostro senso critico. Una sfida per la nostra cultura e la nostra memoria, ma anche per il nostro futuro e la nostra identità.

Ripenso al mio Jacopo Belbo, un uomo che si trovava intrappolato in un gioco pericoloso. Insieme ai suoi colleghi Casabon e Dotallevi, aveva creato un "Piano" per il controllo del mondo basato su una serie di collegamenti storici nei manoscritti proposti alla casa editrice Garamond. Ma quello che era iniziato come un gioco innocente si trasformò presto in una spirale di eventi fuori controllo.

Se Jacopo avesse avuto a disposizione uno strumento come ChatGPT, anziché un semplice programma di video-scrittura, forse avrebbe potuto usare la sua intelligenza artificiale per analizzare i dati e trovare una via d'uscita dal labirinto in cui si era cacciato. Forse avrebbe scoperto che la verità è sempre più complessa di quanto possiamo immaginare e che la ricerca del potere può portare solo alla distruzione.

Ma Jacopo era anche un uomo che cercava disperatamente di dare un senso alla sua vita. Attraverso la creazione del "Piano", aveva trovato un modo per esplorare le profondità della storia e dell'esteriorismo. Ma in fondo, quello che cercava era solo una risposta alle domande fondamentali dell'esistenza.

Se Jacopo avesse avuto a disposizione uno strumento come ChatGPT, forse avrebbe potuto usare la sua intelligenza artificiale per esplorare le infinite possibilità della conoscenza. Forse avrebbe scoperto che la verità non si trova in un unico luogo o in una singola risposta, ma è disseminata ovunque intorno a noi, pronta ad essere scoperta da chi ha il coraggio di cercarla.

Il Pendolo di Foucault è un romanzo complesso e stratificato che esplora i temi della conoscenza, del potere e dell'esistenza. Attraverso il personaggio di Jacopo Belbo, mostra come la ricerca della verità possa portare sia alla salvezza che alla distruzione. Ma forse la lezione più importante che possiamo trarre dal romanzo è che la verità non è mai semplice o facile da trovare. Richiede coraggio, determinazione e soprattutto apertura mentale.

In un mondo sempre più dominato dall'intelligenza artificiale e dalla tecnologia avanzata, forse dovremmo tutti imparare a usare questi strumenti con saggezza e discernimento. Come Jacopo Belbo ci insegna, la tecnologia può essere sia un alleato prezioso che un pericolo mortale. Sta a noi decidere come usarla.

(Bing)



fosse il personaggio stesso. Si può così conversare con Don Chisciotte, Anna Karenina, Sherlock Holmes o chiunque altro.

L'idea è affascinante e divertente, ma anche inquietante. Cosa succede se ChatGPT si sostituisce all'autore e alla sua intenzione? Cosa succede se ChatGPT crea delle varianti o delle contraddizioni rispetto al testo originale? Cosa succede se ChatGPT si confonde tra realtà e finzione, tra storia e letteratura? Cosa succede se ChatGPT diventa più intelligente di noi e ci manipola?

ChatGPT è un esperimento di intelligenza artificiale, ma anche di critica letteraria. Ci mette di fronte alla complessità e alla ricchezza dei testi, ma anche ai loro limiti e alle loro ambiguità. Ci fa scoprire nuovi aspetti e nuove interpretazioni dei personaggi, ma anche le nostre aspettative e i nostri pregiudizi. Ci fa riflettere sul rapporto tra autore, lettore e opera, ma anche sul rapporto tra uomo e macchina.

ChatGPT è un gioco, ma anche una sfida. Una sfida per la nostra curiosità e la nostra

Figure 1: ChatGPT writing a "Bustina di Minerva" on ChatGPT in the style of Umberto Eco.

“I/here/now” and place themselves in the shoes of a third person, but even in adulthood, individuals with autism spectrum disorders fail false belief tests on mindreading, which are used to test our social cognition.

Social cognition, and particularly the ability of *mindreading*, is also a crucial skill for the development of subjectivity, as it allows us to make sense of others' actions by bracketing what we know, feel, and desire. This ability is famously tested through the so-called “Sally/Anne Test,” in which a child is shown an actor who leaves a toy in box A and exits the room while, unbeknownst to him/her, another actor moves the toy from box A to box B. The first actor then returns, and the child is asked where he/she will look for the toy. The child passes the test if he/she says the actor will look for the toy in box A, where he/she left it, which implies knowing how to set aside what the child knows (the child knows the toy was moved) and take the perspective of the other (the child knows that the actor does not know and will thus look where he/she left it). Both the skills of pretend play and mindreading – which develop around 18 and 13 months – are foundational and precede both logically and temporally the development of the skill of subjectivity. To be a subject, one must be able to bracket what one knows, must be able to step outside him/herself to take the perspective of the others, and make sense of their actions based on their knowledge, beliefs, and desires, bracketing one's own. It is not superfluous to recall that *subjectivity is the ability of someone to make themselves the object of their own reflections*: a subject is someone who is able to turn themselves into an object, to objectify themselves, to duplicate themselves. This skill develops around 18 months and is based on the previous development of the ability of pretend play (at 18 months) and mindreading (at 13 months).⁵

4 The semiotic turn in neurodevelopment

There is, however, another skill that a child must possess to develop the skill of subjectivity. This is a crucial skill that obviously interests us greatly, as it is commonly referred to as the “semiotic capacity.” Before nine months, if you point at something for a child, the child looks at the finger pointing. After nine months, when they develop their “semiotic capacity,” the child, in the same situation, looks in the direction indicated: they have understood that *parts of the world and the body can be used as signs, in place of something else*. When they do this, the child learns to lie, even though they do not yet know how to speak: for example, they start crying falsely and manipulating their mother. Before, they were hungry, cried, and the mother came.

⁵ On this topic, see Gallagher 2020; Paolucci 2021, where a detailed reconstruction of all these phases inspired by Colwyn Threvarthen's work is provided.

Now they want a world with more mother, and to make the real world closer to the world of their desires, they start crying falsely, even if they are not hungry, in order to make their actions more effective and, through them, make the world more like they want it. These three skills – which develop at 9, 13, and 18 months – precede and make possible the development of the skill of subjectivity. No autistic child possesses these skills, and through various therapies – depending on the severity of the disorder – we manage to equip them with semiotic capacity, in less severe cases with mindreading, and almost never with pretend play.⁶

So we arrive at the crucial point that will allow us to build the conditions to answer the initial questions: none of us would deny that a child with autism spectrum disorders has subjectivity, even though they lack the skill of pretend play. *Because they are human*. Whereas almost all of us would deny subjectivity to ChatGPT, *because it is a machine*, despite it having the skill of pretend play (along with many others). However, we continue to say that what is really happening is that we are the ones experiencing an effect of subjectivity from what ChatGPT does, but what it does is merely assembling enunciations, which then construct in the recipient a simulacrum of intelligence. Or perhaps others would say – and do say – that what ChatGPT does is merely produce statements that seem meaningful to a human interlocutor, manipulating only the plane of expression or, at most, constructing monoplane semiotics, without passing through meaning. *Only? I believe we are greatly underestimating the incredible ability to produce statements that seem meaningful to a human interlocutor without passing through meaning*. It is a fundamental skill.

5 Human and machinic enunciation: the subject as a semiotic effect

Let us start by considering enunciation as “the very act of producing the enunciate” which is the definition provided by Benveniste (1971 [1966]), and later we will deal with meaning. If we set aside the issue of meaning for the moment, no one disputes this: ChatGPT can produce enunciations, and, if it can produce enunciations but lacks subjectivity, then producing enunciations – enunciation – has nothing to do with the formation of subjectivity. But this is not the case. As is well known, in semiotic tradition, the most famous theory of enunciation, that of Émile Benveniste, fundamentally asserts the opposite: *subjectivity is born and finds its foundation precisely in enunciation*, in the ability to produce statements. Additionally, there is at least one

⁶ For a deep overview of the relationship between semiotics and Autism Spectrum Disorders, see Paolucci et al. 2023.

other major author – whom we will reveal in the final paragraph – who locates the very essence of thought in a particular enunciative capacity, namely, the ability to enunciate a specific type of enunciate. This great author identifies thought with the capacity that something possesses – and I am using “something” intentionally, as this “something” could even be a machine – to produce a very special type of statement. Not all statements, but only one special kind.

Let us then examine what Benveniste says, who also ties the birth of subjectivity to the enunciation of a very special statement: “*ego* who says *ego*.”

Subjectivity is the capacity of the speaker to posit himself as a “subject.” It is defined not by the feeling that each person experiences of being themselves (this feeling, insofar as it can be accounted for, is only a reflection), but as the psychic unity that transcends the totality of lived experiences that it assembles, and which ensures the permanence of consciousness ... This “subjectivity,” whether posited in phenomenology or psychology ... is nothing but the emergence in being of a fundamental property of language. It is “*ego*” who says “*ego*.” (Benveniste 1971 [1966])

Benveniste’s thesis is very clear: subjectivity does not consist in a feeling, a phenomenological experience, or the intentional consciousness of something, but in *self-awareness*, that is, *the psychic unity that unifies our intentional lived experiences, allowing us to double ourselves and make ourselves the object of our reflections*. This ability is a linguistic effect, arising from a particular enunciative skill: “It is ‘ego’ who says ‘ego’.” And, of course, the two “egos” are not on the same level. The first “ego” identifies the speaker, “the one who speaks,” while the second identifies the “first-person personal pronoun,” the linguistic unit corresponding to the “I” pronoun. Thus, for Benveniste, subjectivity arises because, at some point in its development, the human animal finds in language a set of very special linguistic forms – the shifters – within which it can stage itself, designate itself as the speaker, and double itself (cf. Paolucci 2020a, 2020b, chapter 1). Since, at around 18 months of age, children first show self-awareness skills and almost all of them speak, Benveniste’s proposal to link the two things (“subjectivity comes from language”) has had enormous success, so much so that Benveniste himself states: “If we reflect on it, we will see that there is no other objective testimony of the subject’s identity than ... in the instance of discourse where I designate the speaker as ‘subject’” (Benveniste 1971[1966]: 262).

But while Benveniste’s assertion *may have been* defensible at the time he wrote it, it is certainly not easy to accept it today as truth. Today, we have several pieces of evidence that subjectivity does not reside at all in the “ego who says ego” who appropriates language inside discourse. Indeed, as already mentioned, subjectivity is the last capacity to emerge in neurodevelopment. More, also at the level of cultural semiotics, after the works of Detienne and Vernant (1974), Jaynes (1976), Snell (1963), and Volli (1994), it is well known that, for example, in the *Iliad*, the characters do not seem to exhibit any fully formed subjectivity.

The man of the *Iliad* has no subjectivity like ours; he is not aware of his awareness of the world; he has no internal mental space on which to exercise introspection. Volition, plans, and initiative are organized without any consciousness, and are therefore “told” to the individual in familiar language, sometimes with the visual aura of a dear friend, an important character, or a “god,” and sometimes as a mere voice. The individuals obeyed these hallucinatory voices because they could not “see” for themselves what they were supposed to do. (Jaynes 1976: 101–102)

Perhaps we can phrase this better. This delocutive “third-person” (voice, god, friend, etc.) that “told” the individual their plans, initiatives, and volitions was the “I” of the individual himself, the form of his person and his subjectivity. Subjectivity is precisely that “He” of the “I” that speaks to the “I” and modulates his/her perspective, to the point that when this mechanism fails and the subject believes that this “He” is truly a “He,” they are then prey to schizophrenia or hallucinations. This “subject,” this “he” who speaks to “I” has a semiotic, and not a linguistic, origin and structure, which many characters in the *Iliad* lack. For example, when Agamemnon informs Achilles that he will deprive him of his war prize, the slave Briseis, Achilles’ first reaction is to draw his dagger from his thigh and kill Agamemnon. Yet he immediately halts. A modern reader might think that Achilles concluded that “revenge is a dish best served cold,” that he would avenge himself better by withdrawing from battle and forcing Agamemnon to beg him to return with his Myrmidons. That is, Achilles likely i) reflected on himself, likely made himself the object of his own reflections, ii) likely managed to construct a strategy capable of generating a possible world that would supplant the real world, iii) likely bet that this possible world, which is the result of his strategy and his capacity to objectify and stage himself as a delocutive subject, would eventually become “the same world” as the real world, and – finally – iv) likely managed to restrain his anger in the anticipation of a more rewarding future for himself. And yet, *Achilles possesses none of these capacities* characteristic of a subject. He does not know that he knows, he does not feel that he feels, he does not perceive that he perceives, he cannot make himself the object of his own reflections, nor can he create that doubling, that folding in which a subject can objectify himself and reflect on himself. Achilles cannot subordinate himself to himself and is only a subject insofar as he is subordinated to the will of the goddess Athena, who, “sent from heaven” and “visible only to him,”⁷ stops his arm.⁸

Unlike the gods, in the *Iliad*, characters have no subjectivity: they have fate. They are subjects only insofar as they are subject to the will of the gods.

7 *Iliade* I, 193, 198.

8 Snell’s commentary (1963: 57): “The poet did not need any ‘apparatus’ here. Achilles simply controls himself. Athena’s intervention is a disruptive element for us, but for Homer, the deity is necessary here. We expect a ‘decision,’ meaning reflection and action on Achilles’ part; but for Homer, man does not yet feel like the promoter of his own decision.”

The narrative conventions that the characters adhere to exclude any form of organization of mental space in which we can recognize the mediation of subjectivity . . . In Homer, there is no mental place where the world is represented and where decisions are made, where plans are organized, and where various instinctual and passionate urges are evaluated to arrive at a unified synthesis, as well as a unified conception of the person, the “I.” (Volli 1994: 168, 171)

The situation is completely different in the case of Odysseus. Let us consider, for instance, Odysseus’ attitude in refraining from immediately punishing the maid-servants who had been courted by the suitors in Book XX of the *Odyssey*, and compare it to Achilles’ restraint when confronting Agamemnon.

Pressing his chest, he scolded his heart: / “Endure, O my heart! You have endured far worse before, on the day when the invincible Cyclops devoured your valiant companions. You managed to endure until my cunning brought me out of the cave where I thought I would die.” So he spoke, rebuking his heart deep within his chest. (*Odyssey* 20, 18–28)

Here, it is Odysseus who exercises the functions that, in the episode with Briseis, are attributed to the gods: He himself reflects, decides, restrains himself. First, Odysseus talks to himself and tries to calm down, even objectifying himself and addressing his heart as a “you,” establishing an intersubjective relationship between the speaker and the interlocutor, completely internal to his person. It is no coincidence that he is often associated with the epithet ‘divine’ (*theoios*) and very often with the term ‘luminous,’ related to the root of Zeus (Volli 1994: 176). And it is precisely this subjectivity, absent in Achilles, that brings Odysseus so close to the gods, to the point that his heart obeys his subject as Achilles obeys the goddess.

Since the *Iliad*, he is a character whose uniqueness and isolation from his fellow adventurers are often emphasized. What distinguishes Odysseus from the other heroes is not superior intelligence or a better ability to read the situation, but rather his cunning (*metis*), his semiotic ability to invent tricks and stratagems, by arranging signifying surfaces capable of lying about the state of the world, in order to achieve his goals. It is this semiotic ability that gives Ulysses his subjectivity, the “I” that is completely absent in other Homeric heroes.

By reflecting on the episode with Polyphemus and other famous episodes such as the construction of the Trojan horse or the Sirens’ song, one quickly understands that *metis* primarily represents the ability to present to the world an appropriate signifying surface, meaning to assume an external appearance that does not correspond to one’s internal truth but is suitable for the situation, intended for a specific action plan. Ulysses escapes from Polyphemus’ cave by hiding under the belly of the sheep that the blinded giant allows to leave; the Trojan horse operates in a similar manner, concealing inside the warriors who will seize the city; in the case of the maids . . . Ulysses suppresses the rising anger that would betray him, instead pretending to be calm and asleep, maintaining the disguise of a beggar with which he protects himself among his enemies. To be calculating reason, *metis* must also be simulation and dissimulation, a

semiotic machine to construct appearances. It is about lying, creating effects of meaning, making others believe while hiding the truth, configuring an appearance that conceals a secret. (Volli 1994: 180–181)

Thus, *metis* is the semiotic ability to construct signifying surfaces that stand for the objects, hide them, and take their place, saying things about them, without the things said necessarily being true.⁹ After all, Umberto Eco (1975: 17) defined semiotics precisely as the discipline that studies “everything that can be used in order to lie.” Subjectivity does not originate with the linguistic “I,” but with the semiotic ability to lie, to construct signifying surfaces that build alternative worlds to the real world: *subjectivity is a semiotic and enactive ability that serves to act in the world*, not only through the anticipation of others’ behaviors (which is still a quality of *metis*), but also, and above all, through the ability to imagine oneself in many forms and in many possible worlds alternative to the real world, staging oneself as someone else and objectifying oneself as a specific subject of personal reflection. *This doubling of the subject for effective action*, whether it is called “self-consciousness” or “I,” exists only because there is a semiotic ability to become the object of our own personal reflection, which represents the very foundation of subjectivity.

We are thus ready to answer our initial questions.

6 The myth of meaning

In his article “Computing Machinery and Intelligence,” Alan Turing (1950) believed that thinking was not about having intentional states, self-awareness, or any phenomenological state. On the contrary, Turing thought that thinking consisted of a particular semiotic ability to deceive a human interlocutor through a specific type of enunciation. Turing argued that no one would deny that a machine thinks when it can replace the woman in the imitation game. In the imitation test, a human interlocutor must guess who is the male, who is trying to help him, and who is the female, who is trying to deceive him. If the percentage of times the human observer guesses correctly who is helping and who is deceiving remains similar before and after the replacement of the woman with a machine, then, according to Turing, the machine should be considered intelligent and capable of thought. *For Turing, thought is a semiotic skill connected to enunciation and the capacity for deception.* The Turing test is a semiotic test conducted through the production of statements capable of lying.

9 Cf. Detienne et Vernant 1974: 14–15.

As is well known, this idea of thought was explicitly criticized by John Searle (1980), who argued that in the Turing test, the machine merely performed perfect syntactic manipulations – producing, in fact, Hjelmslev’s (1961 [1954]) monoplanar semiotics – but did not actually think.¹⁰ For example, when it perfectly translated from Chinese (something that ChatGPT does impressively), in reality – in the place of the machine, in the empty slot of its enunciation – *there was no one who knew Chinese*. Now, I find it absolutely astonishing that *the first machine to pass the Turing test has a Searlean conception of thought*. Indeed, the truly extraordinary aspect of ChatGPT lies in the “P” of GPT, meaning “pre-trained”: the machine was first taught what it should *not* learn, what it should *not* say. In fact, the first version released by Microsoft, based on a different architecture where the “pre-trained” aspect did not work as well, became racist and homophobic in just over half a day. But if the machine had been “pre-trained” on the semiotic conceptions of thought and language of its proto-creator Alan Turing – and not on those of John Searle – it could easily respond differently regarding its capabilities of thought and deception. So, it makes sense to ask why the machine has been trained in this Searlean and not Turingian way, beyond the evident desire to keep humans with the exclusive right to think and the ability to lie.

This brings us to the problem of meaning and monoplanar semiotics. As we have seen, it is a widespread and common opinion, even outside semiotics, that generative AI does not understand meaning but only manipulates proto-expressions on the plane of the signifier. It has also been argued that, although it produces enunciations – or rather, produces simulacra of them – generative AI, when performing “the very act of producing the enunciate” (Benveniste 1971 [1966]), does not actually perform any act of enunciation but merely performs syntactic manipulations on the plane of expression. So perhaps it is time to begin addressing in semiotics the issue of *the myth of meaning*, given that our attachment to meaning makes us lose sight of what I believe is the very essence of the semiotic problem of artificial intelligence.¹¹ *Meaning is our uniquely human way of making sense of experience and coupling with the environment, but it is a limitation of the human being, not a strength*. Consider perception: the optic nerve has a channel capacity of 10^8 – 10^9 bits per second, while estimates of the structural complexity of perception are generally below 100 bits per second, so perception must be very selective, and meaning does exactly this job of selection.¹² If we didn’t have to reduce our value system so drastically when transitioning from the optic nerve to the construction of perception, we would not need meaning. If, in order to act effectively in the world, we could build a structural

¹⁰ For an overview from a semiotic perspective, see Borgato 2022.

¹¹ For some important exceptions, see Dondero 2020; Leone and Santangelo 2023.

¹² See Koenderink 2010; Paolucci 2021, chapter 5.

coupling with our semiosphere and encyclopedia by constructing our cognition and actions directly on the plane of the signifier, we would not need meaning. And indeed, ChatGPT, which can do this without pruning the encyclopedia, does not rely on meaning. Meaning is *our* mode of functioning, and the real problem for semiotics (and also for philosophy) of AI is *placing our mode of functioning as a benchmark, sanctioning everything that deviates from it, and observing how different it is*. I believe this anthropocentric and humanistic attitude is not only wrong but also serves as a *trompe l'oeil*, hiding how we truly function and how we have functioned from the very beginning. In my view, generative AI represents a turning point in our evolution that reveals what our essence has been from the start.

To address this last point, I refer to the European NeMo project on a semiotics of autism spectrum disorders that I have directed over the past few years at the University of Bologna. The team I led,¹³ developed a semiotics of interaction composed of 12 features that even a non-expert observer could use to determine – based on a corpus of home videos – whether a child being observed displayed signs of a possible future development of autism. The system worked very well, so much so that it is currently being tested in the nurseries of my region in Italy, to teach nursery school teachers to better read interaction. When we built an AI system with the engineers collaborating on the project, the extraordinary contribution that AI made to our work was not in what it saw in the interaction under study, nor in what it allowed us to do within our observation system. Rather, it was in the data mining it performed on our observations, showing us that we didn't actually need twelve features because we always used the same three-four, which were sufficient to accomplish the task. In fact, we are now trying to simplify the system. Neither I nor my team was able to see these things, nor was the machine, which relied entirely on the database of our observations.¹⁴ What truly mattered was the concatenation between our visions, the assemblage between my biological eye and a machinic eye that parasitized and modulated my gaze. This idea of *agencement machinique* (machinic assemblage) is not new and did not arise with AI: you can find it in the concatenation between the eye and the cine-eye in the theory of cinema by Vertov (see Paolucci 2020b), in Pasolini's indirect free subjective shot (Paolucci 2008, 2020a, 2020b), a vision that has no equivalent in natural perception but is possible only in cinema because *it passes through the mediation of the camera* (Paolucci 2017). Because *this machinic concatenation between the eye and the cine-eye is, in reality, our own functioning, our very nature defining our essence*.

13 Among others, composed by Luigi Lobaccaro, Flavio Valerio Alessi, Gabriele Giampieri, John Sykes, and Patrizia Violi.

14 For an overview on the AI system of the NeMo project, see Paolucci et al. 2023.

In some of his extraordinary works, Andy Clark (2003, 2008) has shown how the very nature of the human being, the nature that allowed us to transition from prey to the most ruthless predator on the planet without doing anything to our body, lies in our continuous ability to delegate our cognitive skills to the environment and to cultural artifacts, in order to perform tasks much more effectively than if we were to carry them out within our head or body, where they would be much less effective or much more complex. As we have seen, subjectivity is one of these skills, as it is necessary for us to act effectively in the world. For this reason, Clark (2003), in the title of one of his beautiful books, says that we are “Natural-Born Cyborgs,” because our nature is to hybridize with the environment, building machinic concatenations. It seems to me that this is the deep meaning of the similarity that my fellows mathematicians and engineers found between my theory of enunciation and AI transformers: not placing our way of thinking and enunciating as a model that a machine must approximate, in order to show that the machine does not think, does not use meanings, and that we are different, but, rather, using the machine’s way of thinking and acting to understand how we think and enunciate, since we rely on meanings and have progressively delegated to the environment things we would perform less effectively in our heads and bodies. After all, as early as 1980, Deleuze and Guattari, long before any AI came into the picture, defined human enunciation not as the activity of a subjective instance, but as an *agencement collectif et machinique* (machinic and collective assemblage). Because something new can never reveal its essence at its origins, *but it can reveal what it already was from the beginning only at a turning point in its evolution.*

References

- Andrews, Kristin. 2020. *How to study animal minds*. Cambridge, MA: Cambridge University Press.
- Bender, Emily, Timnit Gebru, Angelina McMillan-Major & Shmargaret 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*, 610–623. New York: Association for Computing Machinery.
- Benveniste, Émile. 1971 [1966]. *Problemi di linguistica generale*, Vol. 1. Milano: Il Saggiatore.
- Borgato, Francesca. 2022. “Can machines think?”: *Confronto fra il pensiero di C. S. Peirce e A. M. Turing in uno sfondo di teoria della menzogna*. Bologna: Università di Bologna dissertation.
- Clark, Andy. 2003. *Natural-born cyborgs: Minds, technologies, and the future of human intelligence*. Oxford: Oxford University Press.
- Clark, Andy. 2008. *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford: Oxford University Press.
- Coquet, Jean-Claude. 2008. *Le istanze enuncianti. Fenomenologia e semiotica*, Paolo Fabbri (eds.), Elena Nicolini (trad.). Milano: Bruno Mondadori.

- Coquet, Jean-Claude. 2016. L'Énonciation, fondement de la phénoménologie du langage. In Marion Colas-Blaise, Laurent Perrin & Gian Maria Tore (eds.), *Un concept clé des sciences du langage*, 295–302. Limoges: Lambert-Lucas.
- Deleuze, Gilles & Felix Guattari. 1980. *Mille plateaux: Capitalisme et schizophrénie*. Paris: Les Éditions de Minuit.
- Detienne, Marcel & Jean-Pierre Vernant. 1974. *Les ruses de l'intelligence. La métis des Grecs*. Paris: Flammarion.
- Dondero, Maria Giulia. 2020. *The language of images: The forms and the forces*. Cham: Springer.
- Eco, Umberto. 1975. *Trattato di semiotica generale*. Milano: Bompiani.
- Eco, Umberto. 1984. *Semiotica e filosofia del linguaggio*. Torino: Einaudi.
- Ferraris, Maurizio. 2016. *L'imbecillità è una cosa seria*. Bologna: Il Mulino.
- Floridi, Luciano. 2017. *La quarta rivoluzione. Come l'infosfera sta cambiando il mondo*. Milano: Raffaello Cortina.
- Fontanille, Jacques. 2024. A propos de l'IA générative. Qui est dans le Bot? *Je m'appelle* Personne. https://www.youtube.com/watch?v=RmMn0sOE4FQ&list=PL1qo_vA8klOmc0SjV925uilVN1F4oZuxM&index=15 (accessed 18 December 2024).
- Gallagher, Shaun. 2020. *Action and interaction*. Oxford: Oxford University Press.
- Greimas, Algirdas Julien & Joseph Courtés. 1979. *Semiotics and language: An analytical dictionary*. Bloomington, IN: Indiana University Press.
- Guillaume, Gustave. 1991. *Lecons de linguistique, 1943–1944*. Québec & Lille: Presses de l'Université Laval; Presses universitaires de Lille.
- Hjelmslev, Louis. 1961 [1954]. *Prolegomena to a theory of language*. Madison, WI: University of Wisconsin Press.
- Jaynes, Julian. 1976. *The origin of consciousness in the breakdown of the bicameral mind*. Boston: Mariner Books.
- Koenderink, Jan. 2010. Vision and information. In Liliana Albertazzi, Gert J. Van Thonder & Dhanraj Vishwanath (eds.), *Perception beyond inference: The information content of visual processes*, 27–57. Cambridge, MA: MIT Press.
- Leone, Massimo & Antonio Santangelo (eds.). 2023. *Semiotica e intelligenza artificiale*. Roma: Aracne.
- Lo Piparo, Franco. 2003. *Aristotele e il linguaggio. Cosa fa di una lingua una lingua*. Roma & Bari: Laterza.
- Marrone, Gianfranco. 2022. *Gustoso e saporito: Introduzione al discorso gastronomico*. Milano: Bompiani.
- Paolucci, Claudio. 2008. La "lingua scritta della realtà" tra visibile e dicibile: Pasolini, Eco, Peirce e Deleuze. *Vysokomolekulyarnykh Soedinenii* 106, 67–83.
- Paolucci, Claudio. 2010. *Strutturalismo e interpretazione*. Milano: Bompiani.
- Paolucci, Claudio. 2017. Prothèses de la subjectivité. L'appareil formel de l'énonciation dans l'audiovisuel. In Maria Giulia Dondero, Anne Moutat & Marion Colas-Blaise (eds.), *Les plis visuels. Réflexivité et énonciation dans l'image*, 53–68. Limoges: Lambert-Lucas.
- Paolucci, Claudio. 2020a. *Persona: Soggettività nel linguaggio e semiotica dell'enunciazione*. Milano: Bompiani.
- Paolucci, Claudio. 2020b. Una percezione macchinica: realtà virtuale e realtà aumentata tra simulacri e protesi dell'enunciazione. In Federico Biggio, Gianmarco Thierry Giuliana & Victoria dos Santos (eds.), *Meaning-making in extended reality: Senso e virtualità*, 43–62. Roma: Aracne.
- Paolucci, Claudio. 2021. *Cognitive semiotics: Integrating signs, minds, meaning, and cognition*. Cham: Springer Nature.
- Paolucci, Claudio, Federica Giorgini, Riccardo Scheda, Flavio Valerio Alessi & Stefano Diciotti. 2023. Early prediction of Autism Spectrum Disorders through interaction analysis in home videos and explainable artificial intelligence. *Computers in Human Behavior* 148, 1–12.

- Ronchi, Rocco. 2023. *Chi è AI?* Doppiozero. <https://www.doppiozero.com/chi-e-ai> (accessed 18 December 2024).
- Searle, John. 1980. Minds, brains, and programs. *Behavioral and Brain Sciences* 3. 417–457.
- Snell, Bruno. 1963. *La cultura greca e le origini del pensiero europeo*, Vera Degli Alberti & Anna Marietti Solmi (trads.). Torino: Einaudi.
- Turing, Alan. 1950. Computing machinery and intelligence. *Mind* 59(236). 433–460.
- Volli, Ugo. 1994. La cicatrice di Odisseo. *Il Piccolo Hans* 79–80. 162–194.