What Is Distinctive About Human Intelligence in the Context of Artificial Intelligence?

A Philosophical Approach With Reference to Robert B. Brandom's Semantic Inferentialism

Robert Kublikowski

Abstract Since antiquity, humans have traditionally been characterised as *animal rationale* or *homo sapiens*. Such a definition takes into account, on the one hand, the physical aspect, referring to the human body, but also, on the other hand, the mental aspect. With this in mind, the present article seeks to develop a philosophical approach to the problem posed in its title.

Various aspects of human language use and cognition are, in the era of the information revolution, being increasingly taken over by AI. So what still remains specific to humans—or, more precisely, to human, natural intelligence, given this dynamically developing context?

The article addresses this question in a series of steps. Initially, it considers whether rationality is a good candidate for a uniquely human trait. In defence of the distinctiveness of human intelligence, and using Robert B. Brandom's semantic inferentialism, it then points to our ability and skill in understanding, as well as the normativity of language use and cognition. The ensuing discussion focuses on the normative categories of deontic status and deontic attitude, and related notions of commitment and entitlement that these in turn imply.

Keywords artificial intelligence; commitment; entitlement; natural intelligence; semantic inferentialism; understanding

✍ Robert Kublikowski, John Paul II Catholic University of Lublin, Poland ✉ robert.kublikowski@kul.pl
0000-0002-5199-3941

Forum Philosophicum 30 (2025) no. 2, 49–64
ISSN 1426-1898 e-ISSN 2353-7043
Subm. 9 July 2025 Acc. 11 September 2025
DOI: 10.35765/forphil.2025.3002.03

Acknowledgements
For valuable comments concerning the topic of this article I am very grateful to Prof. Robert B. Brandom (The University of Pittsburgh), Prof. Andrzej Bronk (The John Paul II Catholic University of Lublin), Prof. Piotr Kulicki (The John Paul II Catholic University of Lublin), Prof. Agnieszka Lekka-Kowalik (The John Paul II Catholic University of Lublin), Prof. Monika Walczak (The John Paul II Catholic University of Lublin), Dr Piotr Biłgorajski (The John Paul II Catholic University of Lublin), Dr Marcin Grabowski (The John Paul II Catholic University of Lublin), Dr Tomasz Łach (The John Paul II Catholic University of Lublin), Dr Łukasz Sarowski (The John Paul II Catholic University of Lublin), Dr Kamil Szymański (The John Paul II Catholic University of Lublin), PhD student Justyna Horbowska (The John Paul II Catholic University of Lublin) and PhD student Damian Szczęch (The John Paul II Catholic University of Lublin). I am also grateful to two anonymous reviewers.

Since antiquity human beings have traditionally been characterised as *animal rationale* or *homo sapiens*. This definition takes into account, on the one hand, the physical aspect, referring to the human body and its classification within biological anthropology, but also, on the other hand, the mental aspect, which belongs to the domains of psychology and philosophy. With this kept firmly in mind, the present article seeks to develop a philosophical approach to the problem posed in the title.

At first glance the task seems simple: we have specific bodies and minds that allow us to experience emotions, make free choices and, above all, use reason (intellect) in ways that display varying degrees of intelligence in respect of thinking, language use and cognition. Our distinctive culture also encompasses society, morality, religion, art, law, politics and the state. These remain our specifically human—both natural and cultural—traits.

However successive aspects of human language usage and cognition—particularly in the era of information revolution—are being "appropriated" or taken over by artificial intelligence (AI). So what remains as distinctive of human beings—or, more precisely, of human, natural intelligence—in the context of rapidly developing AI?

The question will be addressed here in a series of steps. First, I will consider whether rationality is a good candidate for a uniquely human trait. In defence of the distinctive character of human intelligence—and drawing on Robert B. Brandom's semantic inferentialism—I will then appeal to our ability and skill in understanding, as well as to the normativity of language use and cognition. Subsequently, I will focus on such normative categories as deontic status (commitment and entitlement) and deontic attitude (including the undertaking and attributing of commitments, as well as the attributing of entitlements).

1. Is Rationality What Is Distinctive About Human Intelligence?

It seems that, according to the received definition of a human being as *animal rationale*, our distinctive trait is indeed rationality, which characterises both non-linguistic and non-cognitive actions, as well as linguistic and cognitive ones.

One can distinguish between practical and theoretical rationality. The former consists in striving for rational agency—that is, actions that satisfy our needs or desires. The latter refers to the rationality of cognition, aimed at attaining a system of true beliefs about the world (Audi 2004, 17 ff.). Human rationality pertains to the experience of emotions, to decision-making, as well as to thinking, language use and cognition. Cognition includes empirical (sensory) cognition: perception, intuitive (immediate, non-sequential) cognition, discursive (sequential) cognition, as well as creativity and memory.[1]

Surprisingly, emotions can be understood as either irrational or rational. At times, they serve as a prelude to the formation of thoughts (Eemeren 2015, 204). Emotive passion—such as enthusiasm, etc.—supports decision-making and intellectual engagement.

We human beings pride ourselves on our intelligence, and one of its defining features is that we can remember our earlier thinking and reflect on it (Dennett 2013, 26). Intelligence underlies rationality, which is expressed in our thinking capacity and skills—especially in remembering and engaging in reflection.

Our rational linguistic and cognitive actions are evident in our ability to analyse, synthesise (summarise), pose questions, define, reason, discuss, etc. However, AI is also able to use language or get to know things in a surprisingly rational and correct way.[2] One only need mention such spectacular applications of it as automatic translators, data analysis and decision support in astronomical observations, in medical diagnostics or military operations, etc.[3]

1. Our rationality is also revealed in the creativity of cognition. It often consists in—previously unimaginable—questioning of the rules of the system to which the cognition pertains. This may involve the challenging of hidden—sometimes false—assumptions or principles underlying a given research program or scientific theory (Dennett 2013, 45–7).

2. It is important to remember that rationality is originally attributed to human, natural intelligence and only secondarily to AI, which is a human, technological artifact.

3. For more examples of AI applications, see, e.g., Bostrom (2014), Floridi (2014), and Russell and Norvig (2021).

Let us focus on three examples of AI applications, and specifically of Natural Language Processing (NLP) applications which use large language models (LLMs): namely, IBM's Watson, Apple's Siri and OpenAI's ChatGPT.

Watson is used in some online customer services, and has been adapted for medical applications, such as assessing cancer treatment options. It goes beyond simply answering straightforward questions, as does Siri. It is also capable of handling complex riddles as featured in the game show *Jeopardy!*, in which players are not given direct questions but rather clues, and must infer the question that fits the clue. For example: "On May 9th, 1921, this airline opened its first passenger office in Amsterdam. Its name consists of three consecutive letters of the alphabet." The correct answer—in the form of a question—is: "What is KLM?" Watson is capable of meeting such challenges. Unlike Siri, Watson's *Jeopardy!* version does not have access to the Internet (though the medical version does), and it does not understand the structure of conversation. It also cannot obtain answers through logical reasoning. Instead, it relies on parallel statistical searches through a vast but closed database. This database includes documents—countless summaries and source texts, as well as *The New York Times*—providing factual information on a wide range of topics. For *Jeopardy!* Watson's searches are guided by hundreds of specially designed algorithms that reflect the probabilities within the game. Watson can also learn from the answers of its opponents. However, it still makes mistakes. Even in everyday fact-finding tasks, humans often rely on judgments about the adequacy of usage, which remains beyond Watson's reach. For example, one clue in *Jeopardy!* required identifying two of Jesus' disciples whose names are among the ten most popular babies' names and end with the same letter. The correct answer was "Matthew and Andrew"—and Watson gave it immediately. A human gave the same answer—but only after first considering "James and Judas" and then rejecting this, because he or she thought "Judas" was not a popular name for a baby. Watson would be incapable of making that kind of reflection. (Human judgments of adequacy or relevance in language usage are often much more subtle than in this example.) Relevance is the linguistic/conceptual version of the well-known "frame problem" in robotics: the difficulty of determining what is important and what is not in a given situation. It may be that the frame problem will never be fully solved by a non-human system—perhaps due to the complexity of "the frame problem," or because relevance is rooted in our specifically human form of life (Boden 2018, 81–82).

Siri is a personal assistant—a speaking chatbot—that can quickly answer a wide variety of questions. It has access to the entire Internet, including

Google Maps, Wikipedia, the constantly updated *The New York Times* and many local services concerning taxis and restaurants. It also uses the powerful WolframAlpha program—a tool for answering questions—which can use logical reasoning to infer, rather than merely to find, answers to a wide range of factual questions. Siri accepts spoken questions from a user (gradually adapting to the user's voice and dialect) and responds using web searches and conversational analysis. Conversational analysis studies how people organise topic sequences in a conversation and how they structure interactions such as explanation and agreement. This allows Siri to consider questions such as: "What does the speaker want?" and "How should I respond?," and, to a certain degree, adapt to the interests and preferences of the individual user. In summary, it appears to be sensitive not only to topic relevance, but also to personal relevance as this pertains to language use. This seems impressive—at least on the surface. However, Siri is easily led astray and often gives silly answers. And if the user strays beyond known facts, Siri gets confused (Boden 2018, 80–81).

Despite some deceptively impressive examples—such as Watson, Siri, or cases of machine translation—modern computers do not understand what they "read" or "say." Google's search engine retrieves terms and estimates the plausibility of their usage, but this evaluation is statistical. Search engines—and, more generally, NLP systems—are able to find relationships between words and/or texts, but they do not possess understanding (Boden 2018, 56).

ChatGPT does not retain human-like conscious thinking or understanding, but it increasingly simulates such capacities in a way that is specific to its own architecture. It learns linguistic usage patterns—relatively effectively—from large datasets. Its functioning amounts to a complex form of language processing. ChatGPT employs algorithms to analyse expressions into units smaller than sentences (so-called tokens, which include words, parts of words, punctuation marks, etc.). However, it does not treat these tokens as textual elements in the human sense; rather, it generates approximate internal representations of these elements and their meanings in the form of numbers—more precisely, semantic or numerical vectors (known as embeddings) situated within a multidimensional mathematical space. For instance, the vectors corresponding to semantically similar words—such as "cat" and "dog"—are located close to each other, whereas the vectors for "cat" and "democracy" are placed much farther apart. In this way, ChatGPT is able to "compare" semantic relationships between expressions: both similarities and differences. It predicts (or estimates) which unit is most likely to appear next in a given sequence. For example, in the

sentence "The sun rises in the . . . ," the most probable next token is "east." Additionally, ChatGPT applies the so-called attention mechanism, which allows it to take linguistic context into account so as to maintain coherence within the topic being developed. For instance, in the sentence "Anna told Maria that her dog had run away," the pronoun "her" is ambiguous. It is unclear to whom the possessive pronoun refers. The model attempts to determine which referent is statistically most probable in the given context. A key issue, however, is that if ChatGPT's training data contains false information, the model may retrieve and process that falsehood. Put simply, ChatGPT does not have "direct" access to reality, and this can lead to errors, including so-called hallucinations. A partial solution to this problem is Internet access, which can help test relevant information. Nevertheless, ChatGPT requires critical control and correction.[4]

LLMs—such as ChatGPT, etc.—have been parameterised in such a way that, given an input sequence, they calculate (predict) its most probable continuation (i.e. an output sequence). However, this does not mean that LLMs understand (or interpret) the text, as they do not operate directly on its content (Landgrebe and Smith 2025, 10). People using AI-powered, conversational systems—like ChatGPT—learn the "style" of these chat-bots and the kinds of mistakes they make. LLMs can certainly impress us with their conversational capabilities. However, experienced users know what to say in order to reveal the difference between human, general intelligence (which enables us to perform a wide range of tasks) and AI—a computer program that is only "capable" of performing specific tasks (Togelius 2024, 37)[5].

Comparing our achievements to those of AI, it seems that AI is catching up with us—or even surpassing us in some areas (e.g., the speed of answering questions). However, as has already been noted, a fundamental issue remains: the aforementioned AI applications do not understand language—or, to put it more cautiously, they do not understand language in a human way. This seems to mark, at least for now, one of the key differences between such systems and natural human intelligence. Let us therefore now turn our attention to the concept of human understanding.

4. ChatGPT's Language Understanding, based on: https://chatgpt.com/ (retrieved July 4, 2025).

5. In computer science, a distinction is made between the as yet unattained artificial general intelligence (AGI), which aims to simulate human intelligence, and the currently available narrow AI systems that specialise in performing specific tasks.

2. An Attempt to Defend the Distinctive Character of Human Intelligence

We humans possess properties that are distinctive, especially in the context of rapidly developing AI: this is the hypothesis I hope to defend here, with reference to the position called semantic inferentialism worked out by Robert B. Brandom.[6]

When it comes to our own self-cognition, we should pay more attention to what we are capable of doing, rather than our origins or our biochemical structure (Brandom 1998, 3 ff.). Emphasising the pragmatic, dynamic aspect—namely, our capabilities—is justified. However, it is worth remembering that our origins—understood as our beginnings—are something fundamental, as is the biologically grounded structure dependent on these. They set the scope of our potential and actual actions, whereas AI, by contrast, as a human artifact, operates on the basis of the data it has been provided with and the programmed methods for processing that data.

Among the cognitive capacities constituting human mentality one can distinguish that of sentience—the reception of sensory stimuli—and that of sapience (or wisdom). The former is a biological property shared by other animals and humans alike, whereas the latter is connected with human intelligence and understanding (Brandom 2001, 157). What is important here is sapient consciousness or awareness—a sort of practical mastery or kind of know-how (Brandom 1998, 88 ff.; see also Brandom 2019).

Several questions arise. Firstly, can AI achieve understanding? Secondly—and a more important question for this article is the following one: what is this understanding that is distinsctive of human natural intelligence?[7]

6. There are also other approaches to defending the distinctive character of human intelligence. Dreyfus (1992) argued that complex mental processes cannot be fully represented by the logical apparatus on which computers are based. Therefore, AI is not capable of surpassing natural intelligence. Meanwhile, Landgrebe and Smith (2025) claim that it is mathematically impossible for AI to equal or exceed human natural intelligence. In other words, human intelligence will retain its distinctive character and its advantage. They support this claim with the following arguments: (1) human intelligence is a capacity of the human brain and central nervous system, which is a complex, dynamic system; (2) such systems cannot be expressed in mathematical language in a form that would enable their operation within a computer.

7. Incidentally, one may ask whether AI is capable of receiving sensory stimuli. Of course, it is possible to record data through a camera connected to a computer, which is equipped with RAM and a hard-drive memory. However, is this already a case of perception—i.e. perception analogous to that of humans?

2.1. Human Understanding and the Normativity of Language Use and Cognition

Understanding is the fundamental goal of our social, normative and discursive linguistic-cognitive practices (Brandom 2001, 6). Grasping an expression is manifested in our distinguishing its correct from its incorrect uses (Brandom 1998, 13–14, 32). Words such as "correct," "appropriate," "proper"—and their opposites—really serve to indicate the normative aspect of our rationality, intelligence and engagement through language and cognition.[8] By contrast, AI is programmed to operate in accordance with rule-governed norms.

Understanding requires knowing what someone thinks or says about something, as well as what the relation between these two approaches is (Brandom 2001: 158).[9] An additional factor is how we do it (Brandom 1998: 120). Here, the distinction between the content of cognition and the object of cognition is emphasised.

The grasp of an expression manifests itself in its correct use in appropriate circumstances, and in controlling the consequences of such usage (Brandom 1998, 120). For example, a parrot, when pointing at a red object, "knows" the circumstances of using the term "red," but it does not "know" the consequences of its use: namely, that if something is red, then it is colourful (Brandom 2001: 62–66; see also 70–71, 148). In other words, it concerns what the extra-linguistic and linguistic context of an utterance is, and what follows from this. The extra-linguistic context might be, for example, the room where the cage with the parrot stands, and the red scarf lying on the table that the parrot is pointing at. The linguistic context will be the sentence "Something is red," occurring alongside the sentence "Such a thing is colourful."

Understanding is the ability to "navigate" in the space of reasons, or in terms of the cognition of reasons (Brandom 1998, 5).[10] The feature of rationality that qualifies humans as sapients can be identified with the

8. For more on normativity, see, e.g., Hattiangadi (2007).

9. Brandom is the originator of the anaphoric theory of reference and truth, according to which both reference and truth are understood within language rather than as relations between language and the world (see Brandom 1998: 275 ff.). Consider the following sentences: "The National Gallery in London is located relatively close to the Houses of Parliament. It contains very valuable collections." The word "It" in the second sentence refers back to "The National Gallery in London." This is an example of anaphoric reference (i.e. reference internal to language). The same applies to the category of truth: "The Houses of Parliament are situated on the Thames. This sentence is true." Here, the second sentence refers to the first, again demonstrating an anaphoric structure. In a similar way, LLMs operate using the aforementioned attention mechanism, which enables them to analyse the relations between expressions within a linguistic context.

10. For more on reasons, see, e.g., Skorupski (2010).

capacity to participate in the game of giving and asking for reasons (Brandom 2001, 81; see also 1998, 230). The conscious human mind has a distinctive capacity for rational action: i.e. using reasons. This concerns action broadly understood: both extra-linguistic and extra-cognitive, as well as linguistic and cognitive. It is precisely consciousness—awareness of one's own rational activity—that constitutes the fundamental difference between human intelligence and AI.

Participation in the game of giving and asking for reasons involves—and this goes right to the core of our linguistic, cognitive, social and discursive practices—the inferential articulation of a claim. It is in inference (reasoning) that a reason for a given claim is obtained (Brandom 2001, 161–65). Understanding a term—knowing its meaning, where this is tantamount to grasping its concept—requires competence in handling the inferences in which the given term is used. Understanding requires knowledge in the practical sense—i.e. a kind of knowing how: the skill that enables one to distinguish which claims are inferentially connected to the claim in which a given term occurs (Brandom 2001, 48). For example, to explicate to someone how to understand the word "friend," we can present the following inferential consequences of a sentence containing this word:

If *a* is a friend of *b*, then *a* is honest towards *b*.
If *a* is a friend of *b*, then *a* cares about *b*.
If *a* is a friend of *b*, then *a* is faithful to *b*.
If *a* is a friend of *b*, then *a* does not act to the detriment of *b*, etc.

Understanding, then, is a kind of practical competence in respect of knowing how to participate in the game of giving and asking for reasons that support a given proposition, how to establish what the reason for something is, how to distinguish good and bad reasons as well as what the score of the ongoing linguistic-cognitive practice (the game) is, and how to change the score of such a game in which both "parties" have some commitments and entitlements (Brandom 2001, 165). What is at stake here is a inguistically-cognitive, social, normative, and discursive game of giving and asking for reasons, played out in the real world. In specific instances, this may take the form of games such as bridge, checkers, or chess. Understood in this way, such a practice requires the existence of human communities and conscious participants (players). It is therefore not the same as a game of chess played between a human and a machine.

One might critically object that AI is also subject to certain commitments and entitlements—for example, those defined by the rules of chess:

the obligation to make legal moves and the entitlement to move when it is one's turn. The crucial difference, however, lies in the fact that such a machine-player lacks consciousness—as well as everything that is inextricably bound up with the latter.

Our understanding of the concept to be explicated therefore just consists in our practising its inferential use. This concerns the ability to determine what someone is still committed to when they apply that concept. It concerns knowledge about what would entitle someone to do something and what would prevent such entitlement (Brandom 2001, 11). The categories of commitment and entitlement will be elaborated in the next part of the article.

2.2. Normative Categories: Deontic Status and Deontic Attitudes

In Brandom's pragmatic theory, the primary concepts are the normative concepts of deontic status and deontic attitude. Deontic status is a result of someone's attitude towards certain actions, or their effects. Such a status consists of a commitment and an entitlement. The deontic attitude manifests itself in undertaking or attributing a commitment, and attributing an entitlement. Undertaking a commitment is connected with attribution of an entitlement (Brandom 1998, 157 ff.; see also Scharp 2005, 206 ff.). Let us now analyse more precisely what deontic status and deontic attitudes amount to.

2.2.1. Deontic Status: Commitments and Entitlements

The formation and development of concepts—meaning of terms with which we think and speak—takes place by making explicit what is implicit in practices of language use that previously went unquestioned. Meaning, though explicit, can be further explicated in the context of social cooperation. In various conversational situations, interlocutors formulate claims, as well as arguments for or objections against them, and also consider their possible consequences, along with ways of obtaining the entitlements needed to recognise certain claims (Brandom 2001, 149).[11] In linguistic practice, we

11. From early childhood—driven by curiosity about reality—we ask questions such as "What is this?", while simultaneously pointing at the object of our inquiry. In response, we receive successive names, and intuitively begin to form abstract concepts related to the concrete world around us. Through this process, we learn to use appropriate terms for newly recognised, similar things. In training AI, we follow an analogous approach by presenting various examples (patterns). LLMs—unlike humans—have only indirect access to reality, relying on information gathered from various sources (such as the Internet, etc.). By contrast, robots (e.g., humanoid robots)—if equipped with sensors, cameras etc.—have some degree of contact with the real world.

not only form our utterances but also explicate or clarify them. This is done by giving a direct presentation in a clearer form, as well as in discussion, when our claims—even justified ones—are questioned.

We can also speak about content that remains implicit within explicit claims: that is, those implicit inferential consequences of a given claim that have not yet been drawn out. In the context of the network of our inferential practices, expressing acceptance of, or commitment to, one proposition amounts to an implicit expression of acceptance of, or commitment to, another proposition that follows from that claim (Brandom 2001, 18).

For a particular speech act, certain conditions are set under which—according to the practices of the language community—someone is committed or entitled to perform that speech act. What the act changes with respect to other language users, and how the act impacts on their commitments and entitlements, is also determined (Brandom 2001, 129). An important aspect of a discursive practice will be the interpersonal communication through which what individual interlocutors are committed and entitled to comes to be fixed. Entitlements to commitments are reciprocally transferred (Brandom 2001, 165). Subsequent speech acts modify the linguistic practice of the community in question. (This could be a political debate, a marital quarrel, etc.) Individual behaviours modify previous commitments or entitlements. For example, a married couple might be arguing about where to go on vacation. He strongly believes that a trip to London for the Wimbledon Championships is the best option, while she does not like sport, and does not want to go to London. She dreams of going to Paris for sightseeing and shopping. The husband proposes, in conciliatory fashion: "Let's go to London and Paris! If you go to London, you don't have to watch tennis. You can see the sights and go shopping. And I—as a reward—will go shopping with you every day in Paris without complaining!"

An instance of correct linguistic communication will consist in providing someone with a sufficient number of clues from which to infer what a given person intends to commit themself to by making the individual claims that they put forward, and what they are entitling someone else to do through undertaking such commitments. An error in recognising these components will be tantamount to an error in recognising the inferential commitments involved (Brandom 2001, 64).

A given inferential "move" can be justified or entitled by other "moves," can entitle further "moves," and exclude yet other ones (Brandom 2001, 162). For example, the commitment associated with asserting that *a* is a dog does not imply a commitment to asserting that *b* is a mammal, but it does imply a commitment to asserting that *a* is a mammal. Similarly the judgment

that *a* is a dog is not incompatible with the judgment that *b* is a fox, but it is incompatible with the judgment that *a* is a fox (Brandom 2012, 295).

As can be seen, commitments and entitlements are fundamental, social, normative aspects of discursive practices (Brandom 1998, 159 ff.). Let us illustrate the categories analysed here with the following example: someone goes to the British Library to open an account. On the one hand, they fill out and sign a form, thereby committing them to behaving according to the library's rules (i.e. committing them to not making noise, not damaging the book collection, not smoking in the reading room, etc.). By possessing a valid card, on the other hand, such a person obtains an entitlement to use the library's resources. This example will be detailed below with categories such as undertaking or attributing a commitment, as well as attributing an entitlement.

2.2.2. Deontic Attitudes: Undertaking or Attributing Commitments, and Attributing Entitlements

To think or say that things are thus-and-so is to undertake a specific type of inferentially formulated commitment towards the claim expressed. It concerns expressing a claim about a given thing in a correct premise in subsequent inferences. It also concerns the entitlement to use such a claim as a premise (Brandom 2001, 11). Commitments are rational when undertaken on the basis of appropriate entitlements obtained on the basis of the reasons one is prepared to entertain (Brandom 2001, 80).

Someone holding a certain belief applies it when undertaking further cognitive commitments. For example, if someone believes that a perceived object is red, then this belief imposes on them a commitment towards the belief that the perceived object is colourful—that it is red, more precisely scarlet, and that it is not green (Brandom 2001, 108–09; see 45–49).

A participant in the game of giving and asking for reasons understands the discursive significance of some speech act if they can attribute to that act the appropriate commitment regarding the claim used in a reasoning. In addition, they must attribute an entitlement to that commitment. Furthermore, the "player" attributes truth to the claim by the very fact of its use (Brandom 2001, 165–69). Let us recall that the attribution of truth to a sentence just consists in its assertion.[12]

Claiming is identical to acknowledging or undertaking a commitment regarding a given proposition. Meanwhile, undertaking a commitment

12. Assertion—even if it actually concerns a true claim—expresses knowledge on condition that someone making it understands the claim (Brandom 1998, 214).

related to a proposition is an act that entitles the attribution of that commitment both to somebody and to the proposition. Undertaking such a commitment can occur through using the proposition as a premise in practical reasoning (Brandom 2001, 173–77).

Here are some examples of practical inferences that clarify the normative aspect of a social, communicative and discursive practice:

(1) Only opening an umbrella will protect me from getting wet. Therefore, I ought to open the umbrella.
(2) I am a bank employee and I am going to work. Therefore, I ought to wear a necktie.
(3) By repeating a rumour, I might harm someone without a reason. Therefore, I ought not to repeat the rumour.

The use of the word "ought" is meant to express the significance of the conclusion, understood as an instance of undertaking a certain commitment (Brandom 2001, 84–85). Someone who evaluates the inference "I am a bank employee and I am going to work. Therefore, I ought to wear a necktie" will find it correct for any person *a* who makes such an inference. The evaluator undertakes a cognitive commitment regarding the inference: namely, that *a* is a bank employee. Such a commitment differs from attributing to someone a certain desire. In this case, the norm (obliging the bank employee to wear a necktie), which is an implicit assumption of the inference, is linked to holding a certain status: namely, the status of a bank employee. It is about holding this status rather than expressing a certain desire or preference. Whether someone has a reason to wear a necktie simply depends on whether they have the relevant status or not (Brandom 2001, 91).

A (discursive) cognitive practice is an activity of deontic "scorekeeping" (Brandom 2001, 81). In updating the "score" acquired by a participant in the linguistic-cognitive game, assertion-making plays a pragmatically significant role (Wanderer 2008, 123). Individual speech acts will change a person's linguistic-cognitive commitments and entitlements (Brandom 2001, 81). In communicating—through successive instances of "scorekeeping"—various new discursive commitments and entitlements are acknowledged by them and attributed to others, with certain previous commitments and entitlements being rejected (Loeffler 2018, 185).

The socio-historical, discursive process of "scorekeeping" forms a system of acts of undertaking or attributing commitments and entitlements in ongoing interactions. Each participant in such a process constantly

registers and evaluates the actions of the other one (Loeffler 2018, 203–04). Holding a belief is a kind of a commitment understood as taking a certain position within an inferentially connected network. If someone—holding a certain belief—undertakes a specific cognitive commitment, it affects the acceptance of the consequences of that belief as well as the rejection of beliefs that are inconsistent with it (Brandom 2001, 118–19). Individual participants, or at least their acts within the linguistic-cognitive game, are mutually connected, forming a network (system, structure) whose anticipated value is taken to reside in the coherence of beliefs and related commitments or entitlements.[13]

The earlier example concerning the library can be extended by stating that the librarian attributes to someone—who opens an account—the commitment to follow the library regulations. That person then undertakes this commitment, becoming responsible for their behaviour on library premises. At the same time, the librarian attributes to the reader appropriate entitlements as regards using the library resources. Smoking in the reading room would be regarded as a misunderstanding of the regulations, and of the commitments attributed and undertaken. In other words, such behaviour would be inconsistent with those commitments.

Conclusions

What distinguishes us from AI are our natural bodies. However, humanoid robots are increasingly impressive in their appearance and behaviour.[14] Besides that, we are distinguished by our minds: emotions,[15] free decisions[16] and especially consciousness and reason responsible for rational thinking, the use of language and sensory, intuitive[17] and discursive cognition, as

13. AI—for example, ChatGPT—thanks to the aforementioned attention mechanism, is capable of adhering to the rule of coherence in a generated text. However, AI lacks consciousness, and therefore has no beliefs.

14. See, *inter alia*, Actroid from Osaka University in cooperation with Kokoro Company Ltd. (2005), Honda's Asimo (2011), Nadine from Nanyang Technological University (2014), Hanson Robotics' Sophia (2016), Tesla's Optimus (2023), Boston Dynamics' Atlas Electric (2024), and Human Plus from Stanford University (2024), in: https://en.wikipedia.org/wiki/Humanoid_robot (retrieved July 4, 2025).

15. AI is capable of recognising emotions and adjusting the tone of its speech, as well as its facial expressions, to the emotional state of the interlocutor.

16. Human beings possess free will—naturally grounded and culturally developed. In contrast, AI-controlled vehicles and similar systems are said to exhibit autonomous (independent) behaviour, based on software developed by humans.

17. Intuition and creativity, as distinctive characteristics of human intelligence, constitute a subject that would require a more comprehensive analysis than is possible within the scope of the present article.

well as creativity and memory. However, it is not the case that only we are capable of operating in terms of language, cognition (including both sensory and discursive cognition) and memory: AI also is.

An important difference between AI and our own natural intelligence is our distinctively human capacity for understanding. Human understanding is connected to pragmatic and normative categories such as inference,[18] commitment, the undertaking and attributing of commitments or entitlements, and also the attribution of entitlements related to claims, concepts, etc.

Intelligent machines, on the other hand, seem to understand, but in fact this is the effect of using LLMs, which employ machine learning—applying probabilistic calculations and statistics—executed with access to large databases. For example, if we enter the letter "y" into an Internet search engine, AI will complete this as the phrase "youtube." This happens not because the search engine understands that this is what we mean and want to watch on YouTube. The search engine performs this completion because previously—many times when typing "y"—we most often added "outube," and then used the YouTube channel. This is therefore not yet a humanly intelligent act of understanding. One can, however, speak of artificial or machine understanding.

Someone enthusiastic about technology might claim that the topic of the distinctive character of human, natural intelligence, and its distinguishability from AI, is unimportant, as the latter achieves such spectacular results. Yet it is possible that an analysis of the unique characteristics of our own intelligence (such as our understanding, etc.) may prove to be a source of inspiration for the further development of AI. Who knows what the future holds for this rapidly developing field?

References

Audi, Robert. 2004. "Theoretical Rationality: Its Sources, Structure, and Scope." In *The Oxford Handbook of Rationality*, edited by Alfred R. Mele and Piers Rawling, 17–44. Oxford: Oxford University Press.

Boden, Margaret A. 2018. *Artificial Intelligence: A Very Short Introduction.* Oxford: Oxford University Press.

Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies.* Oxford: Oxford University Press.

Brandom, Robert B. 1998. *Making It Explicit: Reasoning, Representing, and Discursive Commitment.* Cambridge, MA: Harvard University Press.

18. Reasoning (inference) is not our exclusive domain, as AI is also capable of performing inference: i.e. analysing related circumstances and presenting the consequences of a given statement. The difference, though, lies in the fact that in our case such inferences are related to additional elements, such as beliefs, commitments, entitlements, etc.

—. 2001. *Articulating Reasons: An Introduction to Inferentialism*. Cambridge, MA: Harvard University Press.
—. 2012. *Between Saying and Doing: Towards an Analytic Pragmatism*. Oxford, New York: Oxford University Press.
—. 2019. *A Spirit of Trust: A Reading of Hegel's Phenomenology*. Cambridge, MA: Harvard University Press.
Dennett, Daniel C. 2013. *Intuition Pumps and Other Tools for Thinking*. New York: W.W. Norton & Company.
Dreyfus, Hubert L. 1992. *What Computers Still Can't Do: A Critique of Artificial Reason*. Cambridge, MA: MIT Press.
Eemeren, Frans H. van. 2015. *Reasonableness and Effectiveness in Argumentative Discourse*. Cham: Springer.
Floridi, Luciano. 2014. *The 4th Revolution: How the Infosphere Is Reshaping Human Reality*. Oxford: Oxford University Press.
Hattiangadi, Anandi. 2007. *Oughts and Thoughts: Rule-Following and the Normativity of Content*. Oxford: Oxford University Press.
Landgrebe, Jobst, and Barry Smith. 2025. *Why Machines Will Never Rule the World: Artificial Intelligence without Fear*. 2nd ed. Abingdon, UK: Routledge.
Loeffler, Ronald. 2018. *Brandom*. Cambridge: Polity Press.
Russell, Stuart, and Peter Norvig. 2021. *Artificial Intelligence: A Modern Approach*. 4th ed. Hoboken, NJ: Pearson.
Scharp, Kevin A. 2005. "Scorekeeping in a Defective Language Game." *Pragmatics and Cognition* 13 (1): 203–06. https://doi.org/10.1075/pc.13.1.14sch.
Skorupski, John. 2010. *The Domain of Reasons*. Oxford: Oxford University Press.
Togelius, Julian. 2024. *Artificial General Intelligence*. Cambridge, MA: MIT Press.
Wanderer, Jeremy. 2008. *Robert Brandom*. Montreal: McGill-Queen's University Press.

Netography

ChatGPT's Language Understanding based on: https://chatgpt.com (retrieved July 4, 2025).
Humanoid robot based on: https://en.wikipedia.org/wiki/Humanoid_robot (retrieved July 4, 2025).