



Searching for answers: meaning-making in AI-infused information seeking

Jens-Erik Mai and Sille Obelitz Søre

DOI: <https://doi.org/10.47989/ir31263023>

Abstract

Introduction. This paper investigates the conceptual implications of the shift from searching for documents to searching for answers in AI-supported information-seeking systems. As large language model (LLM) technologies increasingly provide direct responses rather than directing users to sources, foundational questions arise concerning meaning-making, intentionality, and what it means to have a ‘conversation’ with an information system.

Method. We undertake a conceptual analysis grounded in philosophy of language and philosophy of technology, drawing particularly on Wittgenstein’s notions of language games and forms of life, Coeckelbergh’s theory of technology games, and Dennett’s intentional stance. Relevant literature on AI alignment, cognitive authority, and information evaluation is examined to frame the shift toward answer provision.

Analysis. Using these theoretical resources, we analyse the conditions under which humans treat AI systems as intentional or meaningful interlocutors, despite LLMs lacking intentions, understanding, or participation in language games. The analysis clarifies how meaning-semblant outputs may nevertheless be interpreted as meaningful within information practices.

Results. The investigation shows that while users engage LLMs as if they were intentional agents, LLM responses do not constitute meaning-making in the human sense. The shift to answer provision relocates interpretive and evaluative work to users, who must navigate meaning, trust, and cognitive authority without genuine source attribution.

Conclusion. To understand AI-provided answers in information seeking, it is necessary to distinguish between human meaning-making and machine-generated outputs and to reconsider how evaluation, intentionality, and trust operate in answer-oriented information environments.

Introduction

Philosophers and technologists have long wondered: *Can a machine think?* (cf. e.g. Wittgenstein, 1958, §359; Turing, 1950, p. 433; Searle, 1980, p. 422; Floridi, 2015) – and with the introduction of ChatGPT and other LLM-based systems to the broader public audience, that discussion has again ignited. LLM-based chatbots engage in conversations, produce texts, and respond to questions, make comments and jokes, engage in satire, and provide advice in ways not previously experienced from machines. Therefore, today the question is not only: *can a machine think*, but also *can a machine engage in a meaningful conversation*, and more specifically, enter into a question-and-answer session? The fact that LLM-based chatbots generate texts and statements based on prompts will likely ‘change the way we think and experience the writing process and ourselves as writers. In sum, humans and technologies are entangled with one another’ (Coeckelbergh & Gunkel, 2024, p. 2223) and as such LLM-based technologies will change (or have already changed) work and practice in many corners of society, incl. information seeking, analysis, and evaluation.

In a recent paper, Sundin (2025) – using Vickery’s 1961 terminology – demonstrates how LLM-based technologies shift information seeking and retrieval from a focus on seeking ‘information about documents’ to seeking ‘information recorded in documents’ (Sundin, 2025, p. 292); from searching for links to relevant documents to searching directly for answers. Sundin is concerned that LLM-based technologies pose a challenge to information literacy theory and practice, especially as it relates to evaluating the quality of information. The credibility of information has typically been assessed in a combination of *who* says it and *what* is said – people find information credible when they trust the source of the information and when they find the information to be believable (Savolainen, 2023). In today’s media landscape, however, it is less clear who the sources of the information are when a chatbot generates an answer. Although some systems might include links to sources, it remains unclear how those and other sources were used to generate the specific text. Sundin argues that a fundamental question to pose, therefore, is: ‘what happens when AI-infused information systems increasingly provide answers rather than directing people to sources?’ (Sundin, 2025, p. 292).

In this paper, we take a step back to discuss what happens when we have conversations with a machine – in other words, to be able to address Sundin’s fundamental question, we need first to address the more foundational question of: *what happens when we have a conversation with an AI-infused information system in a ‘question and answer’ session?*

The paper outlines the conceptual background needed to understand the emerging shift from document-oriented to answer-oriented information seeking. We begin by introducing Coeckelbergh’s notion of technology games as a framework for analysing how technologies acquire meaning within human practices. Building on this, we turn to the tradition of examining conversations with machines, highlighting how LLM-based systems invite users to adopt an intentional stance despite lacking intentions, understanding, or participation in language games. We then draw on Wittgenstein’s analyses of language, meaning, and forms of life to contrast human meaning-making with the statistical operations of LLMs, and use Searle’s Chinese Room argument to illustrate the limits of rule-following as understanding.

Through this theoretical synthesis, we develop an account of what it means to ‘ask’ and ‘receive’ an ‘answer’ from an AI system. Finally, we apply this conceptual foundation to the contemporary shift identified by Sundin – from searching for documents to seeking system-generated answers – and discuss the implications for information evaluation, cognitive authority, and the nature of knowledge production in AI-infused information practices.

Technology games

As a conceptual framework to understand and analyse technology, we will employ Coeckelbergh's (2018) notion of 'technology games'. While Wittgenstein uses technology as a metaphor to understand language,

Think of the tools in a tool-box: there is a hammer, pliers, a saw, a screw-driver, a rule, a glue-pot, glue, nails and screws. – The functions of words are as diverse as the functions of these objects. (And in both cases there are similarities.) (Wittgenstein, 1958 §11).

Coeckelbergh (2018) invites us 'to turn the metaphor around' and instead explore what it means 'to compare technology to language' (Coeckelbergh, 2018, p. 1510). Coeckelbergh demonstrates how technology can be analysed and explained along the lines of language games embedded in forms of life:

Technologies, considered in their use, are part of activities and games. What may be called 'technology games' shape and give meaning to particular uses. Particular uses of technology are made possible by games and forms of life. They are only meaningful and use-full on the basis of these transcendental conditions, which structure and limit them (Coeckelbergh, 2018, p. 1511).

This approach to understanding and analysing technology understands technology as part of activities and that its meaning and purpose arise from use, just as language is part of activities and meaning arises in use. If we wish to understand one of the tools in the toolbox, such as the hammer, we need to understand the form of life, purpose, and activities that the hammer is part of. The hammer and its meaning are not present outside the specific usages; it enters a specific activity, with a specific purpose in that context, and performs a specific role. The hammer is encountered with specific expectations by the actors in that context; those expectations are formed by their experience and use of other kinds of tools, the saw, screwdriver, or ruler, etc. 'Technologies are part of a form of life, part of what we do and how we do things; they are part of what we are' (Coeckelbergh, 2018, p. 1513).

Therefore, in this sense, both technology games and language games are social games that go together with social and cultural habits and ways of doing things; they both shape and form what is possible and what to expect. 'Such technology games are not to be understood as entirely distinct from language games: the game involves both the use of language and the use of tools. Technology games are also language games' (Coeckelbergh, 2018, p. 1513). It is important to note that technology games are played by humans in their use of technology, just as language games are played by humans in their use of language. Thus, it is not technology in itself, or language in itself, that produces meaning. It is the use of technology by humans and the use of language by humans that produce meaning.

Using technology games as a frame to understand technology, means that technology is to be analysed within the context in which it is embedded. We cannot understand and analyse technology independent of use and activities. This also means that the social, historical and cultural norms and biases of the activity is part of the analysis of technology, 'When we play a particular technology game and engage in a particular activity with technology and with others, we inherit these wider meanings and these historical patterns and 'have' to follow them – this is the normativity' (Coeckelbergh, 2018, p. 1515).

As such, there are two different kinds of meaning: the meaningful uses of technologies as conditioned by the technology games and the meaning of words established in use, and through language games, language games and technology games become intermingled. Meaning in language games is embedded in the technologies and thus, when using technology, meanings surface in different ways. This is often what happens in terms of biases, discrimination, and the

like, which surface and become systematic in the implementation of technological systems in various practices.

Therefore, to critically analyse technologies, that analysis will extend into the larger (infra)structures in which the technology is embedded,

if there are such things as 'values' at all and if they are part of our 'culture', then these values only have normative and semantic significance as part of forms of life and as part of everyday activities and games; that is, they only have meaning and only exist in use; values only exist when they are lived (Coeckelbergh, 2018, p. 1516).

Coeckelbergh's notion of technology games provides a conceptual framework for understanding and analysing technologies through a sociocultural philosophy of technology with a pragmatic understanding of knowledge and meaning-making that are grounded in specific practices and forms of life. In this understanding, technologies in themselves do not create linguistic meaning, but linguistic meanings can be embedded in technology games.

When AI-infused information systems are employed to seek answers to questions, the user of the AI-infused information system engages in technology games and makes sense and meaning of the technology in a particular context. The user is furthermore engaged in a contextual language game to make sense of their information need and evaluate potentially useful answers in their context and specific situation.

Using Coeckelbergh's notion of technology games as a conceptual framework, we will, in this paper, propose a theoretical foundation for understanding how AI-infused information systems provide answers to users' information-seeking activities. The first step is to explore some of the challenges in engaging in conversations with machines.

Conversations with machines

Sundin (2025) finds that AI-infused information systems might lead to a 'de-skilling of searching for and evaluating information' (p. 298), by transferring analytical skills from user to system, and as such, he argues there is a 'transfer of agency' (p. 298), in which the user loses control and becomes more dependent on the system. This leads Sundin to suggest that 'verifying the originator [of the information] becomes meaningless' (p. 298), because, as he argues, 'there is no originator with an intention, at least not in the traditional sense' (p. 298), and then he asks: 'But does it matter?' (p. 298).

In this paper, we will argue, yes, *that matters*. We will argue that this question is, in fact, the central question to address to understand what happens in the shift from searching for documents to searching for answers. Once both human and machine act as if they have a meaningful conversation, it is important to determine what it means to have a meaningful conversation. This becomes even more important if, as Sundin suggests, we rely on AI-infused systems to provide us with knowledge about the world 'beyond the narrow range of my own personal experience' (Sundin, 2025, p. 292; Wilson, 1983, p. 13).

While the idea of having conversations with machines is novel to the information-seeking situation, as Sundin argues, the notion of having conversations with machines has a long history. For instance, the ELIZA program, which was developed in the 1960s, responded to prompts such as 'I am not feeling well today' or 'I hope to...' with empathic, psychologist-like, follow-up questions ('How long have you not felt well?', 'What would it mean to you if...'). Weizenbaum, who developed ELIZA, noted how people 'became emotionally involved with the computer and how unequivocally they anthropomorphized it' (Weizenbaum, 1976, p. 6). More recently, Wittkower (2020) reports a similar experience when Alexa, Amazon's digital assistant, moved into his family's kitchen in the form of an Amazon Echo device. Wittkower notes that the language his family uses

to describe their interactions with Alexa includes that Alexa ‘thinks’, ‘experiences’, ‘listens’, ‘mishears’, ‘understands’, ‘misunderstands’, ‘interprets’, etc., which might lead ‘to some amount of fallacious projection of thoughts and understanding and even personality’ and force us to address ‘the mindedness of technical systems that clearly do not actually have minds’ (Wittkower, 2020, p. 363).

Such experiences have led to the question: *Can a machine think?* The most common and often discussed test to determine whether a machine can think or possess intelligence is the Imitation Game, or *Turing Test*, originally proposed by Alan Turing in his 1950 paper ‘*Computing Machinery and Intelligence*’ (Turing, 1950). The test is designed to assess a machine’s ability to exhibit intelligent behaviour equivalent to, or indistinguishable from, that of a human. There have been many discussions of the interpretation of the test, including whether the test, in fact, tests machine intelligence (cf. e.g. Goncalves, 2023).

Our interpretation of Turing’s paper and the Imitation Game follows Danziger’s (2022) social-technological interpretation of Turing’s proposal. Danziger argues that the Turing Test is not a test of machine intelligence, but a way to understand ‘humans’ attitude toward machinery’ (Danziger, 2022, p. 18) – hence Danziger’s preference for the original depiction as an *Imitation Game*. In other words, the Imitation Game is not about the capabilities of machines to display intelligence, but about humans and their ‘reaction to the existence of sophisticated machines, about our general attitude toward machinery’ (Danziger, 2022, p. 17). Thus, to explore how AI-infused information systems transform or reshape the information-seeking process, how to evaluate the quality of answers provided by an information search chatbot, or how to assess whether intentionality plays a role in this assessment and evaluation, we need to investigate and understand humans’ attitudes and intentional stance towards machines.

For starters, Fisher (2024) has argued that LLMs lack intentions, and therefore ‘cannot be speakers who say things’ (Fisher, 2024, p. 67) and Coeckelbergh & Gunkel (2024) have, along similar lines, argued that LLMs ‘spit out seemingly intelligible content, but their statements not only mean nothing’ (p. 2222); they ‘write without speaking’ (p. 2226). Nonetheless, as Bender, McMillan-Major, Gebru, and Shmitchell (2021) remarked in their much-cited paper on LLMs’ ability to speak and understand, ‘the tendency of human interlocutors to impute meaning where there is none can mislead both NLP researchers and the general public into taking synthetic text as meaningful’ (Bender et al, 2021, p. 611). In other words, conversations may seem meaningful, but they do not create or produce meaning.

These conversations are merely what Titus (2024) calls ‘meaning-semblant’ behaviour (Titus, 2024, p. 2) – i.e., the LLMs conduct operations and produce outputs that look meaningful to the human but are outside the human practice of meaning-making. Thus, while users may engage in what appears or feels like meaningful conversations with chatbots, chatbots in themselves do not play the technology game – they do not take part in social practices, norms, habits, and rituals, and therefore do not produce or generate meaning. It is humans who play the technology games when they engage the chatbots. Likewise, chatbots do not engage in language games; humans do.

However, even if there is no originator with an intention, as Sundin (2025, p. 298) suggests, people do speak to and interact with LLM systems as if they produce meaning. The reason may be that people anthropomorphize the system, as Weizenbaum (1976) noted, or fallaciously project an amount of thoughts and understanding onto the system, as Wittkower (2020) noted. In any case, to better understand the AI-infused information seeking question-and-answer process, we need to understand humans’ ‘reaction to the existence of sophisticated machines’ (Danziger, 2022, p. 17) and, not least, whether they can engage in meaning-making, incl. the ability to make mistakes, errors, and differentiate between true and false.

Intentional systems

To explain mistakes and errors made by LLMs, it has been suggested to say that they ‘hallucinate’. The metaphor is employed in the industry and literature to explain how LLMs produce responses to prompts, regardless of whether the response is factually correct. The system merely sees things that are not there or that others would have seen differently, as IBM explains: ‘AI hallucinations are similar to how humans sometimes see figures in the clouds or faces on the moon’ (IBM, n.d.). Hicks, Humphries, and Slater (2024) argue that hallucination ‘is an inapt metaphor which will misinform the public, policymakers, and other interested parties’ (Hicks et al., 2024, p. 26). They suggest instead that LLMs should be considered bullshitters (in Frankfurt’s (2005) sense of the word) because they are indifferent towards truth, knowledge, and the world. Fisher (2024) disagrees and argues that, given that LLMs lack intentions, they also lack the ability to be indifferent – they can neither be indifferent nor engaged in the world, ‘because they lack the requisite mental states to perform speech acts like assertion, or to express or communicate particular contents’ (Fisher, 2024, p. 67).

Hicks, Humphries, and Slater (2024) base their analysis on the suggestion that ChatGPT is an ‘intentional system’ in Dennett’s sense of that concept – i.e., ‘a thing whose behaviour is predictable by attributing beliefs and desires (and of course rationality) to it’ (Dennett, 1983, p. 345). Users of the system will adopt an intentional stance towards the system to characterize, understand, or predict the system’s ‘behaviour’ – the output. The argument is not that the technologies have actual, human-like intentions or other intentional attitudes such as ‘wants’, ‘desires’, or ‘beliefs’, etc. The suggestion is that it is useful for humans, according to Dennett, to adopt an intentional stance towards the systems and technologies to explain and understand what they do.

Dennett defines intentional systems as,

a system whose behavior can be (at least sometimes) explained and predicted by relying on ascriptions to the system of beliefs and desires (...). We ascribe beliefs and desires to dogs and fish and thereby predict their behavior, we can even use the procedure to predict the behavior of some machines (Dennett, 1976, p. 179).

The idea being that we ascribe intentions to entities, animals, and machines, although we are aware that they do not in fact possess human-like intentions. We do so to make sense of the entities within our forms of life and within our language games.

As a case to illustrate his point, Dennett employs a chess-playing computer,

For instance, it is a good, indeed the only good, strategy to adopt against a good chess-playing computer. By assuming the computer has certain beliefs (or information) and desires (or preference functions) dealing with the chess game in progress, I can calculate – under auspicious circumstances – the computer’s most likely next move, *provided I assume the computer deals rationally with these beliefs and desires* (Dennett, 1976, p. 179. Italics in original).

For Dennett to make sense of the computer – to understand the computer’s responses and abilities to move chess pieces on the chessboard – he needs to think of the computer as behaving rationally according to its beliefs and desires, even if he accepts that it has none. Again, to understand the chess-playing computer, Dennett meets the computer with expectations and language ‘as variations of older patterns’ (Coeckelbergh, 2018, p. 1513) that were there already before Dennett encountered the chess-playing computer.

Therefore, while computers and technological systems do not have intentions, we might approach them as such, with an *intentional stance* and by considering computers as rational and

engaged in the technology/language game (although we, as humans, are the ones playing those games),

The computer is an intentional system in these instances not because it has any particular intrinsic features, and not because it really and truly has beliefs and desires (whatever that would be), but just because it succumbs to a certain stance adopted toward it, namely the Intentional stance, the stance that proceeds by ascribing Intentional predicates under the usual constraints to the computer, the stance that proceeds by considering the computer as a rational practical reasoner (Dennett, 1976, p. 179).

While Dennett uses old-fashioned symbolic AI-based chess-playing computers to illustrate his point, Hicks, Humphries, and Slater (2024) suggest that we might likewise consider LLMs and other connectionist and language-based AI technologies to be intentional systems. In fact, we suggest that we might not be able to consider these systems as anything less than intentional, given the tight connections between language skills and the development of a theory of mind for other entities (Astington & Baird, 2005) as well as the role language plays in entering the community of minds (Nelson, 2005). Nelson explains that,

entering the community of minds is a developmental process made possible through language. It makes a special case for the emergence of the representational function of language that allows children to go beyond their own private thoughts and beliefs to consider the thoughts and beliefs of others (Nelson, 2005, p. 26).

Nelson further argues ‘that language is the most important general function that leads to higher-order cognitive processes, including the processes involved in theory of mind’ (Nelson, 2005, p. 26), and as such, our ability to think of others (systems, animals, humans) as intentional systems is dependent on language and our language competencies. Thus, seemingly good language competences as those displayed by LLMs, such as ChatGPT, might suggest to us the existence of higher-order cognitive processes, although they are not actually present.

Wittkower (2020) suggests that the human disposition to anthropomorphize, coupled with a system design that requires the adoption of the intentional stance towards AI systems, is problematic, as it deflects our attention from the fact that these systems and technologies do not, in fact, have knowledge, desires, beliefs, etc. They are nothing but very sophisticated statistical models capable of producing very convincing text, speech, pictures, video, code, etc., that in some cases can seem indistinguishable from human-produced content but might be nothing more than (or is) utter bullshit. They are not persons in any meaningful sense of the word, as they cannot be held morally accountable for their ‘actions’ (their output) as they lack communicative understanding and consciousness (Dennett, 1976).

Unlike many other non-human artifacts (dolls, rocks, and books), however, chatbots can produce texts and appear to engage in speech acts, and we might therefore mistakenly perceive them as having intentions and the ability to produce meaning and truth-claims. Given chatbots’ lack of intentions and their lack of ability to ‘care about anything’ (Hicks et al., 2024, p. 6), they cannot be said to be able to lie (have the intention to deceive) or engage in truth-claims. As such, chatbots are best understood as displaying or having ‘indifference towards the truth of the utterance’ (Hicks et al., 2024, p. 5). A LLM system such as ChatGPT is ‘not designed to produce true utterances; rather, it is designed to produce text which is indistinguishable from the text produced by humans’ (Hicks et al., 2024, p. 6).

Unfortunately, chatbots have real-life consequences in many ways – because unclear, vague terminology and metaphors are used to explain their capabilities, and they are decontextualized, flattened, and bullshitting technologies. Further, LLM-based chatbots have become so good at producing linguistic and semantically convincing outputs that humans are often misled to believe

that the chatbot is an active player in the language game. These beliefs are further supported by the design of chatty, communicative interfaces that require us to interact with the systems as if they have minds – intentions, beliefs, and knowledge about the world (Søe & Jørgensen, 2026). In Wittkower's (2020) words, the system requires us to think of it as intentional – 'a mind' – even though it is not. We are the only ones playing the language and technology games.

Chatbots' output is best understood as being *meaning-semblant*, not the production of meaning. Thus, the concern lies in our attitude towards technology, and how we approach it, and what we expect of it. Now that we have unpacked the notions of intention and intentionality, we can return to exploring how AI-infused information-seeking systems might produce answers to questions.

To ask questions

When we pose a question, we do so from an intentional stance where we have specific expectations towards the reply we receive – and we pose the question from a perspective of a knower. When we ask a question, we must already know something. It is not possible to ask questions or wonder about the world without some conception of the world, 'One has already to know (or be able to know) something to be capable of asking a thing's name. But what does one have to know?' (Wittgenstein, 1958 §30).

According to Wittgenstein, there are two basic foundational requirements for knowing: i) to doubt and ii) to be part of a practice. Successful communication rests on the premise that humans – through a good amount of conscious and unconscious training – learn to use words in the same way. As we saw earlier in this paper, Wittgenstein compares the use of words to the use of tools. While there are various uses of a hammer, we have agreed on a few correct or acceptable ways a hammer could be used. The functions of words are as diverse as the functions of tools. The meaning of words is settled in what Wittgenstein called 'language games': 'I shall also call the whole, consisting of language and the actions into which it is woven, the 'language game' (Wittgenstein, 1958 §7).

As such, words (and language) do not have meaning in relation to what they refer to or signify; words are assigned their meaning through their use. Obermeier (1983) summarizes the position: 'Words are used for their meaning, and the meaning depends on the use in a given context' (Obermeier, 1983, p. 341). Therefore, the meaning of words changes when the context changes; meaning is tied to activities and the usage of the words in that context, 'When language-games change, then there is a change in concepts, and with the concepts the meaning of words change' (Wittgenstein, 1969 §65).

Wittgenstein names such discourse communities – where words are used and defined – 'forms of life'. To speak a language is part of an activity, and in that sense, what we do besides using language cannot be separated from our use of language. Judgments of right and wrong, true and false, and meaning in general are based in forms of life. The meaning of words, statements, and dialogues is determined within a particular form of life; therefore, judgments of whether a language game is understood correctly take place within a form of life.

Furthermore, to know and to doubt are closely linked; one can only know something about that which one can also be in doubt about. It only makes sense to talk about knowledge in contexts where it is possible to be in doubt; where one can ask: "How do you know?" and the answer to that presupposes that this can be known in that way' (Wittgenstein, 1969 §40). But how do we establish that the truth of a proposition is certain? Wittgenstein invites us back to the foundation for establishing the meaning of words and statements; the language-game: 'Giving grounds, however, justifying the evidence, comes to an end; – but the end is not certain propositions'

striking us immediately as true, i.e. it is not a kind of seeing on our part; it is our *acting*, which lies at the bottom of the language-game' (Wittgenstein, 1969 §204).

Therefore, to know something, to understand the meaning of a word, or to determine the truth-value of a statement is not about deciphering single words or statements, but about the larger contexts the words and statements are part of: 'What I hold fast to is not one proposition but a nest of propositions' (Wittgenstein, 1969, §225). It is the language-game and the form of life. If all propositions in the nest are doubted, the whole worldview tumbles down. It is not that we, in principle, cannot doubt any and all of the propositions; it is just that we simply do not do it – it does not occur to us as a possibility. Wittgenstein (1969) explains this in a series of four paragraphs:

§341. That is to say, the questions that we raise and our doubts depend on the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn.

§342. That is to say, it belongs to the logic of our scientific investigations that certain things are in deed not doubted.

§343. But it isn't that the situation is like this: We just can't investigate everything, and for that reason we are forced to rest content with assumption. If I want the door to turn, the hinges must stay put.

§344. My life consists in my being content to accept many things.

(Wittgenstein, 1969 §§341-344)

Language use and meaning-making is a form of negotiation. You can question some aspects, blindly accept others, and engage in the formation of yet other aspects. You cannot question the entire foundation of a conversation. The form of life provides the conditions of possibility that constrain the language game – the conditions for the rules – and the playing of the game, the use of words can slowly change the game and the form of life. In other words, if you want to be understood when you say something, you must play by the rules of the language game. Then, in playing, you can ever so slowly, in collaboration with the other players, start changing the rules, and thereby the form of life that from the outset conditions the rules and the language game.

As Coeckelbergh (2018) explains, it is analogous to the river and the bedrock – they each condition and constrain each other, but ever so slightly, the river will be able to change the bedrock, thereby changing its cause.

As our language use sometimes cuts new banks, change is possible. But this change may be imperceptible and slow; the banks that guide the river are relatively stable, some of the rocks are hard. When we learn language, there is already a river-bed, there is already a language and a form of life (Coeckelbergh, 2018, p. 1509).

There is a context we need to understand and appreciate in order to understand the language used in that context. At the outset of his *Philosophical Investigations*, Wittgenstein describes what he calls a 'complete primitive language' for communication at a construction site, to consider how language works, and especially how meaning is generated and attached to words used in the language:

The language is meant to serve for communication between a builder A and an assistant B. A is building with building-stones: there are blocks, pillars, slabs and beams. B has to pass the stones, and that in the order in which A needs them. For this purpose they use a language consisting of the words 'block', 'pillar', 'slab', 'beam'. A calls them out;—B brings the stone which he has learnt to bring at such-and-such a call.—Conceive this as a complete primitive language (Wittgenstein, 1958 §2).

According to Wittgenstein, language, meaning-making, and following rules are not merely about attaching a word ('slab') to a particular object ('a specific stone') and numbly following a rule ('fetch the stone called slab').

The scenario – and this is exactly Wittgenstein's point – is not a good description of how human communication, meaning-making, and language work. In the scenario, the underlying rules for acting 'correctly' and in correspondence with the accepted norms are made explicit, formalized, and moved away from the use of language and the practice in which the language is situated. One needs no interpretative flexibility to act in the situation; one merely needs to correctly employ a set of explicit rules. One does not need knowledge about which stones are correct to use when constructing stone buildings; one merely needs to look in the rulebook and apply the rules correctly, but that does not mean that meaning was produced or exchanged or that the actors understood what was going on or why.

Questions are always posed within a specific context and situation, and these factors are crucial in evaluating whether the answers are useful and meaningful. Therefore, to understand the processes of meaning in AI-infused information-seeking, we must first explore what meaning itself entails; what it means to understand something.

To understand something

Wittgenstein notes that 'to imagine a language means to imagine a form of life' (Wittgenstein, 1958, §19); we can only imagine and understand – with emphasis on *understand* – the language that the builder and assistant use, when we can imagine and understand the activities they are engaged in. In Wittgenstein's scenario, there is no understanding of the language used, only the application of rules. If we want understanding, we need the context, the practice, the form of life.

Likewise, we can only understand a chatbot, its language, and the meaning-semblant output if we understand the form of life that the chatbot is part of. Understanding how to use language (and technology) is about understanding that one 'must be able to do certain things' (Wittgenstein, 1969 §534).

Although Wittgenstein considers how humans use language, we can use that framework to discuss how humans use technology (i.e. the aforementioned technology games) – there is also an already-there language, past experiences, and expectations about what to encounter when we use technologies, such as chatbots or AI-infused information-seeking systems. People approach and encounter technology with a set of expectations and from within a specific form of life.

This is the point which Searle demonstrated with his famous thought-experiment called the '*Chinese Room argument*' (Searle, 1980). The gist of the Chinese Room argument is to show that 'understanding' is more than mere application of rules or instructions. In other words, the fact that the assistant, in Wittgenstein's construction site example above, brings the correct stones when called for does not mean that the assistant 'understands' what a stone is or what purpose it has.

In the Chinese Room thought-experiment, Searle imagines he is locked in a room and asked to manipulate Chinese symbols according to a set of instructions he receives in English to answer questions posed in Chinese. He imagines that after a while, he is 'so good at following the instructions' that his answers are 'absolutely indistinguishable from those of native Chinese speakers' (Searle, 1980, p. 418). The question then is whether the Chinese Room 'understands' Chinese. Searle argues that it does not, because he does not understand a single Chinese word. Proponents of Strong AI suggest that the Chinese Room as a system 'understands' Chinese, because 'understanding' has to do with recognizing patterns and system output and not with individual parts (such as Searle) of the system. Proudfoot (2002) agrees with Searle's conclusion but suggests that when Searle acts inside the Chinese Room, he is 'meaning-blind' (Proudfoot,

2002, p. 172) in the sense that he merely manipulates symbols by means of rules and, as such, therefore has no understanding of the questions and answers; in that sense, the system as a whole is meaning-blind.

In Wittgensteinian analyses of Searle's Chinese Room, Proudfoot (2002) has argued that pure symbol-manipulation is insufficient and 'does not constitute understanding' (Proudfoot, 2002, p. 170) and Obermeier (1983) argues that a key difference in Searle's and Wittgenstein's analyses of intentionality, is that it for Searle is a biological phenomenon and for Wittgenstein a linguistic phenomenon. Obermeier further argues that 'projection of a meaning onto an utterance is not the intention itself' (Obermeier, 1983, p. 340), instead 'what is important for 'understanding' is the function of an utterance in a given situation' (p. 341).

There is no doubt that ChatGPT does well in the imitation game, and we might even approach what McLuhan (1964) called 'the final phase of the extensions of man – the technological simulation of consciousness, when the creative process of knowing will be collectively and corporately extended to the whole of human society' (McLuhan, 1964, p. 3-4). As such, we may be at a stage that exceeds what Turing had in mind, when he predicted that: 'Nevertheless I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted' (Turing, 1950, p. 442).

Therefore, to understand meaning-making in the context of today's LLM technologies, we need to consider not just whether answers are as intelligent as those provided by a human. To understand the changes to information-seeking activities, as predicted by Sundin (2025), we need to explore both our 'attitude toward machinery' (Danziger, 2022, p. 18) and how we may project 'thoughts and understanding and even personality' to LLM technologies (Wittkower, 2020, p. 363). This requires a foundational understanding of what we take understanding and meaning-making to be. In this paper, we have argued that meaning-making is not something that takes place outside of language; instead, meaning is best understood as something that is established through the use of language and within the practice in which it is used. It is the notion of meaning as 'its use in the language' – in the playing of language games (Wittgenstein, 1958 §43).

However, humans and machines work with language differently (Pérez-Escobar & Sarikaya, 2024; Søre, 2021). LLMs produce linguistic outputs based on statistics and calculation of probabilities; humans use and develop language when playing language games as part of their practice, their form of life. This difference is the origin of several of the ethical implications we meet when introducing LLMs in our human practices. We take the outputs of LLMs to be meaningful because they are linguistically and semantically convincing, but they are, in fact, only meaning-semlant. LLMs do not play language games – humans do. And humans further play technology games when they interact with LLMs.

In sum, human knowledge, meaning-making, and ethics are tied to particular contexts and practices, to forms of life, whereas chatbots generate texts independent of such contexts and practices. They inhibit human values and biases because these are statistically present in their training data, but they are not able to question, challenge, or endorse specific views, unless specifically trained or hard-coded to do so. To speak about what one knows is to engage in a language game and to participate in a form of life; humans do that, machines cannot.

Conclusion

We meet and encounter LLMs in social-technological contexts and with expectations about the technology based on our experiences with other technologies. We have an already-there expectation about what we will experience, and we make sense of the technology within a

specific form of life. We talk about the technology in an anthropomorphic manner, and we may say that it ‘understands’ our questions and the way we think, that it might itself ‘think’ and ‘have a good sense of humour’.

It is (almost) impossible to distinguish a text produced by ChatGPT from a text produced by a human – it plays the *Imitation Game* very well and acts as if it understands the exchange of thought; it appears to be understanding language and engaging in meaning-making. We might be concerned when it provides wrong answers, perhaps it ‘lies’, ‘deceives’, ‘has bad intentions’, or perhaps it merely ‘hallucinates’. The machine chats and seems to be playing the social game very well. But it does so because it is very good at statistics and predicting the next tokens in a sequence of text – in this sense, statistics ‘is the new magic’ (Coeckelbergh, 2020, p. 94). While it is possible to engage in a seemingly meaningful (or meaning-semblant) conversation with a chatbot, it is not possible for the chatbot to produce meaning and exchange meaningful thoughts and utterances with us. We might mistake machines as speakers, when in fact, there is only text and no intentions; no ability to distinguish true from false. It is all bullshit. While we play language games within a form of life – the chatbot merely assembles texts based on statistical analysis. It is a meaning-semblant conversation, and it might feel meaningful to us, but there is no production of meaning in the ordinary sense. All meaning is ascribed to the conversation by humans after the production of the text by the LLM.

Information science has a long and established tradition for developing theories and practices for seeking information in the form of documents. However, a shift towards seeking ‘information recorded in documents’ (Sundin, 2025, p. 292) requires engagement with questions concerning what it means to have meaningful conversations, what it means that systems are intentional, and how meaning is formed and information is established. It further means that there is a shift in abilities required for humans dealing with information systems – a shift from evaluating the quality of sources to evaluating the quality of answers. Furthermore, it requires evaluating answers when it is unclear where the answer stems from, on what grounds it was generated, and whether hallucinations are part of the answer requires knowledge that otherwise must be taken at face value.

Taken together, these findings prompt key questions about information-seeking practice, literacy, and models, including: how does meaning matter in information-seeking? is meaning-making absent in the information-seeking process? and in which ways does intention play into information-seeking evaluation? How one conceptually approaches and thinks about these questions will define not only how one might implement information-seeking in practice, but also, in the end, will define what kind of practice and scholarship the information-seeking field will become.

In this paper, we have provided a conceptual foundation for information systems that provide answers to questions. When we rely on AI-infused systems to provide us with knowledge about the world ‘beyond the narrow range of my own personal experience’ (Sundin, 2025, p. 292, Wilson, 1983, p. 13), we need to be able to trust the system, to treat the system as a cognitive authority, to use Wilson’s term (Wilson, 1983). That requires that we must first sort out the kind of interaction we have with the system, what we expect from the system, and whether to place meaning-intentions in the answers the system supplies – that is, to figure out which technology game we are playing. In other words, to address Sundin’s question: ‘what happens when AI-infused information systems increasingly provide answers rather than directing people to sources’ (Sundin, 2025, p. 292), we must first ask ourselves: *how is meaning present or absent in the question-and-answer session we have with the system?*

Acknowledgements

Research for this paper has, in part, been supported by the project *Absence of information in decision-making processes* (ABSENCE), generously funded by the *Independent Research Fund Denmark* (grant no. 2097-00015B, PI: Sille Obelitz Sørensen).

About the authors

Jens-Erik Mai is Professor of Information at the Department of Communication, University of Copenhagen. Mai received his Ph.D. from the University of Texas at Austin, and his research interests include fundamental questions about data and information in the digital society. He can be reached at: jemai@hum.ku.dk.

Sille Obelitz Sørensen is Associate Professor in Philosophy of Information and Technology at the Department of Communication, University of Copenhagen. Sørensen received her Ph.D. from the University of Copenhagen, and her research interests include conceptual questions regarding the relations between information, humans, and technologies. She can be reached at: sille.obelitz@hum.ku.dk.

References

- Astington, J. W., & Baird, J. A. (Eds.). (2005). *Why language matters for theory of mind*. Oxford University Press.
- Bender, E. M., McMillan-Major, A., Gebru, T., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FACCT '21)*, 610–623. ACM. <https://doi.org/10.1145/3442188.3445922>
- Coeckelbergh, M. (2017). *Using words and things: Language and philosophy of technology*. Routledge.
- Coeckelbergh, M. (2018). Technology games: Using Wittgenstein for understanding and evaluating technology. *Science and Engineering Ethics*, 24(5), 1503–1519. <https://doi.org/10.1007/s11948-017-9953-8>
- Coeckelbergh, M. (2020). *AI ethics*. MIT Press.
- Coeckelbergh, M., & Gunkel, D. J. (2024). ChatGPT: Deconstructing the debate and moving it forward. *AI & Society*, 39, 2221–2231. <https://doi.org/10.1007/s00146-023-01710-4>
- Danziger, S. (2022). Intelligence as a social concept: A socio-technological interpretation of the Turing test. *Philosophy & Technology*, 35, Article 68. <https://doi.org/10.1007/s13347-022-00561-z>
- Dennett, D. C. (1976). Conditions of personhood. In A. O. Rorty (Ed.), *The identities of persons* (pp. 175–196). University of California Press.
- Dennett, D. C. (1983). Intentional systems in cognitive ethology: The “Panglossian paradigm” defended. *Behavioral and Brain Sciences*, 6(3), 343–390. <https://doi.org/10.1017/S0140525X00016393>
- Fisher, S. A. (2024). Large language models and their big bullshit potential. *Ethics and Information Technology*, 26, Article 67. <https://doi.org/10.1007/s10676-024-09802-5>

- Floridi, L. (2015). Singularitarians, Altheists, and why the problem with artificial intelligence is H.A.L. (Humanity At Large), not HAL. *APA Newsletter on Philosophy and Computers*, 14(2), 8–11.
- Frankfurt, H. (2005). *On bullshit*. Princeton University Press.
- Goncalves, B. (2023). Can machines think? The controversy that led to the Turing test. *AI & Society*, 38, 2499–2509. <https://doi.org/10.1007/s00146-021-01318-6>
- Hicks, M. T., Humphries, J., & Slater, J. (2024). ChatGPT is bullshit. *Ethics and Information Technology*, 26, Article 38. <https://doi.org/10.1007/s10676-024-09775-5>
- IBM. (n.d.). What are AI hallucinations? Retrieved November 2025, from <https://www.ibm.com/topics/ai-hallucinations> (Archived by the Internet Archive at <https://web.archive.org/web/2025101160702/https://www.ibm.com/think/topics/ai-hallucinations>)
- McLuhan, M. (1964). *Understanding media: The extensions of man*. McGraw-Hill.
- Nelson, K. (2005). Language pathways into the community of minds. In J. W. Astington & J. A. Baird (Eds.), *Why language matters for theory of mind*. Oxford University Press.
- Obermeier, K. K. (1983). Wittgenstein on language and artificial intelligence: The Chinese-room thought experiment revisited. *Synthese*, 56(3), 339–349. <https://doi.org/10.1007/BF00485470>
- Pérez-Escobar, R. A., & Sarikaya, D. (2024). Philosophical investigations into AI alignment: A Wittgensteinian framework. *Philosophy & Technology*, 37, Article 80. <https://doi.org/10.1007/s13347-024-00761-9>
- Proudfoot, D. (2002). Wittgenstein's anticipation of the Chinese room. In J. Preston & M. Bishop (Eds.), *Views into the Chinese room: New essays on Searle and artificial intelligence*. Clarendon Press.
- Savolainen, R. (2023). Assessing the credibility of information sources in times of uncertainty: Online debate about Finland's NATO membership. *Journal of Documentation*, 79(7). <https://doi.org/10.1108/JD-08-2022-0172>
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417–424. <https://doi.org/10.1017/S0140525X00005756>
- Sundin, O. (2025). Theorising notions of searching, (re)sources and evaluation in the light of generative AI. *Information Research*, 30(CoLIS), 291–302. <https://doi.org/10.47989/ir30CoLIS52258>
- Søe, S. O. (2021). A story of surveillance? Past, present, prediction. *SATS*, 22(1), 11–22. <https://doi.org/10.1515/sats-2020-0021>
- Søe, S. O., & Jørgensen, R. F. (2026). Engaging with non-minds and hybrid others. Philosophical perspectives on AI and automated decision-making. In H. P. Olsen et al. (Eds.), *Artificial intelligence, humans and the law*. Routledge (Artificial Intelligence, Law and Society).
- Titus, L. M. (2024). Does ChatGPT have semantic understanding? A problem with the statistics-of-occurrence strategy. *Cognitive Systems Research*, 83, 101174. <https://doi.org/10.1016/j.cogsys.2023.101174>

- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433–460.
<https://doi.org/10.1093/mind/LIX.236.433>
- Vickery, B. C. (1961). *On retrieval system theory*. Butterworths.
- Weizenbaum, J. (1976). *Computer power and human reason: From judgment to calculation*. W. H. Freeman.
- Wilson, P. (1983). *Second-hand knowledge: An inquiry into cognitive authority*. Greenwood Press.
- Wittgenstein, L. (1958). *Philosophical investigations* (G. E. M. Anscombe, Trans.; 3rd ed.). Blackwell.
- Wittgenstein, L. (1969). *On certainty* (G. E. M. Anscombe & G. H. von Wright, Eds.; D. Paul & G. E. M. Anscombe, Trans.). Blackwell.

Copyright

Authors contributing to *Information Research* agree to publish their articles under a [Creative Commons CC BY-NC 4.0 license](https://creativecommons.org/licenses/by-nc/4.0/), which gives third parties the right to copy and redistribute the material in any medium or format. It also gives third parties the right to remix, transform and build upon the material for any purpose, except commercial, on the condition that clear acknowledgment is given to the author(s) of the work, that a link to the license is provided and that it is made clear if changes have been made to the work. This must be done in a reasonable manner, and must not imply that the licensor endorses the use of the work by third parties. The author(s) retain copyright to the work. You can also read more at: <https://publicera.kb.se/ir/openaccess>