

Preprint. Forthcoming in: Durt, Christoph. 2025. “The Meaning of Context and Co-Text for Human Understanding and Large Language Models.” Edited by Christoph Durt and Krämer, Sybille. *Philosophy & Digitality*, Special Issue on LLMs and the Patterns of Human Language Use, 2 (1).

The Meaning of Context and Co-text for Human Understanding and Large Language Models

Abstract

Context serves two distinct yet interrelated functions: (1) it provides a framework for interpreting symbolic expressions, and (2) it forms the core of Large Language Model (LLM) computation of numerical relationships between tokens. Each function relies on different features of context, and the widespread failure to distinguish them has given rise to confusion concerning the ability of LLMs to “understand” context and meaning. The paper distinguishes two kinds of context: (1) a broad sense, including the world we experience and live in, and (2) numerical relationships to other text parts. To clearly demarcate the two senses, the concept of “co-text” will be used for the second. LLMs transform co-text to produce text that is meaningful to humans, but this does not mean that LLMs understand meaning. Understanding the meaning of text requires embedding it in the broader context of human language use. Since LLMs do not do that by themselves, the correct question about LLM understanding is not whether they can understand context, but to what extent computations of co-text can compensate for missing context. The paper concludes with an outline of an answer: LLMs can compensate to a large extent for missing context because the patterns they derive from human language use constitute co-texts that are intertwined with the context of sense-making.

Keywords: AI, context, co-text, Large Language Models (LLMs), meaning, modeling, understanding

1. Introduction

How can Large Language Models (LLMs) process and produce meaningful text—that is, text that serves as a sensible and often informative response to a prompt? This question often gets answered by alleging that LLMs “understand meaning” (Manning 2022) or do something that “amounts to understanding” (Agüera y Arcas 2022). However, such claims raise questions about the nature of understanding, meaning, and simulation of understanding. Others argue that LLMs are merely “stochastic parrots” that repeat text fragments without understanding their meaning (Bender et al. 2021). Yet this explanation alone fails to clarify how the repetition of text fragments can generate output that meaningfully responds to a prompt. The question of how LLMs process and produce meaningful text extends beyond the technical architecture of these models to encompass fundamental issues: the nature of digital states, computation, language use, text, and meaning comprehension, along with their interconnections—topics that have been thoroughly investigated by numerous philosophers of the last millennia.

Context plays a crucial role in both LLM computation and human understanding. On the one hand, numerical relationships to other words form part of a word’s context. On the other hand, context constitutes a frame of meaning; the meaning of any expression depends on the context in which it is used. But are these the same types of “context” in both cases? To avoid conflating different things and processes, this paper distinguishes two different types of context: context in the broad sense and a narrow subset of context, which is called co-text.

The claim that LLMs understand language overlooks that even difficult written tests take place in limited contexts. Unlike much intelligent human activity, the context of LLMs is text—even when the text is a video file. Like in the Turing Test, the point is to transform input text into output text. The mere capability to transform text into text is not sufficient for understanding; text can only be understood in relation to its use in meaningful contexts. The claim that LLMs understand language rests on the confusion of a part of context with the whole of context. The difference between context and co-text explains how LLMs produce text that is meaningful to humans: by transforming co-text that humans can understand in meaningful contexts.

2. Context and Co-text

‘Context’ stems from Latin *contexere* “to weave together,” from *con* (with) and *texere* (to weave, to make). *Texere* is also the root of ‘text’ and ‘technology,’ which hints at fundamental semantic similarities. All three concepts describe the activity of assembling artifacts to create something new. The result displays regular patterns that form a structure—one that simultaneously separates us from the world and mediates between ourselves and the world. Context, text, and technology all trace back to the Indo-European **tetk-* (to weave, join, fit together, braid, interweave, construct, fabricate, build). Text creates a pattern of expressions interwoven with context. ‘Context,’ in its broadest sense, determines the meaning of text and other symbolic expressions. Context is not a side issue that may or may not be considered in addition to the main question of textual meaning. Rather, context is indispensable for meaning and sense-making.

Context proves indispensable because words and phrases derive their meaning from the context in which they are used. The most obvious examples of the context-sensitivity of meaning are synonyms: ‘bat’ can mean a type of sports equipment or an animal. Even seemingly unambiguous statements like “This is a very insightful paper” can, depending on its context, express genuine praise or cutting sarcasm. The reason why any expression can shift its meaning when interpreted in different contexts is straightforward: words and symbols carry no inherent meaning—they only gain meaning through use in specific contexts. In other words: “Without context, words and actions have no meaning at all” (Bateson 2002, 14). Without context, there is no meaning and hence no understanding of text.

Context in a broad sense—including behavior and imagination (see section 4)—exhibits differences that make a difference to meaning. However, ‘context’ can also be understood much more narrowly, referring to the distribution of words across a text corpus. Since these distributional relationships can be accounted for with multidimensional vectors, and these are at the core of LLM computation, LLMs compute contextual relationships of this narrow kind. In this narrow sense, LLMs might be said to “learn” context, but that doesn’t necessarily mean they do so the same way humans do. Nor does it imply that LLMs learn other kinds of context, or that they “understand” context analogously to humans. Such suggestions overlook crucial differences.

To disentangle the confusion, two meanings of context must be clearly distinguished. LLMs operate with context in a narrow sense, namely numerical relationships between tokens. Humans, particularly specialized linguists, sometimes use statistical context in similar ways, but doing so

without computers would be a tedious exercise of questionable value. Normally, humans approach text in a very different way: by relating it to a much broader context. To differentiate between the narrow contexts that LLMs operate with and the broader contexts that humans can become aware of, this paper distinguishes between two meanings of context. This distinction will throw light on the question of how far calculations of narrow context can substitute for understanding in a broader context.

The distinction between the broad and narrow senses of context will be marked with a technical term for the narrow subset: “co-text.” Many authors writing about context do not differentiate between context and co-text, but those who do typically define co-text as the surrounding text, and context as the relevant co-text combined with other pertinent features. In the linguistic literature, the context of an utterance includes “not only the relevant co-text (i.e., the relevant surrounding text) but also the relevant features of the situation of utterance” (Lyons 1995, 271). Accordingly, in this paper, ‘co-text’ simply denotes the text surrounding a word, token, or symbol within a text or text corpus.

Co-text forms part of context, but it differs fundamentally from other parts of context. Since LLMs by themselves only transform symbols into other symbols, they relate symbols only to co-text and not to the broader context. The text corpus they train on includes enormous amounts of texts that have been used in meaningful contexts, such as a dictionary entry on bats, reflecting real contextual language use. LLMs do not model language but language use.

More specifically, current LLMs model written language use: LLMs have direct access only to written text, even when it represents graphic, audio, or video files. This is not the case for humans: text and co-text include non-written signs and expressions, and these are embedded in a broader context, whereas LLMs map relationships between words or tokens in their training data. When the model’s weights adjust through repeated training iterations, more general patterns are abstracted from concrete co-texts and weights are induced that befit numerous co-texts. Since all these operations consist of transformations of co-text, the concept of co-text is key to understanding LLMs.

The complex weights are co-textual stochastic patterns consisting of repetitive numerical regularities. The stochastic patterns derive from the use of writing across numerous contexts and hence mirror frequent patterns of language use. The patterns of relationships derived from training data are then refined in further training to produce the correct output for different tasks.

Through training, the weights become, on the one hand, more “average”: they abstract from individual uses, leading to more common word choices and expressions. On the other hand, they become higher-dimensional, enabling the generation of output that precisely fits the co-text of input text. Mapping patterns in numerous contexts of language use to enormously complex stochastic patterns creates an LLM that predicts frequent text continuations precisely adjusted to the co-text of the pre-processed input prompt. Co-text isn’t just pivotal for LLM processing; co-text is fundamentally all they have. At the core of the foundational model (Liang et al. 2022), LLMs only have access to the relationships between words in input data and process them with weights determined from the context of their training data.

The question remains whether co-text alone is enough to understand at least some meaning. Precisely this is the claim of a semantic theory that long predates LLMs, as well as all but the most rudimentary forms of Deep Learning technology. Distributional Semantics argues that word usage relationships reveal meanings—“You shall know a word by the company it keeps!” (Firth 1957, 11). Co-text consists of the distribution of words or tokens in a text corpus, from which LLMs extract stochastic patterns, which they then use to transform input co-text (prompts). Since co-text refers to the co-occurrence of words or tokens in texts, it closely relates to collocations, defined as

“actual words in habitual company” (Firth, 1957, p. 9). The concept of co-text denotes what Distributional Semantics considers essential for meaning: the relationships between words in text.

This approach “is based on the hypothesis that the meaning of a linguistic expression can be induced from the contexts in which it is used” (Boleda and Herbelot 2016, 623). By “context,” Boleda and Herbelot mean co-text. Distributional Semantics reduces “context” to co-text and asserts that “context” in the sense of co-text alone suffices to determine linguistic meaning. Like the later Wittgenstein, Distributional Semantics focuses on language use. However, since it only examines word distributions and disregards pragmatics and other aspects of language use beyond co-text, it represents a “minimal version of this theory of use in linguistics” (Krämer 2025). The later Wittgenstein, by contrast, means by “use” much more than just distributions in a text corpus.

Historically, Distributional Semantics lacked the computational power to process large amounts of language use. Today, however, LLMs show that calculations on distributions in a text corpus can produce both grammatical and meaningful text. By utilizing enormous amounts of stochastic relationships between text items—i.e., co-text—they generate text that would require contextual understanding if humans produced it. However, whether the distributional co-text that LLMs use is sufficient for understanding meaning remains to be determined. If the hypothesis that co-text suffices to understand at least some meaning proves true, then LLMs understand at least some meaning when they correctly compute co-text. Let’s examine one of the most compelling descriptions of this claim.

3. Learning Meaning from Co-text

An intuitive example that is supposed to support the claim that LLMs understand meaning by learning co-text is the learning of the meaning of the word *shehnai*:

“[I]f I have held an Indian *shehnai*, then I have a reasonable idea of the meaning of the word, but I would have a richer meaning if I had also heard one being played. Going in the other direction, if I have never seen, felt, or heard a *shehnai*, but someone tells me that *it’s like a traditional Indian oboe*, then the word has some meaning for me: it has connections to India, to wind instruments that use reeds, and to playing music” (Manning 2022, 135, italics in the original).

Humans can learn word meanings in multiple ways, including from information about how a word relates to other familiar words. We can even learn something about the meaning of *shehnai* without an explicit definition simply by seeing the word used in a story. In either case, we never gain complete knowledge of a word’s use—our understanding of meaning remains always partial. No single person can know everything about the use and reference of the word *shehnai*. Learning a word’s or expression’s meaning isn’t a matter of all or nothing. One can (and typically does) only partially understand its meaning and yet still use the word correctly in at least some contexts.

What does this mean for LLMs? Manning believes that the ability of humans to learn part of a word’s meaning by learning its relation to other words justifies the claim that semantic meaning consists of understanding networks of connections between linguistic forms. He concludes that LLMs also learn meaning:

“Using this definition whereby understanding meaning consists of understanding networks of connections of linguistic forms, there can be no doubt that pretrained language models learn meanings. As well as word meanings, they learn much about the world.” (Manning 2022, 135)

Through training on co-text, LLMs can produce a definition of *shehnai* and, in this weak sense, “learn meanings” and “much about the world.” This sense of learning, however, merely involves acquiring linguistic forms that someone who understands their meaning and reference can interpret. LLMs undoubtedly “learn meanings” in this limited sense. But Manning argues for a much stronger

claim: that machines use writing as a knowledge store “just like people” (ibid.). This could simply mean that machines, like people, learn from text—LLMs do this by adjusting their output based on textual input. Or it might mean something more substantial: that machines learn just like people do (by understanding textual relationships as meaningful). Manning’s longer quote resolves this ambiguity. Since he believes LLMs learn meaning by understanding networks of connections between linguistic forms, he concludes that LLMs, just like people, understand these networks as semantically meaningful.

However, Manning’s argument contains a fundamental flaw. Even though it is reasonable to say that some human understanding of meaning consists of understanding networks of connections, it’s not self-evident that LLMs *understand* networks of connections. This would presuppose that “understanding” linguistic forms consists merely of the ability to operate with them in ways that make sense to humans. If human learning of meaning from linguistic forms involves more than just computing form, then the fact that humans can learn semantics from linguistic forms doesn’t prove that machines also acquire semantic knowledge from linguistic form. If humans learn meaning from linguistic forms by understanding their context rather than calculating co-text, then the fact that humans learn from co-text does not warrant the conclusion that LLMs can do the same. To address this issue, let’s consider two distinct questions: (1) Would a human with access to nothing but word relationships truly learn their meaning? (2) What enables humans to gain semantic knowledge from word relationships, and could LLMs possibly do the same?

Regarding (1), we can imagine a human with access to only linguistic forms and ask whether that suffices to learn meaning. Like an LLM, the person would learn nothing but numerical relations between numerical tokens. While this would involve tedious work and might prove practically impossible, here only the theoretical possibility matters. The situation would resemble that of the human in the “Chinese Room” thought experiment (Searle 1993), who receives nothing but symbols that are meaningless to him. As in Searle’s thought experiment, the key question becomes whether the mechanical transformations of symbols lead to any true understanding of their meaning. Access to only the symbols and transformation rules alone proves insufficient for understanding their meaning. To reasonably argue that understanding exists in these numerical and symbolic relationships, one must look beyond the narrow context of symbols alone. One would need to consider their meaning in the context of the broader system (Haugeland 2003) and, ultimately, the entire language in which the symbols carry meaning, not just the co-textual relationships. Humans only understand meaning when they view symbols within this broader context.

Someone without knowledge of Chinese gains little understanding from “reading” a Chinese book, even if she knows all the syntactic relationships between the characters. Likewise, a person who knows about ‘shehnai’ merely through its relationships to other words or tokens, without knowing what any of these words or tokens mean, would not be in a position to understand the meaning of “shehnai.” Analogously, a machine that only processes numbers cannot learn the meaning of what those numbers might represent under some interpretation. Such interpretation requires more than converting numbers into other numbers or symbols into other symbols.

The fact that we can make sense of LLM output does not imply that the LLM itself understands the meaning of any of the numbers it operates with. It simply processes states, and the speakers of the language interpret these states as having a certain meaning. Since users readily interpret LLM output, it easily appears as if the LLM would understand its meaning. This appearance is amplified by the fact that LLMs model an important aspect of interpretation. Interpretation means seeing—putting the *interpretandum*, the text needing interpretation, in a specific context. Since co-text is an important part of context and LLMs model co-text, they can seemingly “understand” the relevant context. In this sense (but not others), they can deal with relevance.

Relevance has long been flagged as a problem for Symbolic AI (Dreyfus 1992). LLMs, in contrast, appear to handle relevance fine since they use statistical relationships that frequently align with how people understand context. However, LLMs face a crucial limitation: the numerical relations they operate with are neither meaningful nor relevant by themselves nor through their relation to other co-text. Numerical patterns become meaningful and relevant only when they are interpreted within a meaningful broader context.

Manning's argument relies merely on the claim that LLMs learn and understand *some* meaning. But even that holds true only for those who can already understand meaning, and not for people or things with access to form alone. LLMs learn relationships and patterns between tokens taken from written language use and use them to generate output that those who understand meaning can interpret and understand. Section 5 elaborates this point and argues with regard to question (2) that interpretation involves understanding in a context beyond mere co-text. LLMs only handle co-text and therefore do not understand meaning, though this does not mean LLMs merely parrot human language use.

4. The Missing Context

Space permits only an outline of an answer to question (2): What allows humans to gain semantic knowledge from word relationships, and do LLMs have the same access? If, as discussed in the last section, one cannot learn a word's meaning from a definition or story if one does not already know the language in which it's expressed, what does that knowledge involve? As language speakers, we build upon the meaning, composition, and context of familiar words to understand unfamiliar ones. Even a baby who knows no words may already understand some context that makes words meaningful—including others' behavior and shared situational interactions. Things and interactions are already meaningful to us before we learn words, and this understanding is a prerequisite for learning linguistic meaning.

The pre-verbal use of signs for communication symbolic language can build on is sometimes also called a language. Thomas Reid distinguished between “artificial” and “natural language” (Reid [1764] 1997, 50–53), and the distinction between “cultural” and “natural language” has recently been applied to LLMs (Stuart 2024). Speaking of two kinds of languages, one of which is non-conventional and non-verbal, can be a bit confusing. The pivotal point can be expressed clearer with the concept of context: context beyond co-text is necessary to make a system of signs and rules into a language. Using the above example, we can learn some of *shehnai*'s meaning from a dictionary only because we already understand at least some of the words in its dictionary definition. We gain semantic information from co-text precisely because we comprehend its context. Learning its meaning by hearing, seeing, or feeling the instrument presupposes again understanding some context, such as the instrument taking part in a musical performance.

In the linguistic subfield of pragmatics, the contributions of the broader context to meaning have been discussed since Austin (Austin 1962) and Grice (Grice 1995), typically with reference to speech acts. Because language is embedded in meaningful communicative interactions, Wittgenstein introduced the concept of language games—a concept he developed from his earlier notion of calculus, which he later reconceived as a specific type of language game (Durt 2018). His account developed from language use as a calculus and a reference theory of meaning to a recognition of a plurality of uses in contexts that include forms of life and the world (Wittgenstein [1953] 2009). The later Wittgenstein abandoned his earlier belief that all language use can appropriately be described by modeling the world through representational co-text. While communicative interactions may involve exchanging text, even then there exists not just co-text but a broader communicative context in which the exchange happens, such as receiving a message from a friend.

The context of written text also includes awareness of the situation in which a text was written, its purpose, and its position within broader discourse. The author’s intentions play a role in interpretation—whether this is justified or not. Despite the long-proclaimed “death of the author” (Barthes [1967] 2020) and despite—or because—the role of the author might have been overly reified in the Western tradition (Gunkel 2025), people still consider authorial intent when interpreting textual meaning.

People may furthermore wonder what a text can tell them in the context of our own interests and knowledge, and people tend to understand expressions in the context of their personal experience and their culture. Our emotions and moods may further shape interpretive context. Society and even physical laws provide additional context for language games (Wittgenstein [1953] 2009). Worldviews, religion, analogies, and models understood as meaningful representations offer yet more contextual layers. Context in all these senses is interwoven with co-textual forms of language use but encompasses far more than just co-text.

Context allows humans to learn meaning from very little language use. Unlike LLMs, which calculate relationships between tokens using multidimensional vectors derived from massive text corpora, humans often grasp word meanings from just a single example. When a child learns the word “cat,” the child already has a rough idea of what a cat is. Humans recognize examples as instances of concepts or things that already hold meaning for us in the context of the world we live and communicate in. The furry animal already exists in the child’s world, and the child can likely see its features as meaningful variations within the categories of thing, living being, and animal. Moreover, the word is part of a communicative act the child may have already learned in a different context. Verbal communication builds upon non-verbal communication that predates verbal communication.

Symbolic expressions carry meaning because they serve functions within the broader context of communication, language, and behavior. They occupy a place in the world we communicate and live in. The world already holds meaning for us before we use language. Language doesn’t create meaning *ex nihilo*, but further sharpens and refines existing meaning, enabling richer forms of communication, interaction, and sense-making.

Whereas, for humans, co-text is part of a broader context, LLMs lack direct access to any of the aspects of context described above, apart from co-text. Even if we added sensors and motors to an LLM, all the LLM would get from these is more co-text: data that relates to other data. It makes no fundamental difference for the LLM whether it receives sensor data or synthetic data: for the LLM, it is all just co-textual input. The only context LLMs operate with is co-text. For LLMs, there is nothing outside of co-text.

Because LLMs lack direct access to context beyond co-text, they do not understand language. They also cannot understand co-text itself, since that again would require access to context. LLMs simply process co-text in ways that humans can understand within the corresponding context. Whether we consider context in its broader or narrow sense, calling LLM processing “understanding” is misleading. The right question is not whether LLMs understand context, but rather: to what extent can co-text processing compensate for a lack of contextual understanding when generating meaningful and appropriate language?

5. Compensating for Missing Context

Context can be explicitly described in text, and the description can be modeled in text and co-text. Yet, trying to linguistically account for all context would prove hopeless. Context is too extensive

and differs too fundamentally from text and co-text—any description would inevitably fall short. Nevertheless, operations on co-text can substitute broader contextual understanding to a much larger degree than most would have imagined before the transformer architecture was applied to massive text corpora. So how can operations on co-text compensate for missing context?

A simple answer is to claim that we humans supply the necessary context when we prompt an LLM or interpret its output. Understanding the meaning of the output is a critical human contribution—without it, interactions with LLMs would serve no purpose. While interpreting LLM-generated text is natural and necessary, a philosophical problem emerges when humans project their own understanding onto the machine. Researchers have long documented the human propensity to project understanding onto computers (Weizenbaum 1966).

If understanding is merely projected into computers, LLMs could be “stochastic parrots” that repeat text fragments without understanding their meaning (Bender et al. 2021). This view aligns with the previous section’s conclusion that LLMs lack genuine understanding. However, LLMs do far more than simply repeat text fragments. They extract patterns of language use, which form the foundation of their ability to generate meaningful text. LLMs use stochastic patterns not only to parrot existing text, but also to recombine tokens in novel ways that are meaningful to humans. These complex recombinations of patterns in human language use mirror meaningful semantic relationships that go well beyond mere parroting. While LLMs do not understand text, reducing their functioning to parroting oversimplifies their actual capabilities.

The challenge in recognizing the power of stochastic pattern recombination runs deep. The most common reason people overlook the importance of these patterns stems from an overly simplistic view of how language relates to the world. This becomes clearest when examining Denotational Semantics, probably the most widespread semantic theory and one commonly held by those who claim LLMs merely parrot human language. Unlike Distributional Semantics (sections 2–3), Denotational Semantics maintains that a word’s meaning is contingent on a specific thing beyond text and co-textual relationships: its reference to things or states in the world. Since LLMs arguably lack any intentional relationship to real-world things or states, Denotational Semantics concludes they cannot understand language (Bender and Koller 2020; cf. also Søgaard 2023).

Denotational Semantics offers a straightforward answer to the “symbol grounding problem” (Harnad 1990)—the puzzle of how inherently meaningless symbols acquire meaning: symbols are grounded in the world because they represent things or states of affairs in the world. If LLMs lack an intentional relationship with the world and hence cannot grasp what symbols actually mean. The core insight of the symbol grounding problem is both important and correct: symbols must be embedded in the world to carry meaning. However, since Denotational Semantics assumes that reference to a world independent of language and sense-making gives text its meaning, it struggles to explain how LLMs can transform text in ways that are meaningful to humans—going beyond mere parroting of text chunks. In its canonic form, Denotational Semantics makes two critical oversimplifications that block our understanding of LLMs: it assumes (1) that grounding involves only one type of relationship to the world—denotation or representation—and (2) that language refers to something in the world that exists independently of language and sense-making.

Oversimplification (1) overlooks the rich variety of language games, creating the same problems that plagued the concept of calculus of the early Wittgenstein, and which he eventually abandoned in favor of that of language game (see section 4). For LLMs, this narrow view fails to recognize that some complex transformations make sense because they model aspects of language games that interweave with the world in ways that go far beyond mere denotation of entities. This oversimplification also leads some to mistakenly see intentionality as either an insurmountable barrier or a problem that could be solved through “grounding” simply by adding sensors and motors

to an LLM. This limited perspective blinds us to how intricately meaning is woven into the diverse patterns of language use.

Oversimplification (2) misrepresents the relationship between text and the world. While text and co-text often refer to something beyond themselves—sometimes to things in the world—this doesn't mean they refer to entities existing outside the context of language. The world we speak and write about is already shaped by language itself. In this sense, Derrida correctly asserts that “*There is nothing outside of the text* [there is no outside-text; *il n'y a pas de hors-texte*]” (Derrida [1967] 2016, 172). By “text,” Derrida doesn't simply mean text or co-text but context, as he later clarifies.¹ The point isn't that only text and co-text exist, independent of context and the world. Rather, text is deeply embedded in context, and the shared world we communicate in forms part of context. Our language use in the world provides the context in which meaning becomes possible. Language doesn't refer to context-independent meaning outside itself but to entities, events, and relationships that gain meaning within the context of language.

The world exists and carries meaning before we use language. But as meaningful, it does not exist independently of our sense-making activity. Once we express ourselves in language, we can build upon and refine existing contexts while creating new contexts shaped by our very expression. Context doesn't exist separately from text—it connects necessarily to text and other forms of language use. The patterns in co-text intertwine deeply with patterns of language use. Context grounds text and co-text by offering a horizon of shared sense-making within common language practices.

6. Conclusion

The paper's application of the distinction between context and co-text to LLMs opens a perspective beyond the dichotomy between mere parroting of language and human-like understanding. LLMs produce text that proves meaningful to humans by recombining co-textual patterns they extract from human language use. While co-text does not encompass the broader context that is vital to language understanding, it remains thoroughly intertwined with context. This means complex transformations of co-text can compensate for much missing contextual understanding, especially when humans interpret the output.

The interrelationship between co-text and the broader context explains both the advantages and limits of the typical accounts of LLM understanding or the lack thereof. Traditional semantic theories highlight important aspects of meaning: denotational semantics shows that language needs embedding in the world to be meaningful, while distributional semantics reveals that meaning emerges through use. Yet this paper has shown that LLMs challenge us to move beyond denotational and distributional semantics. Language use doesn't merely represent the world but is thoroughly interwoven with it. Co-text is pivotal for LLMs, but it needs to be situated within the broader context of meaningful language use. LLMs leverage co-text and thereby demonstrate its important role in sense-making, along with possibilities for compensating for missing context.

¹ “The phrase which for some has become a sort of slogan, in general so badly understood, of deconstruction (“there is nothing outside the text” [*il n'y a pas de hors-texte*]) means nothing else: there is nothing outside context. In this form, which says exactly the same thing, the formula would doubtless have been less shocking.” (Derrida 1988, 136)

7. Literature

- Agüera y Arcas, Blaise. 2022. “Do Large Language Models Understand Us?” *Daedalus* 151 (2): 183–97. https://doi.org/10.1162/daed_a_01909.
- Austin, J. L. 1962. *How To Do Things With Words*.
- Barthes, Roland. (1967) 2020. “The Death of the Author.” *Aspen No. 5+6, Item 3: Three Essays*, June. <https://web.archive.org/web/20200604173422/http://www.ubu.com/aspen/aspen5and6/threeEssays.html#barthes>.
- Bateson, Gregory. 2002. *Mind and Nature: A Necessary Unity*. Advances in Systems Theory, Complexity, and the Human Sciences. Creskill, New Jersey: Hampton Press, Inc.
- Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–23. FAccT '21. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445922>.
- Bender, Emily M., and Alexander Koller. 2020. “Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data.” In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5185–98. Online: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.463>.
- Boleda, Gemma, and Aurélie Herbelot. 2016. “Formal Distributional Semantics: Introduction to the Special Issue.” *Computational Linguistics* 42 (4): 619–35. https://doi.org/10.1162/COLI_a_00261.
- Derrida, Jacques. 1988. *Limited Inc*. Evanston, IL: Northwestern University Press.
- . (1967) 2016. *Of Grammatology*. Translated by Gayatri Chakravorty Spivak. Fortieth-Anniversary Edition. Baltimore: Johns Hopkins University Press.
- Dreyfus, Hubert L. 1992. *What Computers Still Can't Do: A Critique of Artificial Reason*. Cambridge, Mass: MIT Press.
- Durt, Christoph. 2018. “From Calculus to Language Game: The Challenge of Cognitive Technology.” *Techné: Research in Philosophy and Technology* 22 (3): 425–46. <https://doi.org/10.5840/techne2018122091>.
- Firth, John Rupert. 1957. “A Synopsis of Linguistic Theory 1930–55.” *Studies in Linguistic Analysis (Special Volume of the Philological Society)* 1952–59:1–32.
- Grice, Herbert Paul. 1995. *Studies in the Way of Words*. 4. print. Cambridge, Mass.: Harvard Univ. Pr.
- Gunkel, David J. 2025. “Does Writing Have a Future?” *Philosophy & Digitality*, no. LLMs and the Patterns of Human Language Use.
- Harnad, Stevan. 1990. “The Symbol Grounding Problem.” *Physica D* (41): 335–46.
- Haugeland, John. 2003. “Syntax, Semantics, Physics.” In *Views Into the Chinese Room: New Essays on Searle and Artificial Intelligence*, edited by John M. Preston and Michael A. Bishop, 379–92. Oxford University Press.
- Krämer, Sybille. 2025. “How Should the Generative Power of LLMs Be Interpreted? Do Chatbots Based on Large Language Models (LLMs) Understand Linguistic Meaning?” *Philosophy & Digitality*, no. LLMs and the Patterns of Human Language Use.
- Liang, Percy, Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, et al. 2022. “On the Opportunities and Risks of Foundation Models.” arXiv. <http://arxiv.org/abs/2108.07258>.
- Manning, Christopher D. 2022. “Human Language Understanding & Reasoning.” *Daedalus* 151 (2): 127–38. https://doi.org/10.1162/daed_a_01905.

- Reid, Thomas. (1764) 1997. *An Inquiry into the Human Mind on the Principles of Common Sense*. Edited by Derek R. Brookes. Critical ed. The Edinburgh Edition of Thomas Reid / General Ed.: Knud Haakonssen 2. University Park, Pa: Pennsylvania State Univ. Press.
- Searle, John R. 1993. "The Problem of Consciousness." *Consciousness and Cognition* 2 (4): 310–19. <https://doi.org/10.1006/ccog.1993.1026>.
- Søgaard, Anders. 2023. "Grounding the Vector Space of an Octopus: Word Meaning from Raw Text." *Minds and Machines* 33 (1): 33–54. <https://doi.org/10.1007/s11023-023-09622-4>.
- Stuart, Susan Aj. 2024. "Why Language Clouds Our Ascription of Understanding, Intention and Consciousness." *Phenomenology and the Cognitive Sciences*, March. <https://doi.org/10.1007/s11097-024-09970-1>.
- Weizenbaum, Joseph. 1966. "ELIZA—a Computer Program for the Study of Natural Language Communication between Man and Machine." *Communications of the ACM* 9 (1): 36–45. <https://doi.org/10.1145/365153.365168>.
- Wittgenstein, Ludwig. (1953) 2009. *Philosophische Untersuchungen =: Philosophical investigations*. Translated by G. E. M. Anscombe, P. M. S. Hacker, and Joachim Schulte. Rev. 4th ed. Chichester, West Sussex, U.K. ; Malden, MA: Wiley-Blackwell.