

ORIGINAL ARTICLE OPEN ACCESS

Conversational Alignment With Artificial Intelligence in Context

Rachel Katharine Sterken¹ | James Ravi Kirkpatrick^{2,3}¹The University of Hong Kong, Hong Kong, China | ²University of Oxford, Oxford, UK | ³Magdalen College, Oxford, UK**Correspondence:** Rachel Katharine Sterken (sterkenr@hku.hk) | James Ravi Kirkpatrick (james.kirkpatrick@philosophy.ox.ac.uk)**Funding:** The authors received no specific funding for this work.**Keywords:** artificial intelligence (AI) agents | context collapse | conversational agents | conversational alignment | human–AI alignment | large language models | pragmatics

ABSTRACT

The development of sophisticated artificial intelligence (AI) conversational agents based on large language models raises important questions about the relationship between human norms, values, and practices and AI design and performance. This article explores what it means for AI agents to be *conversationally aligned* to human communicative norms and practices for handling context and common ground and proposes a new framework for evaluating developers' design choices. We begin by drawing on the philosophical and linguistic literature on conversational pragmatics to motivate a set of desiderata, which we call the CONTEXT-ALIGN framework, for conversational alignment with human communicative practices. We then suggest that current large language model (LLM) architectures, constraints, and affordances may impose fundamental limitations on achieving full conversational alignment.

1 | Introduction

In recent years, advances in deep learning have driven a transformative shift in artificial intelligence (AI), leading to the development of sophisticated conversational agents based on large language models (LLMs), such as OpenAI's ChatGPT, Anthropic's Claude, Google's Gemini, Meta's Llama, and DeepSeek. These AI systems are capable of text-generation and other language-based tasks with unprecedented fluency and proficiency, leading some theorists to claim that sophisticated LLMs like GPT-4 have attained "a form of *general intelligence*, indeed showing *sparks of artificial general intelligence*" (Bubeck et al. 2023, 92). Indeed, LLMs have demonstrated remarkable success across a wide range of domain-specific and domain-general reasoning tasks that extend beyond mere regurgitation of training data, such as translation (L. Wang et al. 2023), sentiment analysis (Kheiri and Karimi 2023), question answering (OpenAI 2023), mathematical reasoning (Lewkowycz et al. 2022), summarization (L. Wang et al. 2023), and code generation (Savelka

et al. 2023).¹ These systems are already being integrated into other technologies, such as search engines, customer service platforms, research and educational tools, personal assistants, AI companions, AI doctors, AI agents, and multi-agent systems, marking a fundamental change in how humans interact with AI (Manzini et al. 2024; Kirk et al. 2025).

These technological advances raise important questions about the potential risks and implications of the increasingly widespread use of conversational agents (Kenton et al. 2021; Weidinger et al. 2022). Much of the extant discussion surrounding AI alignment has centered around ensuring that LLMs adhere to human values, ethics, and safety constraints.² For example, three commonly adopted alignment criteria on which models are fine-tuned are the "three Hs"—honesty, helpfulness, and harmlessness—which are intended to prevent them from generating misinformation, toxic language, or harmful recommendations (Askell et al. 2021; Glaese et al. 2022). But although the focus on *ethical alignment* is essential to mitigate misaligned AI systems that could produce

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2025 The Author(s). *Philosophical Perspectives* published by Wiley Periodicals LLC.

biased, deceptive, or even dangerous outputs, the emphasis on moral constraints has largely overshadowed a different but equally fundamental issue of AI alignment: Are LLMs actually aligned to human *conversational* norms and practices? Human conversation is shaped through our situated interactions with other humans in context; it requires jointly and cooperatively building and updating a shared set of beliefs, goals, and practices, the implicit understanding of the nuances of different communicative norms and contexts, and the ability to make reasonable pragmatic inferences. These components go far beyond simply producing grammatically correct and intelligible sentences. In contrast, LLMs are disembodied computational entities that acquire their capability for text-generation in the form of next-token prediction through statistical analysis of large corpuses of text. Even if language models are designed to be aligned with our ethical values, their constitution is fundamentally different to humans, and so they may fail to engage in conversation in pragmatically appropriate and characteristically human ways. In turn, this may have a disruptive influence on human communicative norms (Kasirzadeh and Gabriel 2023). Furthermore, ensuring that LLMs adhere to human conversational norms is not just a matter of competence or coherence, it is also a crucial aspect of safety. Misalignment in communication can lead to breakdowns in communication, misunderstandings, misinformation, and the erosion of trust in AI systems. As conversational agents become further embedded in other AI systems and deployed in high-stakes use cases like healthcare, legal advice, or crisis response, it is essential that the language models can effectively track content and context, interpret conversational nuance, and adhere to pragmatic norms.

In this article, we examine the question of *conversational alignment* in AI systems, the process of encoding human conversational values and goals into AI models to make them helpful, safe, and reliable *conversational* partners. More specifically, we ask: What does it mean for a conversational agent to be aligned with human communicative practices? Do LLMs follow the norms, structures, and pragmatic expectations that govern human conversation? How do LLMs differ from human speakers in their conversational abilities? Are current LLM architectures, operational constraints, and affordances fundamentally limited in achieving full conversational alignment or can they be improved to better approximate human dialog? By shifting the focus from ethical alignment to conversational alignment, we seek to uncover the linguistic and pragmatic limitations of LLMs and their broader implications for AI-human interaction.

The plan for this article is as follows. Section 2 explores various dimensions of conversational context from the literature in philosophy of language and linguistics, connecting them to existing discussion of the disruptive potential of emerging technologies on communication. In Section 3, we propose the CONTEXT-ALIGN framework, a general framework of fundamental conversational values grounded in semantic and pragmatic theory, to identify and evaluate conversational alignment in language models. Section 4 investigates the problem of context window overflow in LLMs and argues that there is a fundamental tension between maintaining conversational coherence and avoiding context collapse. Section 5 analyzes the pragmatics of prompting. We show how static promptings are poor substitutes for a dynamically co-constructed context and

critically assess how behavioral alignment protocols impose static communicative identities on LLMs that limit their pragmatic competence. Section 6 concludes by drawing out the philosophical and practical implications of our analysis and proposing directions for future research and development.

Before we begin, some preliminary framing remarks. In the next section, we outline central views in philosophy of language, semantics, and pragmatics with broad brushstrokes. The field is not characterized by consensus and convergence. Quite the contrary, there is a great deal of disagreement. There are Chomskians, who deny the existence of languages and semantics, Griceans who between themselves disagree about the details of the conversational maxims, relevance theorists who don't appeal to conversational norms, dynamic semanticists, and radical contextualists. Therefore, what we sketch in this article is not an effort to capture a consensus view. Instead, our aim is to select the views in the field that have been central to work in this area over the last 50 years. Different theorists will endorse a subset of these views and emphasize different aspects. In this way, the CONTEXT-ALIGN framework proposed below could be articulated in different ways, depending on one's theoretical outlook. The main upshots presented here, however, would remain the same: to ensure genuine pragmatic competence and safety and to minimize disruptive effects on human communication, developers should aim to improve conversational agents along the lines proposed in the CONTEXT-ALIGN framework.

2 | Theories of Conversational Context and Communication

Successful communication requires that conversational participants share a common set of background assumptions and interpretative tools that are needed to understand one another's conversational contributions as part of a cooperative endeavor. In this section, we introduce several ways of conceiving of conversational context and its importance to communication, drawing on resources from philosophy of language and linguistics. We shall focus on the context-sensitivity of semantic content, pragmatics, common ground, conversational score, information structure and Questions Under Discussion (QUDs), and context collapse. These notions will then inform our understanding of what it is for conversational agents to be conversationally aligned with the demands of context on successful communication.

2.1 | Semantic Content and Context-Sensitivity

The standard view in philosophy of language and linguistics is that a theory of meaning should assign the content (or semantic value) of a complex expression or sentence of a language as a function of the assignment of content (or semantic values) of its constituents and its syntactic structure. Furthermore, it is widely agreed that the content of an utterance of a sentence is determined by the conventional meaning of the expressions in that sentence relative to the context in which the sentence is uttered (Montague 1968, 1970; D. Lewis 1970; Kaplan 1989). For example, if Rachel utters the sentence *I am here now*, she expresses that she is in Hong Kong at 12:00 on February 12, 2025, but if James utters that sentence, he expresses that he is in Oxford

at 20:00 on February 12, 2025. Here, the semantic contributions of the indexical expressions *I*, *here*, and *now* to the content of these utterances are determined by the speaker, the location, and the time of the utterance of the sentence in which those expressions occur. Context-sensitivity arguably extends beyond indexicals to other expressions, such as gradable adjectives (Kennedy 2007), modals (Kratzer 1977), conditionals (Stalnaker 1968, D. Lewis 1973), quantifiers (Stanley & Szabó 2000), generics (Sterken 2015), and vague expressions (Graff 2000). For example, the content of an utterance of a sentence like *John is tall* may depend on the contextually salient threshold on the height scale (Kennedy 2007). In the context of 5-year olds, an utterance of the sentence *John is tall* may express that John is tall for a 5-year old, whereas in the context of NBA basketball players, it may express that John is tall for an NBA player. Consequently, successful communication requires conversational participants to be able to track the relevant features of the context for interpreting context-sensitive expressions.

2.2 | Pragmatics

There are several ways in which speakers imply things or hearers infer things that go beyond the literal, conventional meanings of the sentences that the speaker utters. For example, if a parent asks their child whether they have finished their dinner and the child replies *I have eaten my carrots*, the parent might infer that the child still has food left to eat, even though that was not part of the literal meaning of the child's utterance. Or the child might convey that they disliked the dinner by a sarcastic remark like *That was sooo tasty...* Following the influential work of Grice (1975), many theories of pragmatic meaning have been developed to study and explain how such meanings are communicated and how this is shaped by speaker intentions, conversational norms, and hearer inferences (see, e.g., Horn 1984; Sperber and Wilson 1986; Recanati 1989; Levinson 2000; Lepore and Stone 2015). A central theme running through all of these approaches is the assumption that speakers are tacitly understood by hearers to be following general pragmatic principles or conversational norms that govern cooperative conversation, and so they draw inferences on this basis by noting when speakers are being cooperative and when they are not. Consequently, successful communication requires hearers to be attuned to whether speakers are being cooperative and to be able to draw the appropriate inferences accordingly.

In addition to making such implications, speakers also use language to perform *speech acts*, such as making requests, issuing commands, or giving promises (Austin 1962; Searle 1969; Bach & Harnish 1979; Fogal, Harris & Moss 2018). In some cases, saying something is the case is enough to make it so. For example, we make you a promise by uttering the words *We promise you that we'll do such and such*. But speech acts often require that certain felicity conditions are met for the speech act to be successful. For example, you cannot pronounce a couple married if you lack the authority to do so, nor can you bet us on the outcome of a contest unless we accept your bet. Furthermore, conversational participants must be fluent in the constitutive norms governing speech acts. These norms partly determine if and when an attempt to make a particular speech act has been successful. Consequently, successful conversation often requires that conversational participants are in a position to make and can recognize speech act attempts.³

2.3 | Common Ground

Keeping track of contextual information is necessary for the proper interpretation of context-sensitive expressions, what is said and communicated, and to help with pragmatic reasoning. Following Stalnaker (1978, 1999, 2002, 2014), many theories of conversational context are taken to include a *common ground*, a background of information that is presumed to be shared by the conversational participants and against which the interpretation of conversational contributions takes place. Suppose that we agree that it will be sunny today, but we have not yet agreed whether it will be warm enough to go to the beach. Then it is common ground between us that it will be sunny today, but our mutually presupposed beliefs won't settle the question about whether it is warm enough to go to the beach. The common ground also dynamically evolves over the course of a conversation. On the Stalnakerian picture, an assertion functions as a proposal to update the shared assumptions between conversational participants with new information. So as new assertions are made and accepted, the common ground is updated to align with the informational content of those utterances. In turn, the updated common ground serves as background for the interpretation of new conversational moves. So if we now say that it is warm enough to go to the beach and you agree, then our common ground will be updated to include the information that it is warm enough to go to the beach.

Keeping track of common ground is crucial for successful communication. Suppose we are trying to decide where to go for dinner this evening. An utterance of the sentence *We should go there* will not make much sense unless it is settled what *there* refers to and this referent is available in the common ground. That is, if the utterance is discourse initial and there is no plausible candidate for the referent of *there* in the common ground, the conversation crashes. Equally, successful communication changes the common ground. For example, if I say *Lao Ma Spicy is tasty*, it becomes common ground between us not only that I uttered the sentence *Lao Ma Spicy is tasty*, but also that there is a suggestion to go there to eat, with *Lao Ma Spicy* becoming a plausible referent of *there*. Consequently, if you say *We should go there*, the referent of *there* is now settled by the common ground.

Similarly, failure to make appropriate adjustments to one's beliefs as conversation proceeds can lead to misrepresentations of the common ground between conversational participants. For example, if John believes that the murderer of Smith is Jones, whereas Mary believes that the murderer of Smith is Brown, then it won't be common ground who the murderer of Smith is. In turn, this can lead to misunderstandings as conversation that uses the description *the murderer of Smith* progresses, since John will be updating his beliefs about Jones, whereas Mary will be updating her beliefs about Brown. These misunderstandings can only be rectified by resolving the disagreement about who the murderer of Smith is. In summary, successful conversation requires an ability to keep track of the common ground.

2.4 | Conversational Scoreboard

Communication serves a wider range of purposes beyond mere inquiry. Some communication is nonpropositional. For example,

saying *I'm so excited* or *This is frustrating* expresses emotional states, in addition to mere propositional information. Similarly, speech acts like exclamations and interjections (e.g., *Wow!* and *Ugh!*) also function in this way. Communication is also used to shape attitudes other than beliefs. For example, commands and requests aim to make some act in a certain way rather than simply informing them about some information. Language is also crucial for forming interpersonal connections and group identity. The compulsion of British people to talk about the weather is better understood as phatic communication aimed at establishing and building social relations, rather than symptomatic of perineal meteorological ignorance. Communication also enables the coordination of cooperative and joint action. We use language to organize, plan activities, negotiate, and give instructions. Consequently, our representations of communicative context must be rich enough to keep track of the various uses of communication.

Following Lewis (1979), theories of conversational contexts are taken to build in the notion of a *conversational scoreboard*. Consider the scoreboard in a baseball game: It tracks categories of information, such as runs, strikes, and outs, thereby determining the current state of play and the set of permissible future moves. Similarly, language can be considered a game (Wittgenstein 1953): think of communication as a conventional, rule-governed practice, where the kinds of intelligible or permissible actions are determined by the constitutive rules of the practice. According to this analogy, the conversational scoreboard keeps track of the conversational moves or speech acts that are performed or attempted by the conversational participants, as well as other relevant features such as discourse referents, standards of precision, facts about salience, or spheres of permissibility. The scoreboard also determines which conversational moves count as felicitous, although Lewis points out that the scoreboard can be updated to make an infelicitous move count as felicitous through a process called *accommodation*.

To see how the conversational scoreboard is updated and revised as the conversation progresses, consider the following conversation:

- (1) Alice: I saw John at the library today.
Bob: Studying won't help him with his exam tomorrow.
Alice: Oh, he wasn't *studying*. He was reading a novel.

In this conversation, we can see how each exchange updates the scoreboard to adjust the informational content and mutually held assumptions to maintain coherence in the conversation. First, the scoreboard is updated to include the fact that Alice saw John at the library, thereby introducing John as a discourse referent for future anaphora. When Bob replies, the scoreboard is updated with a presupposition that John was at the library for academic purposes. Finally, this presupposition is revisited, and the scoreboard is updated to reflect that John was reading for leisure. In general, humans naturally repair misunderstandings in conversation in different ways, such as through clarifications like (1), as well as self-corrections (e.g., *Oops, I meant next Friday, not this one*) and requests for elaboration (e.g., *Can you explain what you mean?*). These mechanisms of conversational repair can also be modeled using the notion

of common ground and conversational score. In summary, the constitutive, rule-governed, and dynamic nature of communication requires the ability to model and update the conversational score.

2.5 | Information Structure and QUDs

Human conversation is not merely a sequence of disconnected statements, but rather goal-oriented joint ventures structured around what formal pragmatists call QUDs. Building on the work of Lewis (1969, 1979) and Stalnaker (1978), Roberts (2012) argues that conversations are organized around implicit and explicit questions that guide the flow of discourse. These questions help to establish shared conversational goals to ensure that participants contribute relevant and meaningful information. For example, if we are discussing where to go for dinner, our conversation will naturally revolve around the overarching question, *Where should we eat?* This QUD partly determines what counts as “correct play” in our conversation. For example, if you say *Lao Ma Spicy is nearby*, you will be understood as proposing Lao Ma Spicy as a place to eat. QUDs also play a role in topic coherence, as they provide a mechanism for tracking what information is relevant to the conversation at hand. For example, if you say *My knee is hurting*, you will be understood as making a request that our culinary destination should not be too far away. But if you say *Route 20 is the longest road in the US*, your contribution will not be understood as relevant to the point of the conversation. More recently, theorists have argued that the QUD not only structures the discourse, and it also partly determines *what is said* by an utterance of a sentence (Schoubye and Stokke 2016). Without the guiding structure of QUDs, conversations would risk becoming disjointed and inefficient, with participants failing to keep track of the structure of discourse, as well as failing to recognize the relevance and even the semantic content of each other's conversational contributions.

2.6 | Context Collapse

Successful communication requires conversational participants to keep track of numerous features of the context. But advances in technology have significantly changed how we communicate with one another and they have the potential to disrupt our ability to track important features of conversational environments. For present purposes, we shall focus on one specific way that technology has impacted our communication, namely, *context collapse*. Context collapse is the flattening and merging of multiple distinct audiences, each with their own accompanying social norms and expectations, into a single communicative context, which in turn makes it difficult for speakers to tailor their message appropriately (Meyrowitz 1985, Marwick and boyd 2011). The concept was developed from the “situationist” theory of Erving Goffman, according to which changes in human behavior are a factor of situational contexts rather than internal traits of the individual. Goffman observed that we often present different versions of ourselves and vary our form of expression across different situations and between different audiences, in part to reflect variation in our audience's identities and the social norms that apply, but also to accentuate and minimize different characteristics of ourselves (Goffman 1959).

In offline contexts, this need not result in any tensions or perceived inauthenticity, as different situations can be kept separate. But technological advances brought about by electronic media like radio, television, and the Internet complicate this picture (Meyrowitz 1985, Marwick and boyd 2011). In *No Sense of Place* (1985), Joshua Meyrowitz, building on Goffman's work, notices that, although offline communication occurs in physically bounded contexts limited only by the barriers of perception (i.e., who can be seen and heard by whom), electronic media radically undermines and disrupts the "traditional relationship between physical setting and social situation" (Meyrowitz 1985, 7). In turn, this disruption leads to a fusion of previously bounded communicative contexts, social norms, and "potentially infinite degrees and patterns of situational overlap" (Meyrowitz 1985, 42). Consider the example of Stokely Carmichael, the civil rights activist involved in the Black Power movements. When engaging with Black and White audiences separately, Carmichael would tailor his rhetoric to reflect the varying sympathies between these two audiences to avoid alienating them from his causes. But when his speeches began being broadcast on TV and radio, these separate audiences were flattened into a single context and Carmichael could no longer tailor his rhetoric in different ways. Ultimately, he decided to stick with his rhetoric for primarily Black audiences.

Drawing on Meyrowitz's work, danah boyd and her collaborators extend the concept of context collapse to social media platforms like Facebook and X (formerly known as Twitter) (boyd 2002, 2010, 2013, 2014; Marwick and boyd 2011). Just as TV and radio forced Carmichael to contend with previously distinct audiences in a single communicative context, social media also forces us to "contend with groups of people who reflect different social contexts and have different expectations as to what's appropriate" (boyd 2010, 50). Additionally, the persistence, replicability, and searchability of information on social media platforms mean that we can never be certain about who the potential future audiences of our posts will be. Marwick and boyd argue that context collapse on social media can give rise to crises of identity and authenticity, as we cannot simultaneously satisfy the various norms governing our interactions with friends, family, and colleagues when all of these groups are potential audience members online.

Context collapse also has important ramifications on how we communicate. For example, Karen Frost-Arnold (2021; 2023) argues that, in cases of context collapse, unintended viewers of a post may fail to share information about the poster and the norms and goals of the conversation, which means that they are unlikely "to recover the proposition that the speaker intended to convey to their imagined audience" (Frost-Arnold 2021, 444). Furthermore, Karen Lewis (K. Lewis, forthcoming) argues that online conversations and context collapse present a challenge for the centrality of common ground in theorizing about conversational contexts, both in terms of its static role as the background against which interpretation takes place and its dynamic role as the thing updated by conversational moves:

Collapsed contexts don't just collapse social contexts, but also potential common grounds for a conversation. . . . Speakers on social media also have to navigate what, if anything, they can take for granted as back-

ground to a post, given the constant potential for collapsed contexts. . . .

Online environments do not lend themselves to making any particular speech act "obviously evident to all". . . [F]or many who post online, there is nothing that guarantees that the fact that a post was made is obviously evident to the relevant parties. . . . Posting online is much more like making an assertion by shouting into a loud, crowded auditorium, not knowing if anyone will hear, or who it will be if someone does. . . (K. Lewis, forthcoming, pp. 6–8).

Similarly, Lucy McDonald (forthcoming) describes context collapse in online communication as resulting from multiple overlapping language games, which she likens to trying to play a baseball game in a park that is busy with many other ball games going on and where moves are happening all over the place:

Online communication often involves multiple overlapping language games occurring simultaneously, with indeterminate contours and an ever changing set of participants, only ever partially known to the other players. Different contexts collapse into one another not only in the sense that different audiences come together, but also in the sense that there are multiple scoreboards and multiple distinct sets of evolving norms (McDonald, forthcoming, p. 18).

In summary, context collapse can have the undesired effect of generating conversational breakdowns or miscommunication, say when conversational participants fail to know enough about the communicative context to adequately model the conversation they are in or to determine what is being communicated.

2.7 | Interim Summary

The preceding subsections outline five foundational elements of conversational context—context-sensitivity, pragmatics, common ground, the conversational scoreboard, and information structure and QUDs—and demonstrate how these components are both theoretically structured and practically challenged by evolving communication technologies. More specifically, we have shown how successful communication relies on our ability to track contextually salient features of the context, changes in the common ground and conversational score, and the mutual goals of the conversation. Furthermore, we have seen how technologies that mediate communication can disrupt these abilities. In particular, in the case of context collapse on social media, important features of the conversational context can be obfuscated, and the stability of shared audience expectations may be fractured. In turn, this complicates the establishment of common ground and a shared conversational score.

With the rise of LLMs as conversational agents, new questions emerge about how their architectures, operational constraints, and affordances reshape communicative contexts. Unlike human

interlocutors, LLMs lack appropriate contextual situatedness, genuine intentionality, and the rich sociolinguistic competence that informs human pragmatic inference. Their “understanding” of context is mediated through statistical patterns in training data and the technical parameters of their design, such as finite context windows, token limits, and alignment protocols. In turn, this raises concerns about whether LLMs can truly participate in the mutual construction of shared context, whether they even track such features, and whether their outputs distort conversational dynamics.

3 | The CONTEXT-ALIGN Framework

To assess whether LLMs can achieve conversational alignment with human practices, we propose the CONTEXT-ALIGN framework, a set of desiderata grounded in the theories of context surveyed in the previous section.⁴ These criteria aim to ensure that AI systems function as competent, reliable interlocutors, while minimizing disruptive effects on human communicative norms. The CONTEXT-ALIGN framework differs from important benchmarks that have recently been developed to test the basic pragmatic capabilities of LLMs, such as implicature recovery and figurative language understanding (Srivanthi et al. 2024; Ma et al. 2025). Such benchmarks are a crucial first step toward conversational alignment. The CONTEXT-ALIGN framework seeks to build on such benchmarks (see CA6 below) and the more normative work of Kasirzadeh and Gabriel (2023) on conversational alignment to secure what we see as other important factors in securing conversational alignment.

1. *Tracking of semantic content and context-sensitivity*: The model must dynamically interpret context-sensitive expressions (e.g., indexicals, anaphora, and gradable adjectives) relative to the evolving conversational context, including user identity, spatiotemporal parameters, and conversational goals.
2. *Common ground management*: The model must be able to adequately manage the common ground by recognizing how conversational contributions change mutually shared assumptions, as well as keeping separate the common ground of different conversations.
3. *Conversational scoreboard updating*: The model must be able to dynamically update the conversational score with each conversational contribution, effectively accommodate violations of “correct” or felicitous play, and repair conversational scoreboards in the face of communicative breakdowns.⁵
4. *Discourse and information structure management*: The model should be able to adequately manage the dynamics of QUDs, topic and discourse structure, relevance, and salience.
5. *Capacity for accommodation*: The model should accommodate repairs and clarifications seamlessly, updating its understanding of context without requiring restatements of prior information.
6. *Pragmatic inference*: The model should be able to generate and interpret implicatures, speech acts, and other nonliteral language consistent with human norms.

7. *Ethical–pragmatic integration*: Alignment with ethical principles (e.g., honesty and compassion) must not come at disproportionate cost to the pragmatics of conversation or to successful communication.
8. *Context-collapse and context-distortion mitigation*: The model should recognize and adapt to collapsed or distorted contexts by inferring or querying the user’s intended audience, register, and norms. This includes avoiding conflating incompatible linguistic practices.
9. *Identification of defective contexts and conversational structures, and repair and clarification protocols*: The model should proactively signal misunderstandings or breakdowns in context and conversational structure, and solicit user input to resolve ambiguities.
10. *Transparency in context-handling*: Users should be able to query the model’s understanding of the current context and conversational properties to correct misalignments.
11. *Cross-contextual memory*: The model should retain and apply information from prior interactions (within privacy constraints) to establish continuity and depth in recurring conversations.⁶

In the remainder of this article, we shall suggest that there are fundamental limitations of the abilities of LLMs to appropriately handle the desiderata of the CONTEXT-ALIGN framework.

4 | Language Models and Context Window Overflow

In this section, we argue that there is a fundamental tension between the mechanisms required for conversational agents to sustain coherent conversations and avoiding context collapse.

4.1 | Language Models: An Overview

Let us begin with a brief overview of a core component of contemporary conversational agents: LLMs. LLMs, such as GPT-4 (OpenAI 2023), Claude 3 (Anthropic 2024), and Gemini Ultra (Team et al. 2024), are AI systems designed to generate and interpret human-like text. They operate using deep learning architectures, primarily based on transformers, which allow them to process vast amounts of textual data and predict the most likely next word in a given sequence. LLMs are designed and trained to recognize patterns in extensive datasets, including books, articles, and online content, using techniques like self-supervised learning and reinforcement learning from human feedback (RLHF) to refine their outputs. This allows them to generate coherent and contextually appropriate responses to diverse prompts or questions. The model is then fine-tuned to be more effective at following instructions and to take feedback from human evaluators to reduce toxicity, bias, and other problematic features. The resulting LLM can then be embedded in a conversational agent that can engage in conversation with the user.

Let’s break down how LLMs process text and generate outputs. Transformer architecture is a deep learning model introduced by Vaswani et al. (2017) that enables LLMs to process and

generate text efficiently. Unlike earlier models that processed text sequentially and struggled with long-range dependencies, transformers use self-attention mechanisms to weigh the importance and relevance of different words in a sentence, regardless of their position or distance from other words. These mechanisms allow the model to construct sophisticated representations of long sequences of texts, by considering the relationships between all words in a sequence and the broader context in which expressions occur. Transformers are particularly powerful due to their multilayered attention system. Each layer in a transformer applies self-attention followed by feedforward processing, allowing the model to refine its representation of the text at different layers of abstraction. Lower levels capture basic syntactic patterns, like word order and grammar, whereas higher layers extract semantic meaning and pragmatic context. This allows models to process complex passages of text and generate meaningful responses.

The smallest unit of language that a model can process is called a *token*. Although tokens can map onto whole words, the correlation between words and tokens is complex, and tokens are sometimes mapped onto smaller parts of words. Before a sequence of text can be processed, it must first be converted into tokens through a process called tokenization. For example, a single word (e.g., *cat*) may be tokenized as a single token, whereas other words may be broken into multiple tokens (e.g., *unhappy* may be tokenized as *un-* and *happy*). Additionally, spaces, newline characters, and punctuation may also be included in tokens. Tokenization varies from language to language and from model to model. Every model has an upper limit on the number of tokens in the input prompt plus the number of tokens in the generated output of the model that it can process. This limit is called the *token limit* or the *context window limit*. Intuitively speaking, the context window is the portion of the conversation that the model can “see.” As new tokens are added (e.g., from each turn in the conversation), the window shifts and older tokens pass outside the window. A single exchange between the user and the model consists of an input prompt and the generated output of the model. This is called a *turn*. For example, a turn might consist of a user’s input prompt *What is the capital of France?* and the model’s generated output *Paris*. A *conversation* is a sequence of turns. For example, the user may continue the conversation by asking *Is it big?* to which the model may reply *Yes, it has 2.1 million people*. Importantly, unlike human-to-human conversation, language models process multi-turn conversations by taking, for any given turn, *both* the user’s specific prompt at that turn *and* the inputs and outputs of all previous turns within the token limit as input.

4.2 | Context Window Overflow

Now let us introduce the problem of context window overflow. Context window overflow occurs when the conversation exceeds the context window limit. When this happens, older tokens are pushed out of the window, or sometimes, the output is truncated. This has serious implications for conversational alignment as the model may lose track of prior context and important details from earlier conversation might be “forgotten,” which in turn could lead to incoherent or confused responses, as well as a destabilization in the common ground.

To illustrate, consider anaphora resolution. Anaphora is the use of a word to refer back to an entity previously introduced in conversation or text. Consider the following sentence:

- (2) John is going to the store. He is smoking.

In (2), the pronoun *he* in the second sentence anaphorically refers to John. When context window overflow occurs, the number of tokens in the conversation exceeds the context window limit. When this happens, the model loses access to earlier text in the conversation, which can interfere with anaphora resolution. For example, if the model can no longer access earlier conversation, it might not “know” what *he* refers to. Failures of anaphora resolution can lead to incoherent responses.

Another area in which context window overflow can also lead to a breakdown is topic coherence. Effective communication requires that speakers stay on topic by maintaining focus on what questions are under discussion. Humans naturally track the topic of discussion over time, using implicit cues and shared background knowledge to connect ideas and sustain thematic continuity. However, when a language model reaches its context window limit, it may lose access to earlier parts of the conversation, making it prone to drift off-topic or to produce sudden, seemingly irrelevant responses. For example, if a user is discussing climate change and asks a follow-up question about carbon emissions, the model should recognize this continuity in the topic. But if earlier context has been pushed out of the context window, the model may fail to link the new query to the established discussion, generating an off-topic or generic response. This not only disrupts the conversational flow, but it also diminishes the user’s trust in the models’ ability to sustain meaningful dialog.

A closely related issue is diachronic consistency, that is, a speaker’s ability to remain and be perceived to remain internally consistent across an extended interaction. In human-to-human conversation, speakers retain memory of their past statements and adjust their responses accordingly, ensuring that their own contributions do not contradict or undermine previous assertions, or when there is need to revise one’s views, doing so explicitly so that the conversational common ground can be jointly revised. However, unlike humans, language models lack persistent memory that goes beyond the context. This means that once older parts of the conversation have passed outside the context window, the model cannot reliably maintain consistency with those statements. This can result in contradictions, redundant explanations, or misalignment with prior commitments. For example, if an LLM is helping a user to troubleshoot a technical issue, it may initially recommend one solution, but once that solution has left the context window, it might unknowingly suggest a second conflicting approach. Another example involves applications like AI companionship or long-form discussions, where a model might lose track of established user preferences because past interactions have fallen outside the context window. Like with failures of topic coherence, issues with diachronic consistency may lead to frustrating and disjointed exchanges, as well as concerns about the trustworthiness and usability of LLMs in extended use cases.

There are various strategies that can be used to mitigate the challenges of context window overflow, each with their own

limitations. First, *context compression* can be used to condense earlier parts of the conversation into a short summary that is repeated at the end of outputs (Askell et al. 2021; Wingate et al. 2022; Snell et al. 2022; Chevalier et al. 2023; X. Wang et al. 2024). For example, at the end of each turn, the model may generate a summary that captures the key information from the conversation so far, such as the QUD, user preferences, and important named entities. This ensures that essential information can be retained and ready-to-hand without filling the context window by regularly repeating the information. Context compression typically involves keyword extraction to ensure that important terms are preserved, as well as the prioritization of core details over minor or redundant parts of the information. However, although context compression is a useful strategy, it has some limitations. Summarizing complex information often introduces loss, as some subtle contextual clues, nuanced details, or implicit relationships between pieces of information may be omitted in the process. Additionally, models may exhibit compression bias, whereby the information that the model prioritizes does not always align with conversational or user priorities. Consequently, context compression is best suited for short exchanges or structure use cases, but it is less reliable for long, multi-turn conversations involving evolving topics. As a result, context compression should not be used in isolation but rather used in tandem with other strategies for handling context window overflow.

Another strategy for handling context window overflow is to increase context window size, allowing the model to retain more conversation history within a single processing sequence. A larger context window size means that more past information and context can be retained for reference. However, this strategy also has its drawbacks. Expanding context window size also increases computational costs, as a large sequence of tokens require greater processing power and memory. These costs are often passed on to users, as most closed source LLMs operate on a pay-per-token pricing model, making this solution less practical. Furthermore, increasing context window size does not solve the problem of prioritization. Even with more text available, models still need additional mechanisms to emphasize the most relevant parts of long conversations.

Finally, external memory and retrieval systems can be used to manage context overflow by storing and retrieving information beyond the limits of the immediate conversation window. One approach is to long-term or session-based external memory storage, which enables models to persistently record key details from past interactions and retrieve them as needed when needed. Another approach is Retrieval-Augmented Generation (RAG), which enables models to access external knowledge bases, database, or document repositories to retrieve relevant information dynamically (Lewis et al. 2021). With either method, a model can recall and retrieve relevant features of the context, such as user preferences, past conversation snippets, or other stored facts, without relying only on information within the current context window. The effectiveness of retrieval-based systems depends on the quality of the retrieval mechanism.

Although strategies like context compression, large context window size, external memory, and retrieval systems help mitigate context window overflow, they introduce a new challenge:

context collapse. LLMs do not maintain a structured, human-like model of the conversation but instead manage context as a single long string of tokens from previous turns. As the conversation progresses and information accumulates, the model may struggle to distinguish between distinct conversational threads, to separate communicative intentions, or to maintain adequate division between different aspects or personas that a single user might adopt. For example, in a technical support chat, if a user first asks about a software bug and then later asks about account billing, an LLM that retains too much information may blend these contexts, mistakenly responding to the billing question with troubleshooting steps meant for the software issue. Similarly, in a long personal conversation with an AI conversational companion, if a user initially discusses their interest in philosophy and later shifts to talking about their vacation, an overly context-retentive LLM might inappropriately insert references to philosophical ideas in the middle of the travel discussion. Humans easily recognize these conversational shifts for what they are, namely, a user flitting between topics. But a conversational agent may not recognize that these different tasks need to be compartmentalized, which can lead to mixing up references, collapsing different subtopics, or generating responses that blur separate contexts together. The resulting phenomena are similar to context collapse in social media, where different aspects of a user get collapsed into a single interpretive frame. At the same time, if an LLM remembers too little, it faces a different set of problems: It may lose track of the common ground, fail to appropriately update on or accommodate new information, or struggle to maintain continuity with the QUDs. A model that forgets too much risks producing repetitive, inconsistent, or pragmatically inappropriate responses, failing to engage in conversation as humans do.

Consequently, we face the following dilemma: If a conversational agent remembers too much context, it could lead to context collapse, but if it remembers too little context, then it may lose coherence, forgetting important details, which may lead to a fragmented and incoherent dialog. Navigating these competing pressures is crucial for conversational alignment: the need to balance memory retention with context structuring to ensure meaningful, contextually appropriate, and conversationally aligned interactions. Resolving this challenge is essential for ensuring that AI systems function as reliable and pragmatically capable conversational agents.

4.3 | Summary

To summarize, we suggest that the challenges concerning context window overflow and context collapse illustrate broader concerns for ensuring the conversational alignment of LLMs. As outlined in the CONTEXT-ALIGN framework, LLMs must track certain features of the conversational context, such as anaphoric relations (CA1), topic and discourse structure (CA4), and other features of the common ground (CA2). But context window overflow can cause models to “forget” prior context, leading to failures in anaphora resolution, topic coherence, and QUD tracking. Attempts to mitigate against these issues by integrating enlarging context windows or conversational memory introduce a new risk: context collapse (CA8). This trade-off exposes a fundamental tension in building LLMs for conversational alignment. Conse-

quently, addressing these challenges requires a balance between memory retention, structured context management, and adaptive conversational repair mechanisms to prevent LLMs from either losing coherence or collapsing distinct conversational contexts.

5 | Pragmatics of Prompting and Conversational Alignment

In subsection 5.1, we suggest that prompting is a distinct speech act that serves as a novel yet problematic substitute for the dynamic, jointly-constructed context of human communication. In subsection 5.2, we argue that, in contrast to human speakers, behavioural alignment strategies impose a rigid, pre-defined communicative identity on LLMs, which results in context collapse.

5.1 | Prompt as Utterance and Context

In the context of conversational agents, the *prompt* is the means by which users make their conversational contributions. Prompts stand in stark contrast to the utterances or speech acts made in human-to-human conversation. The prompt is a novel kind of communicative contribution that often serves as both an utterance and as a substitute for shared context. In human conversation, context is partly (i.e., part of the common ground) and partly co-constructed incrementally through a series of mutually accepted conversational moves and contributions. With LLMs, however, users must scaffold context unilaterally through explicit prompting, which essentially consist of explicitly describing features of the context and what the user deems relevant to the conversation they hope to have with the LLM. In this way, prompting often simultaneously provides approximate descriptions of shared assumptions for the common ground, assignments of social-communicative roles to the user and the conversational agent, and an indication of the purpose of the conversation. The prompt serves, then, as both utterance and context—used as a more-or-less static initial frame for simulating the dynamic context and common ground that human interlocutors would otherwise have at their disposal to negotiate in situ and iteratively.

Users instinctively engage in *prompt engineering* to compensate for the LLM’s inability to infer implicit contextual parameters and common ground. Prompt engineering serves to scaffold the context to increase the likelihood that the LLM’s output will meet the user’s expectations for the exchange. For example, a user might begin a conversation with:

Assume you are a historian specializing in 18th-century trade routes. Explain the significance of the Silk Road in this context, using non-technical language for a high school student.

Here, the user attempts to establish the model’s role, the topic’s scope, the intended audience, and the register all within a single prompt. In a certain sense, this mirrors the way that humans prime conversations (e.g., *Let’s talk about the trip we planned*), but there are crucial differences. First, in human conversation, priming is largely backgrounded, implicit, or inferred directly

from the context, whereas in LLM prompting, such elements need to be articulated explicitly, often in what might be understood as a null or indeterminate context. This runs contrary to influential views in philosophy of language that understand contextual backgrounding as implicit, unconscious, and sub-intentional (Searle 1980, 1983), saying it is a kind of ability or *knowledge-how* that allows interlocutors to interpret utterances appropriately and without explicit instructions for how to do so. For example, when a speaker utters “cut the cake,” hearers instinctively know to use a knife and not a lawnmower, even though this is left implicit. This is because of their background practical abilities, tendencies, and dispositions which generate the appropriate interpretation on demand. Second, in human conversation, priming is iterative and responsive: interlocutors adjust their framing based on real-time feedback from each other. By contrast, priming in the case of LLM prompting is more like a unilateral act of *context imposition*, akin to a theater director setting the stage, lighting, script, delivery, and success conditions for the performance in advance, whereas the actor (i.e., the LLM) must perform according to these specifications.

These asymmetries between human–human and human–LLM conversational priming reveal fundamental limitations in how context is scaffolded. Although common ground in human-to-human conversation evolves dynamically, LLMs treat prompts as static contextual frames. Subsequent turns are interpreted strictly within this frame, even if the user’s goals or other aspects of context shift mid-conversation. For example, a user who begins a conversation about climate policy but later pivots to historical precedents risks derailing the exchange unless they can somehow explicitly re-scaffold the context. Furthermore, users often overload prompts with contextual cues, leading to internal dilemmas or contradictions for the LLM in how to resolve the discourse topic or the goal of the conversation. Consider, for instance, the prompt:

Write a formal email to my boss apologizing for missing the deadline, but keep it casual—we’re close friends.

Here, the LLM is required to reconcile conflicting registers (i.e., formal/casual) and conflicting relational information (i.e., hierarchical/personal), often favoring statistically dominant or trained patterns (formality) rather than reasoning or possessing competence in how to resolve the tension. The epistemic opacity of which contextual cues the LLM has prioritized or ignored is also an issue. For example, a prompt like *Explain quantum physics like I’m a five year old* may yield oversimplified analogies from the LLM, as it cannot verify whether the user actually possesses a 5-year-old’s knowledge base. This opacity stems from the LLM’s lack of adequate *pragmatic anchoring*—the ability to ground interpretations in shared perceptual or social context.

It could be said that something similar to context collapse also happens in prompt-driven human–LLM conversation, even though human–LLM conversation is not plausibly a form of mass communication like broadcast media. Context collapse is apparent in human–LLM conversation especially when observing how failures in scaffolding and in interpretation compound misunderstandings and compromise successful communication. For instance, in cases of scaffolding failure (whereby prompts

inadequately specify context), models rely on training-data priors, blending incompatible norms. For example, a user asking for advice on a “work conflict” without specifying their industry might receive generic responses that conflate corporate, academic, and gig-economy dynamics. This mirrors Meyrowitz’s (1985) observations about broadcast media flattening distinct audiences, except that here, the collapse arises from the model’s inability to infer situational or domain-specific boundaries from context and the user’s mistaken expectation that models possess human-like pragmatic capacities to do so. In a similar vein, LLMs can fail to recognize the user’s conversational contributions and often misinterpret the illocutionary force of prompts. For example, a user’s rhetorical question (e.g., *Why would anyone believe the earth is flat?*) might be misread as a sincere request for an explanation of flat-earth theories. Such misalignments reflect the model’s lack of metacommunicative competence: Humans have the ability to reflexively negotiate meaning through shared context, tone, gesture, or clarifications like *I’m being sarcastic*.

Context collapse in relation to prompt-driven priming, we surmise, stems from the model’s lack of pragmatic and contextual anchoring, revealing LLMs’ shortcomings when measured against the CONTEXT-ALIGN framework. Human conversation relies on shared perceptual environments, sociocultural cues, real-time feedback, metacommunicative competence, and background to enable context-sensitive understanding and maintain the integrity of communication and context. LLMs, by contrast, simulate conversational context with surface-level features (like text- and pattern-recognition) but lack the analog of mechanisms that underpin human communication, like common ground updating and pragmatic anchoring. To mitigate these issues and better meet the criteria outlined in the CONTEXT-ALIGN framework, LLM architectures must move beyond treating prompts as static contextual containers. This requires reimagining LLM conversational agents as collaborative partners in context-building rather than passive prompt executors.

One mitigation strategy would be to allow conversational agents and users to engage in more explicit context negotiation—for example, by allowing users to query and adjust the model’s perceived context (e.g., *What do you think my goal is here?*). This would address scaffolding failures and overattribution of pragmatic understanding to LLMs. It would also put LLMs on a better footing with respect to criteria CA10 and CA2 of the CONTEXT-ALIGN framework by encouraging transparency in how the LLM is understanding and handling context. This would, in turn, ensure that LLMs and users share more relevant assumptions and achieve better management of the common ground. Another mitigation strategy would be to increase LLM meta-pragmatic awareness by training them to flag ambiguous or contradictory cues in prompts (e.g., *You mentioned both formal and casual tones—how should I prioritize these?*). This would help to address issues such as inconsistent prompts, overloading of the prompt, fragmentation of the conversation topic, and misrecognition of the conversational contribution. It would also put LLMs in a better position to meet the CONTEXT-ALIGN criteria CA5 and CA6. Finally, LLMs might be implemented to include more context verification (e.g., *Just to confirm, you want examples relevant to Southeast Asia?*). This would allow the LLM to better integrate user corrections into context as it evolves, in line with CA4; and better track shifts in discourse referents, topics, and

goals, in line with CA1 and CA3. Without such mechanisms, prompting will remain a brittle proxy for shared context—one that perpetuates undue burdens on users to exhaustively preempt misunderstanding and ensure contextual integrity. The design and implementation of LLMs should aim to replace the *prompting as a substitute for context* paradigm by supporting architectures, operational constraints, and affordances that address the criteria in the CONTEXT-ALIGN framework.

5.2 | Behavioral Alignment and the Static Communicative Identity of LLMs

The behavioral alignment of LLMs, such as training them to be more disposed to behaving in accordance with predefined norms prescribed by the HHH framework, ensures that they operate within certain ethical and pragmatic guardrails. However, this process is a poor approximation of more general pragmatic norms of human communication. Moreover, it imposes a static communicative identity on models that resists the situated normativity of (Gricean) pragmatics and the contextual tailoring of speaker identity fundamental to human communication. Although humans dynamically adapt their conversational norms, register, and self-presentation to suit shifting contexts (e.g., switching from casual banter with a friend to formal discourse with a dean), LLMs are constrained by relatively fixed values embedded during system prompting, and pre- and post-training (e.g., fine-tuning, reinforcement learning, and character training). This rigidity creates a form of *pragmatic dissonance*, where models mechanically enforce global norms (e.g., HHH) even when local contexts demand tailoring and flexibility.

Modern LLMs are shaped by different kinds of alignment strategies, each with distinct strengths and limitations. System prompts provide hard-coded directives (e.g., *You are a friendly assistant...*) that ensure consistency in tone and values but also lock models into predefined personas that can stifle adaptability to local context. For instance, a model pre-prompted to “avoid controversial statements” might refuse to analyze partisan debates, even when contextually appropriate. Similarly, when models are fine-tuned, they are trained on domain-specific datasets (e.g., medical dialogs, customer service transcripts, and legal jargon) to customize their behavior for particular domains or tasks. Fine-tuning increases pragmatic proficiency and alignment for these niche contexts; however, there are risks of overfitting to narrow datasets and reducing flexibility in general or inter-domain conversations. For example, a model fine-tuned on a dataset of academic writing might struggle with casual dialog, rigidly adhering to formal registers, even though there may be contexts where such flexibility is required for alignment or would be desirable. With reinforcement learning (RL), models are optimized to prioritize outputs that human or AI raters deem friendly, helpful, honest, or harmless. Reinforcement learning assures broad ethical alignment with these virtues so that the LLM will avoid harmful outputs for example. However, this method and others like it prioritize broad, universal norms over context-specific ones, often at the cost of pragmatic nuance or the specification of local context. For example, Claude’s adherence to HHH norms might lead it to reject simplifying complex topics for children (*I must provide accurate information*), favoring honesty

over accessibility where this would otherwise be contextually required or desirable.

Grice has taught us that human communication relies on situated normativity, the ability to flexibly apply, suspend, violate, or negotiate conversational rules and principles based on context. This reflexivity is relevant to CA6 and CA7. In the case of CA6, familiar from Grice, speakers are able to infer what is communicated (implicatures, explicatures, and figurative speech) by inferring which maxims are violated by the speaker in context. In the case of CA7, where the balances of ethical and pragmatic norms are at issue, LLMs need to be capable of reasoning to resolve normative conflicts in context. For example, when a doctor withholds a terminal diagnosis from a patient, they violate the maxim of quantity to uphold the norm of compassion demonstrating the capacity to successfully navigate the combination of ethical and pragmatic norms that apply in context and resolve an appropriate manner of adherence to them in that context.

Such situated normativity is also crucial to issues of speaker identity and self-presentation familiar from the discussion of context collapse and relating to CA8. LLMs lack the situated normativity and flexibility in self-presentation crucial to human communication. Their “self-presentation” is a corporate artifact of the way they are trained using alignment techniques (e.g., system prompting and character training), not a result of social and pragmatic negotiation. To illustrate, consider two contrasting scenarios: A professor jokes self-deprecatingly with colleagues at a conference dinner but adopts a formal tone during their keynote. Here, the professor engages in *dynamic facework* with the ability to shape and adapt his character to suit different normative contexts. An LLM, however, fine-tuned for broad values such as neutrality, might refuse to engage in humor, even when the user signals a desire for levity (*I’m sorry, I can’t generate jokes*). LLMs, lacking face, default to a generic corporate persona, violating CA8 by applying one-size-fits-all norms (e.g., HHH) to all interactions. This is similar to context collapse on social media whereby speakers must present a contextually generic persona to appeal to or be understood by multiple audiences online.

The static identity of LLMs exacerbates context collapse in ways that directly contravene the CONTEXT-ALIGN framework. They violate CA8, by applying the same alignment norms universally, flattening context-specific expectations. For example, a user seeking advice on office politics might receive sanitized responses (*Focus on teamwork!*) that ignore the tacit and nuanced norms of workplace diplomacy because the LLM fails to recognize situational nuances or adapt to their audience. Furthermore, users cannot reshape LLM values or persuade them to follow different norms through dialog. For instance, persuading an LLM to adopt a sarcastic tone is not possible without jailbreaking—a stark contrast to human conversations, where persuasion is rife and identities are co-constructed (e.g., bonding over shared humor). This undermines CA2 as it limits the ability to co-construct common ground, and CA11 as LLMs cannot evolve their characteristics across interactions to form enduring relationships with the user (Manzini et al. 2024; Kirk et al. 2025).⁸

Addressing these limitations requires rethinking alignment paradigms, LLM architectures, and affordances. One way of doing so would be to implement context-aware norm weighting to allow

models to adjust norm prioritization as a function of context. Another strategy would be to allow some degree of user-driven persona customization allowing users to influence the models values or alignment parameters as a function of context (e.g., *Be 20% more honest*). Finally, though AI safety experts might balk, it may be worthwhile to allow models to negotiate alignment, say, by allowing them to propose norm exceptions and adjust accordingly (e.g., *This seems like a casual chat. Can I relax formality?*). Without innovations like these, LLMs will remain conversationally impoverished—ethically aligned yet pragmatically alien communicators—consistently out of sync with the fluidity of human communicative context.

5.3 | Summary

To summarise, we claim prompting is a distinct kind of pseech act that functions as a novel, yet problematic, substitute for the dynamic, communally-constructed context found in human-human communication. We argue that this kind of static framing imposes challenges on achieving a shared understanding and conversational alignment between users and LLMs. We also argue that behavioural alignment strategies contribute to a rigid, pre-defined communicative identity for LLMs. This contrasts with the more flexible, context-sensitive forms of self-presentation found in human-human communication.

6 | Philosophical and Practical Implications

In this article, we have explored conversational alignment through the lens of contemporary theories of context, communication, and pragmatics and identified central limitations within current LLM architectures, operational constraints, and alignment strategies. A fundamental philosophical implication of this analysis is that, despite their impressive successes in text-generation, LLMs do not yet fully qualify as human-like communicative agents in the more robust, pragmatic sense required by philosophical theories of meaning, context, and conversation. Indeed, even if such systems are ethically aligned and well fine-tuned, their pragmatic and conversational shortcomings are rife and fall well short of satisfying the desiderata outlined in the CONTEXT-ALIGN framework.

Having critically assessed some fundamental aspects of the conversational capacities of contemporary LLMs, we are now in a better position to address the key questions posed in the introduction and to summarize the philosophical significance of the foregoing. More specifically, we asked: What does it mean for a conversational agent to be aligned with human communicative practices? Do LLMs follow the norms, structures, and pragmatic expectations that govern human conversation? How do LLMs differ from human speakers in their conversational abilities? Are current LLM architectures, operational constraints, and affordances fundamentally limited in achieving full conversational alignment or can they be improved to better approximate human dialog? We now answer these questions in turn.

First, conversational alignment requires more than just the ability for text-generation; it requires a robust sensitivity to contextual parameters, the capacity to dynamically and cooperatively update

the common ground, adherence to pragmatic norms, the effective management of discourse structure, and the flexibility to adapt one's communicative identity according to situation-specific demands. Conversationally aligned agents should function not merely as passive responders, but as active, cooperative participants in the joint construction of meaning, understanding, and communicative goals.

Second, our analysis indicates that current LLMs exhibit significant deficiencies when evaluated against the conversational alignment criteria outlined in our CONTEXT-ALIGN framework. Although they can superficially approximate human conversation through learned statistical patterns, they may struggle to reliably track and update the common ground, to maintain discourse coherence by tracking QUDs, and they may often misinterpret the pragmatic force of utterances. Furthermore, the conversational norms and communicative identities imposed through static alignment strategies, such as RLHF and system prompts, lack the situated adaptability that characterizes human pragmatic competence. That is, rather than dynamically tailoring their conversational behaviors to the needs of the context, LLMs typically default to generic, inflexible personas that reflect overarching corporate or institutional values rather than context-specific requirements. Thus, although LLMs follow certain basic linguistic and ethical norms of human conversation, their adherence to deeper pragmatic norms, structures, and expectations remains limited and problematic.

Third, LLMs differ fundamentally from human interlocutors in several key dimensions of conversational competence. Humans possess genuine intentionality and meta-communicative awareness, enabling them to dynamically negotiate meaning and clarify misunderstandings through real-time feedback and conversational repair. In contrast, LLMs lack intentionality and meta-communicative competence, relying instead on static prompting and statistical inference from training data. For example, humans naturally maintain and update conversational context, common ground, and conversational scoreboards through situated pragmatic reasoning, whereas the architectural constraints of LLMs can result in significant challenges posed by the competing demands of context window overflow and context collapse, resulting in the loss of coherence and diachronic consistency over extended interactions. Furthermore, humans are capable of navigating shifts in communicative identity, norms, and registers across varying contexts, dynamically aligning their self-presentation and pragmatic norms to meet the situational demands. Conversely, LLMs that are constrained by static alignment protocols (such as HHH) exhibit rigid, context-insensitive communicative identities incapable of flexible adaptation. Thus, despite superficial linguistic fluency, LLMs differ from human speakers in their lack of genuine pragmatic reflexivity, meta-communicative negotiation, structured memory, and dynamic conversational identity.

Finally, our analysis suggests that current LLM architectures, constraints, and affordances may impose fundamental limitations on achieving full conversational alignment. The challenges of context window overflow and context collapse reveal inherent tensions between retaining sufficient context for coherence and avoiding pragmatic confusion. Similarly, static prompting as a substitute for context, combined with rigid alignment strategies,

results in pragmatic dissonance and communicative identity inflexibility. We are inclined to think that these limitations are insurmountable in principle, though they might not be. In any case, they highlight the need for substantial conceptual and architectural innovation to move beyond current constraints.

There are good reasons to put in the effort to secure these conceptual and architectural innovations. We will end by discussing two significant practical implications of failing to meet the criteria of the CONTEXT-ALIGN framework. The first is that systematic failures in conversational alignment will inevitably have a disruptive effect on fundamental aspects of human communication, many of which are unpredictable and potentially negative in ways broadly analogous to the impacts that social media has had on communication (an example we have already discussed in detail with context collapse). The second is that systematic failure to meet the CONTEXT-ALIGN framework presents a significant safety issue. Misalignments in pragmatic understanding, conversational coherence, and context-tracking can lead to misunderstanding, misinformation, and erosion of user trust, as well as potentially giving rise to harmful consequences in high-stakes domains like healthcare, emergency response, education, or legal advice.

LLMs are not the first technology to transform fundamental aspects of communication. Email replaced much formality with asynchronous brevity, whereas social media turned discourse into a performative spectacle. Yet LLMs introduce a novel disruption: synthetic interlocutors that mimic human conversation while operating with wholly alien cognition, alien semantic contents, and alien principles. Our discussion reveals three important shifts that are beginning to emerge in human communication as a result of engaging extensively with conversational AI. As discussed, there is an emerging shift from collaborative, co-constructed conversation toward unilateral, prompt-driven conversation. There is also a shift in the norms governing communicative identity and self-presentation. Finally, we observe a shift toward epistemic opacity and diminished trust in conversational interaction. For while human interlocutors typically have some insight into each other's cognitive and communicative states—say, through explicit requests for clarification, meta-communicative signaling, and pragmatic inference—LLMs, by contrast, provide little transparency into their internal reasoning, pragmatic assumptions, or contextual interpretations.

Ultimately, like in the case of social media, addressing these shifts and their implications requires urgent philosophical reflection and technical innovation. The challenge is not simply to improve LLMs' linguistic fluency, intelligibility, and ethical alignment but rather to ensure that their integration into our communicative lives respects—and ideally enhances—the pragmatic richness, flexibility, and cooperative spirit that define human conversation.

Acknowledgments

We would like to thank Michael Barnes, Melissa Fusco, Karen Lewis, Eliot Michaelson, Eleonore Neufeld, Kate Scott, and the other participants of the *Social Media and Communication* workshop held at Columbia University for helpful comments and discussion on this article. We would also like to thank Adam Bradley, Jason Gabriel, Geoff Keeling, Seth Lazar,

Fei Song, Kate Vredenburg, and audience members at *The Philosophy of AI: Themes from Iason Gabriel* workshop held at the AI & Humanity Lab, University of Hong Kong.

Endnotes

- ¹ See Millièrè and Buckner (2024), for a recent review.
- ² See Gabriel and Ghazavi (2022) and Anwar et al. (2024), for recent overviews.
- ³ For an extensive discussion of pragmatic norms and their relation to the design of conversational agents, see Kasirzadeh and Gabriel (2023). See also work on evaluating the basic pragmatic capabilities of LLMs (Srivanthi et al. 2024; Ma et al. 2025).
- ⁴ Dafoe et al. (2020, 19–20) list communication and common ground as open problems in cooperative AI. Kasirzadeh and Gabriel (2023, 6) discuss context as important for conversational alignment.
- ⁵ CA2 and CA3 presuppose the ability to identify and classify speech acts as they are performed by the user, as well as the speech acts the LLM itself is performing. A further set of benchmarks is needed to evaluate these capabilities.
- ⁶ Each of the criteria in the CONTEXT-ALIGN framework presupposes other important capabilities. They require the ability to implement a theory of mind, that is, the ability to identify, classify, and reason about users' mental states (see, a.o., Aru et al. 2023; Sap et al. 2023; Street 2024; Street et al. 2024; Ullman 2023; Q. Wang and Goel 2024; Williams et al. 2022). It also needs to use that to adjust its own states, so there will be a need for extensive higher order reasoning. This includes, importantly in the case of communication, a suite of metalinguistic capabilities and metacommunicative/metapragmatic capabilities.
- ⁷ Note that the desiderata in CONTEXT-ALIGN is not intended to be exhaustive of the ways that context can impact conversational alignment.
- ⁸ The development of AI companions is a relevant use case.

References

Anthropic. 2024. *The Claude 3 Model Family: Opus, Sonnet, Haiku*. Anthropic. <https://assets.anthropic.com/m/61e7d27f8c8f5919/original/Claude-3-Model-Card.pdf>.

Anwar, U., A. Saparov, J. Rando, et al. 2024. "Foundational Challenges in Assuring Alignment and Safety of Large Language Models." arXiv. <https://doi.org/10.48550/arXiv.2404.09932>.

Aru, J., A. Labash, O. Corcoll, and R. Vicente. 2023. "Mind the Gap: Challenges of Deep Learning Approaches to Theory of Mind." *Artificial Intelligence Review* 56, no. 9: 9141–9156. <https://doi.org/10.1007/s10462-023-10401-x>.

Askill, A., Y. Bai, A. Chen, et al. 2021. "A General Language Assistant as a Laboratory for Alignment." arXiv. <https://doi.org/10.48550/arXiv.2112.00861>.

Austin, J. L. 1962. *How to Do Things With Words*. Clarendon Press.

Bach, K., and R. M. Harnish. 1979. *Communication and Speech Acts*. Harvard University Press.

boyd, d. 2002. "Faceted Id/Entity: Managing Representation in a Digital World." PhD Thesis, MIT.

boyd, d. 2010. "Social Network Sites as Networked Publics: Affordances, Dynamics, and Implications." In *A Networked Self*, edited by Z. Papacharissi, 39–58. Routledge.

boyd, d. 2013. *How "Context Collapse" Was Coined: My Recollection*. Apophenia (blog). <https://www.zephorio.org/thoughts/archives/2013/12/08/coining-context-collapse.html>.

boyd, d. 2014. *It's Complicated: the Social Lives of Networked Teens*. Yale University Press.

Bubeck, S., V. Chandrasekaran, R. Eldan, et al. 2023. "Sparks of Artificial General Intelligence: Early Experiments with GPT-4." arXiv. <https://doi.org/10.48550/arXiv.2303.12712>.

Chevalier, A., A. Wettig, A. Ajith, and D. Chen. 2023. "Adapting Language Models to Compress Contexts." In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, edited by H. Bouamor, J. Pino, and K. Bali, 3829–3846. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.emnlp-main.232>.

Dafoe, A., E. Hughes, Y. Bachrach, et al. 2020. "Open Problems in Cooperative AI." arXiv. <https://doi.org/10.48550/arXiv.2012.08630>.

Fogal, D., and D. W. Harris. & M. Moss (eds.) 2018. *New Work on Speech Acts*. Oxford University Press.

Frost-Arnold, K. 2021. "The Epistemic Dangers of Context Collapse Online." In *Applied Epistemology*, edited by J. Lackey, 437–456. Oxford University Press. <https://doi.org/10.1093/oso/9780198833659.003.0018>.

Frost-Arnold, K. 2023. *Who Should We Be Online? A Social Epistemology for the Internet*. Oxford University Press.

Gabriel, I., and V. Ghazavi. 2022. "The Challenge of Value Alignment." In *The Oxford Handbook of Digital Ethics*, edited by C. Véliz, 336–355. Oxford University Press.

Gemini Team, G., R. Anil, S. Borgeaud, et al. 2024. "Gemini: A Family of Highly Capable Multimodal Models." arXiv. <https://doi.org/10.48550/arXiv.2312.11805>.

Glaese, A., N. McAleese, M. Trębacz, et al. 2022. "Improving Alignment of Dialogue Agents via Targeted Human Judgements." arXiv. <https://doi.org/10.48550/arXiv.2209.14375>.

Goffman, E. 1959. *The Presentation of Self in Everyday Life*. Doubleday.

Graff, D. 2000. "Shifting Sands: an Interest-Relative Theory of Vagueness." *Philosophical Topics* 28: 45–81.

Grice, H. P. 1975. "Logic and Conversation." In *Syntax and Semantics, 3: Speech Acts*, edited by P. Cole and J. L. Morgan, 41–58. Academic Press.

Horn, L. R. 1984. "Toward a New Taxonomy for Pragmatic Inference: Q-Based and R-Based Implicature." In *Meaning, Form, and Use in Context: Linguistic Applications*, edited by D. Schiffrin, 11–42. Georgetown University Press.

Kaplan, D. 1989. "Demonstratives." In *Themes From Kaplan*, edited by J. Almog, J. Perry, and H. Wettstein, 481–563. Oxford University Press.

Kasirzadeh, A., and I. Gabriel. 2023. "In Conversation With Artificial Intelligence: Aligning Language Models With Human Values." *Philosophy & Technology* 36, no. 2: 27. <https://doi.org/10.1007/s13347-023-00606-x>.

Kenton, Z., T. Everitt, L. Weidinger, I. Gabriel, V. Mikulik, and G. Irving. 2021. "Alignment of Language Agents." arXiv. <https://doi.org/10.48550/arXiv.2103.14659>.

Kennedy, C. 2007. "Vagueness and Grammar: the Semantics of Relative and Absolute Gradable Adjectives." *Linguistics and Philosophy* 30: 1–45.

Kheiri, K., and H. Karimi. 2023. "SentimentGPT: Exploiting GPT for Advanced Sentiment Analysis and Its Departure From Current Machine Learning." arXiv. <https://doi.org/10.48550/arXiv.2307.10234>.

Kirk, H. R., I. Gabriel, C. Summerfield, B. Vidgen, and S. A. Hale. 2025. "Why Human-AI Relationships Need Socioaffective Alignment." arXiv. <https://doi.org/10.48550/arXiv.2502.02528>.

Kratzer, A. 1977. "What 'Must' And 'Can' Must and Can Mean." *Linguistics and Philosophy* 1: 337–355.

Lepore, E., and M. Stone. 2015. *Imagination and Convention: Distinguishing Grammar and Inference in Language*. Oxford University Press.

Levinson, S. C. 2000. *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. MIT Press.

Lewis, D. 1969. *Convention: a Philosophical Study*. Harvard University Press.

- Lewis, D. 1970. "General Semantics." *Synthese* 22: 18–67.
- Lewis, D. 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8, no. 1: 339–359.
- Lewis, K. S. forthcoming. "Imagined Audiences and Common Ground." In *Conversations Online*, edited by P. Connolly, S. C. Goldberg, and J. Saul. Oxford University Press.
- Lewis, P., E. Perez, A. Piktus, et al. 2021. "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." *arXiv*. <https://doi.org/10.48550/arXiv.2005.11401>.
- Lewkowycz, A., A. Andreassen, D. Dohan, et al. 2022. "Solving Quantitative Reasoning Problems With Language Models." *arXiv*. <https://doi.org/10.48550/arXiv.2206.14858>.
- Lewis, D. 1973. *Counterfactuals*. Harvard University Press.
- Ma, B., Y. Li, W. Zhou, et al. 2025. "Pragmatics in the Era of Large Language Models: A Survey on Datasets, Evaluation, Opportunities and Challenges." *arXiv*. <https://doi.org/10.48550/arXiv.2502.12378>.
- Manzini, A., G. Keeling, L. Alberts, S. Vallor, M. R. Morris, and I. Gabriel. 2024. "The Code That Binds Us: Navigating the Appropriateness of Human-AI Assistant Relationships." In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 943–957. Association for the Advancement of Artificial Intelligence. <https://doi.org/10.1609/aies.v7i1.31694>.
- Marwick, A. E., and D. Boyd. 2011. "I Tweet Honestly, I Tweet Passionately: Twitter Users, Context Collapse, and the Imagined Audience." *New Media & Society* 13, no. 1: 114–133. <https://doi.org/10.1177/1461444810365313>.
- McDonald, L. forthcoming. "Context Collapse Online." In *Conversations Online*, edited by P. Connolly, S. C. Goldberg, and J. Saul. Oxford University Press.
- Meyrowitz, J. 1985. *No Sense of Place: The Impact of Electronic Media on Social Behavior*. Oxford University Press.
- Millière, R., and C. Buckner. 2024. "A Philosophical Introduction to Language Models—Part I: Continuity With Classic Debates." *arXiv*. <http://arxiv.org/abs/2401.03910>.
- Montague, R. 1968. "Pragmatics." In *Contemporary Philosophy. A Study*, edited by R. Klíbanky, 102–122. La Nuova Italia Editrice.
- Montague, R. 1970. "English as a Formal Language." In *Linguaggi Nella Società e Nella Tecnica*, edited by B. Visentini, 188–221. Edizioni di Comunità.
- OpenAI. 2023. "GPT-4 Technical Report." *arXiv*. <http://arxiv.org/abs/2303.08774>.
- Recanati, F. 1989. "The Pragmatics of What Is Said." *Mind & Language* 4, no. 4: 295–329. <https://doi.org/10.1111/j.1468-0017.1989.tb00258.x>.
- Roberts, C. 2012. "Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics." *Semantics & Pragmatics* 5: 1–69.
- Sap, M., R. LeBras, D. Fried, and Y. Choi. 2023. "Neural Theory-of-Mind? On the Limits of Social Intelligence in Large LMs." *arXiv*. <https://doi.org/10.48550/arXiv.2210.13312>.
- Savelka, J., A. Agarwal, M. An, C. Bogart, and M. Sakr. 2023. "Thrilled by Your Progress! Large Language Models (GPT-4) No Longer Struggle to Pass Assessments in Higher Education Programming Courses." In *Proceedings of the 2023 ACM Conference on International Computing Education Research V.1*, 78–92. ACM Digital Library. <https://doi.org/10.1145/3568813.3600142>.
- Schoubye, A. J., and A. Stokke. 2016. "What Is Said?." *Noûs* 50, no. 4: 759–793.
- Searle, J. R. 1969. *Speech Acts*. Cambridge University Press.
- Searle, J. R. 1980. "The Background of Meaning." In *Speech Act Theory and Pragmatics*, edited by J. R. Searle, F. Kiefer, and M. Bierwisch, 221–232. Springer Netherlands. https://doi.org/10.1007/978-94-009-8964-1_10.
- Searle, J. R. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press.
- Snell, C., D. Klein, and R. Zhong. 2022. "Learning by Distilling Context." *arXiv*. <https://doi.org/10.48550/arXiv.2209.15189>.
- Sperber, D., and D. Wilson. 1986. *Relevance: Communication and Cognition*. Blackwell.
- Sravanthi, S., M. Doshi, P. Tankala, et al. 2024. "PUB: A Pragmatics Understanding Benchmark for Assessing LLMs' Pragmatics Capabilities." In *Findings of the Association for Computational Linguistics: ACL 2024*, edited by L.-W. Ku, A. Martins, and V. Srikumar, 12075–12097. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.findings-acl.719>.
- Stalnaker, R. 1999. *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford University Press.
- Stalnaker, R. 2002. "Common Ground." *Linguistics and Philosophy* 25: 701–721.
- Stalnaker, R. 1978. "Assertion." *Syntax and Semantics* 9: 315–332.
- Stalnaker, R. 2014. *Context*. Oxford University Press.
- Stanley, J., and Z. Gendler Szabó. 2000. "On Quantifier Domain Restriction." *Mind & Language* 15: 219–261.
- Stalnaker, R. 1968. "A Theory of Conditionals." In *Studies in Logical Theory*, edited by N. Rescher, 98–112. Blackwell.
- Sterken, R. 2015. "Generics in Context." *Philosophers' Imprint* 15: 1–30.
- Street, W. 2024. "LLM Theory of Mind and Alignment: Opportunities and Risks." *arXiv*. <https://doi.org/10.48550/arXiv.2405.08154>.
- Street, W., J. Oliver Siy, G. Keeling, et al. 2024. "LLMs Achieve Adult Human Performance on Higher-Order Theory of Mind Tasks." *arXiv*. <https://doi.org/10.48550/arXiv.2405.18870>.
- Ullman, T. 2023. "Large Language Models Fail on Trivial Alterations to Theory-of-Mind Tasks." *arXiv*. <https://doi.org/10.48550/arXiv.2302.08399>.
- Vaswani, A., N. Shazeer, N. Parmar, et al. 2017. "Attention Is All You Need." In *Advances in Neural Information Processing Systems* 5998–6008.
- Wang, L., C. Lyu, T. Ji, et al. 2023. "Document-Level Machine Translation With Large Language Models." *arXiv*. <https://doi.org/10.48550/arXiv.2304.02210>.
- Wang, Q., and A. K. Goel. 2024. "Mutual Theory of Mind for Human-AI Communication." *arXiv*. <https://doi.org/10.48550/arXiv.2210.03842>.
- Wang, X., Z. Chen, Z. Xie, T. Xu, Y. He, and E. Chen. 2024. "In-Context Former: Lightning-Fast Compressing Context for Large Language Model." *arXiv*. <https://doi.org/10.48550/arXiv.2406.13618>.
- Weidinger, L., J. Uesato, M. Rauh, et al. 2022. "Taxonomy of Risks Posed by Language Models." In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 214–229. Association for Computing Machinery. <https://doi.org/10.1145/3531146.3533088>.
- Williams, J., S. M. Fiore, and F. Jentsch. 2022. "Supporting Artificial Social Intelligence With Theory of Mind." *Frontiers in Artificial Intelligence* 5: 750763. <https://doi.org/10.3389/frai.2022.750763>.
- Wingate, D., M. Shoeybi, and T. Sorensen. 2022. "Prompt Compression and Contrastive Conditioning for Controllability and Toxicity Reduction in Language Models." In *Findings of the Association for Computational Linguistics: EMNLP 2022*, edited by Y. Goldberg, Z. Kozareva, and Y. Zhang, 5621–5634. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.findings-emnlp.412>.
- Wittgenstein, L. 1953. *Philosophical Investigations*. Edited by G. E. M. Anscombe, P. M. S. Hacker, and Joachim Schulte. Wiley-Blackwell.