

Commentary on “LLMs and the Logical Space of Reasons”: Inferential Roles, Rule-Following, and Digital Speech Acts

Mirco Sambrotta^{1,*}

¹Independent researcher, Bratislava, Slovakia

*Correspondence should be addressed to Mirco Sambrotta, mirco.sambrotta@gmail.com

Received date: May 07, 2026, **Accepted date:** June 12, 2026

Citation: Sambrotta M. Commentary on “LLMs and the Logical Space of Reasons”: Inferential Roles, Rule-Following, and Digital Speech Acts. Archives of Information Technology. 2026;1(1):43–47.

Copyright: © 2026 Sambrotta M. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract

The emergence of Large Language Models (LLMs) has reignited longstanding philosophical debates concerning linguistic understanding, conceptual content, and the conditions under which a system may be regarded as a genuine language user. The article “LLMs and the Logical Space of Reasons” offers a sustained critique of claims that contemporary LLMs possess genuine understanding, arguing instead that they remain outside the normatively structured practices constitutive of discursive agency. This commentary examines and further develops that central thesis through the lens of an inferentialist account of discursive practice. Drawing on Wittgenstein’s remarks on rule-following and Brandom’s critique of regularism, it argues that current LLMs exhibit only conformity to rule-like regularities rather than genuine responsiveness to inferential norms. Although such systems generate linguistically sophisticated outputs and can often mimic inferentially appropriate responses, their performance remains grounded in statistical regularities rather than participation in practices of commitment, entitlement, and normative assessment. The commentary further explores the implications of this diagnosis for theories of assertion and speech acts. After examining competing accounts such as proxy assertion, proto-assertion, and make-believe approaches to chatbot discourse, it argues that LLM-generated utterances are best understood as simulations of assertion rather than genuine assertions. Building on this idea, the paper proposes the notion of “degenerate speech acts” as a useful conceptual category for understanding forms of communication mediated by artificial systems. Finally, the discussion situates the paper within broader debates concerning sapience, sentience, and artificial agency, suggesting that the challenge posed by LLMs is not merely technological but fundamentally philosophical. The commentary concludes that behavioral sophistication alone cannot establish genuine participation in the logical space of reasons, since discursive understanding depends on normative statuses that current LLMs do not possess.

Inferentialism and the Practice of Giving and Asking for Reasons

The recent success of Large Language Models (LLMs) has prompted a renewed philosophical interest in the nature of linguistic understanding. Their capacity to produce fluent, context-sensitive, and often strikingly coherent outputs has led some to suggest that such systems may count as genuine language users. The paper “LLMs and the Logical Space of Reasons” offers a sustained and philosophically rigorous challenge to this suggestion. Its central claim is that, despite their impressive performance, LLMs do not genuinely understand language, because they fail to participate in the normatively structured practices that constitute the *logical space of reasons*.

This commentary aims to highlight and further develop that

claim. In doing so, it both clarifies the underlying inferentialist framework and extends the paper’s conclusions to the domain of digital speech acts. Rather than merely summarizing the argument, it situates it within a broader philosophical landscape and draws out its implications, particularly with respect to the nature of assertion and the emergence of what may be called “degenerate speech acts.” The guiding idea is that the failure of LLMs to qualify as language users is not simply a matter of architectural limitation or insufficient training, but reflects a deeper absence of normative status, an absence that becomes especially visible when we turn from linguistic understanding to speech acts.

A central contribution of the paper lies in its defense of an inferentialist account of meaning. Against both representationalist and behaviorist approaches, it argues that conceptual and propositional content is constituted by

participation in socially instituted, norm-governed practices. Meaning is not an inner state, nor a mere pattern of use, but something that emerges from the interplay of normative attitudes and statuses within a linguistic community. Drawing on Robert Brandom’s framework, the paper understands conceptual content in terms of its inferential articulation within a “space of reasons” [1]. To grasp the meaning of a claim is thus to locate it within a network of inferential relations (i.e., relations of material consequence and incompatibility) and to be able to navigate that network appropriately. Crucially, these inferential roles are not merely descriptive regularities. They are sustained by practices in which participants treat one another as *committed* and *entitled*, and in which performances are assessed as correct or incorrect. This point is decisive for the overall argument. If meaning is constituted by such practices, then linguistic understanding cannot be reduced to the ability to produce well-formed or even inferentially appropriate outputs. It requires participation in a normative structure in which one’s utterances are not merely produced, but assessed, challenged, and revised [2,3].

Rule-following and the Limits of Statistical Learning

The paper’s analysis of LLMs turns on a classical distinction, drawn from Wittgenstein and developed by Brandom: the distinction between genuinely following a rule and merely acting in accordance with a rule-like pattern. This distinction allows us to resist a powerful temptation in current debates: namely, the idea that behavioral indistinguishability suffices for attributing understanding [4]. LLMs may produce outputs that are indistinguishable, in many contexts, from those of competent speakers. But this does not show that they follow rules in the relevant sense. It shows only that their behavior conforms to patterns that resemble rule-governed activity. The notion of “regularism,” as criticized by Brandom, is particularly useful here. Regularism attempts to explain rule-following entirely in terms of behavioral regularities, thereby collapsing the distinction between being bound by a rule and merely conforming to it. As such, it represents a deflationary account of normativity, one that reduces correctness to mere regularity of performance.¹

The paper argues that LLMs fall precisely into this category: their outputs exhibit regularistic adherence and systematic

conformity to inferential patterns without any grasp of their normative significance. The central claim is that statistical regularities, no matter how sophisticated, cannot amount to inferential normativity, and therefore cannot ground participation in the space of reasons. This diagnosis is further supported by the analysis of LLM training. While fine-tuning methods, such as Reinforcement Learning with Human Feedback (RLHF), Supervised Fine-Tuning (SFT), and Quantized Low-Rank Adapters (Q-LoRA), introduce elements of correction and alignment, they fall short of the kind of normative training characteristic of human language acquisition. Instead, what such systems acquire are statistical distributions and probabilistic associations, rather than the inferential norms that structure discursive practice and the capacity to be responsive to such norms as norms.

In light of this, a core part of the paper is the comparison between human language acquisition mechanisms and LLM training. Human learners acquire linguistic competence through participation in practices of sanction, correction, and guidance, in which performances are treated as right or wrong and adjusted accordingly. Crucially, to treat a performance in this way is not merely to classify it, but to take the speaker to have undertaken a commitment and to be answerable for it: open to challenge, liable to correction, and required to adjust in light of appropriate feedback. Such practices are not merely instrumental; they are constitutive of the norms that govern language use. By contrast, LLMs are trained to optimize the likelihood of word sequences in large corpora. Even when fine-tuning procedures introduce feedback mechanisms, the norms that guide learning remain fundamentally statistical. At most, such systems “learn” (and thereby “know”) which words typically follow others on the basis of distributional patterns; but even if we grant this, the norms they thereby master are of a wholly different kind from the norms of inference to which contentful items are subject. As the paper puts it, such systems capture patterns of token co-occurrence without engaging with the inferential norms that govern language use, underwrite discursive practice, and articulate the space of reasons.

This difference is not merely quantitative, but qualitative. It marks the gap between systems that are initiated into a normative practice (where performances incur commitments and are subject to assessment) and systems that are calibrated to reproduce statistical regularities. The latter may approximate the former with increasing sophistication, but they do not thereby become participants in the same kind of practice. At most, they can be understood as simulating our discursive practice without occupying a position within it. Machine language performance should therefore still be regarded as, at best, a simulation of grasping and understanding conceptual and propositional content. If genuine language use requires such grasp (understood as the capacity to navigate inferential

¹ At the opposite extreme, Brandom contrasts regularism with what he calls “regulism,” according to which rule-following requires the explicit formulation and consultation of rules at each step (1994, Ch. 1). As Wittgenstein’s rule-following considerations make clear, such a picture gives rise to an infinite regress of interpretations (1953, §§198–202). The inferentialist alternative avoids both extremes by treating rule-following as participation in practices structured by implicit normative statuses, in which what counts as correct or incorrect is determined within a social practice of giving and asking for reasons.

norms and to undertake commitments for what one says), then current LLMs remain significantly removed from counting as full-fledged language users. While future developments may narrow this gap, the present landscape continues to exhibit deep obstacles to construing these systems as genuine participants in discursive practice.

Assertion and Its Simulation

The paper’s discussion of assertion in its final section brings the argument to a decisive point. LLM-driven chatbots often produce outputs that resemble assertions: they offer reasons, respond to challenges, and even retract claims when prompted. Yet, as the paper insists, assertion is not merely a matter of producing such outputs. It is a norm-governed practice that requires the capacity to follow inferential norms and to undertake commitments [5]. At this point, the argument moves beyond general considerations about meaning and understanding to a more specific claim: LLMs are not merely non-understanding systems; they are not even genuine asserters. Their outputs may conform to the patterns of assertoric discourse, but they do not stand in the appropriate relation to the norms that constitute assertion. This conclusion is further sharpened through a careful engagement with alternative accounts.

The idea of “proxy assertion” is rejected on the grounds that, in typical LLM contexts, no clearly identifiable principal can be said to bear responsibility for the content.² Indeed, while proxy models of speech acts require a determinate agent on whose behalf the assertion is made [6], such a structure is largely absent in ordinary interactions with LLM-driven systems [7,8]. The notion of “proto-assertion” is likewise set aside. Although it offers a graded account of assertoric capacity, placing LLMs in an intermediate position analogous to early-stage language learners [8], it presupposes a partial sensitivity to inferential norms that, on the present view, is not borne out by the nature of current systems. What LLMs exhibit is not an incomplete grasp of such norms, but responsiveness to a different kind of norm altogether, namely statistical norms governing token co-occurrence.

A further illuminating perspective is provided by recent proposals that interpret chatbot discourse in terms of make-believe. On this view, interactions with LLMs can be understood as a form of prop-oriented imaginative engagement, in which the system’s outputs function as prompts for treating them *as if* they were assertions made by an agent [9,10]. This account helps explain why chatbot outputs so closely resemble genuine assertions without committing us to the claim that they are genuine instances of assertion. Rather, they are best understood as simulations embedded within a practice of make-believe, reinforcing the conclusion that LLMs do not

genuinely participate in the norm-governed space of reasons. In line with the latter, the paper suggests that LLM outputs can be best understood as “simulations of assertion.” They exhibit many of the surface features characteristic of assertoric practice (i.e., offering reasons, responding to challenges, and retracting claims when prompted), yet lack the normative standing required for genuine participation. In this respect, the view also aligns with other accounts that emphasize the limitations of output-based criteria for attribution of assertion [11], insofar as it accepts their emphasis on sanctionability and descriptive function as relevant constraints on any adequate theory of assertion. However, it departs from that framework in holding that even these strengthened conditions remain insufficient in the absence of genuine responsiveness to inferential norms.

Degenerate Speech Acts and Digital Discourse

The paper’s appeal to simulation provides a natural bridge to a broader conceptual proposal. If LLM outputs are not genuine assertions, but nevertheless resemble them in systematic ways, how should they be classified? One promising answer is to treat them as instances of “degenerate speech acts.” These can be understood as a distinctive class of digital speech acts, characteristic of LLM-mediated communication, in which the normative structure of assertion is systematically simulated without being instantiated.

From the perspective of classical speech-act theory, such performances can be interpreted as cases in which the felicity conditions for successful assertion are not fully satisfied [12]. Although the utterances exhibit many of the surface features associated with assertoric practice, there is no agent capable of undertaking the commitments and responsibilities ordinarily incurred by making an assertion. Consequently, they may be understood as a special kind of failure of illocutionary force: utterances that mimic the form of assertions but lack the normative standing required for their successful execution [13]. They are therefore not merely defective assertions, but performances of a distinctive kind of communicative act that preserve the superficial form of assertoric discourse while suspending the normative commitments that ordinarily give assertion its force.

This perspective helps to resolve a tension noted in the literature: the need to explain both why LLM outputs are not genuine assertions and why they so closely resemble them. The notion of simulation, especially when developed in terms of make-believe, allows us to understand interactions with LLMs as involving a kind of prop-oriented engagement, in which users treat outputs *as if* they were produced by a responsible agent. Such interactions can be epistemically and practical valuable, but they do not amount to genuine participation in a norm-governed practice of assertion. More broadly, the concept of degenerate speech acts points to

² Cf. Austin 1962 on speech act responsibility.

an emerging category of communicative phenomena that challenges traditional assumptions of speech-act theory, according to which illocutionary performances are necessarily tied to agents capable of undertaking commitments and bearing responsibility for what they say. As LLM-mediated communication becomes increasingly integrated into educational, professional, and social contexts, understanding these hybrid forms of discourse will become an increasingly important task for the philosophy of language and AI ethics [14,15].

Sapience, Sentience, and Normative Status

The argument acquires an additional layer of significance when considered in light of Brandom’s distinction between “sapience” and “sentience” [1]. While the paper does not frame its central thesis in these terms, its conclusions bear directly on a growing debate about whether discursive capacities can be dissociated from forms of animate life. In Brandom’s sense, sentience is not mere reliable differential responsiveness to environmental stimuli characteristic of artifacts such as thermostats and land mines, but a form of conscious awareness distinctive of embodied creatures: the kind of flexible, purposive responsiveness exhibited by animals as they navigate their surroundings in light of their needs and capacities, best understood in terms of sensitivity to affordances for action [16]. It is this practically oriented awareness that marks the domain of the sentient. Sapience, by contrast, consists in the capacity for conceptual understanding: the ability to grasp and deploy contents articulated by inferential relations, and thereby to participate in the norm-governed space of reasons.

A widespread view takes sapience to presuppose sentience, while some recent proposals suggest that the former may be instantiated independently of the latter [17]. The framework developed in the paper supports a more cautious conclusion. Even if one allows that different forms of normativity may be associated with sentience and discursivity, the crucial question is whether the latter can be instantiated in the absence of any normative accountability whatsoever. The paper’s analysis of LLMs suggests that they cannot. What is missing in these systems is not merely sentience, but the capacity to occupy the normative statuses that make participation in discursive practice possible. This brings out a key methodological point emphasized throughout the paper: neither sentience nor sapience can be inferred from output behavior alone. Just as behavioral indistinguishability does not suffice for attributing conscious awareness, it does not suffice for attributing understanding and participation in the space of reasons. The distinction between the simulation of discursive capacities and their genuine realization is also of direct relevance in the ethical domain, particularly in light of ongoing debates about the risks and ethical implications of advanced AI systems.

Conclusion

On the whole, “LLMs and the Logical Space of Reasons” offers a powerful and carefully argued case against attributing linguistic understanding to current artificial systems. Its central insight is that meaning, understanding, and assertion are inseparable from participation in normatively structured social practices. By situating LLMs’ statistical competence and behavioral sophistication, however advanced, outside those practices, the paper provides a principled explanation of why their impressive performance does not amount to genuine understanding.

The broader implication is that the integration of LLMs into communicative contexts does not simply extend existing discursive practices, but transforms them. What emerges is a hybrid space in which normatively binding and normatively inert forms of discourse coexist. Understanding this transformation requires not only technical analysis, but philosophical clarity about the nature of normativity, responsibility, and rule-following. Until artificial systems can occupy normative statuses, undertake commitments, and be held accountable within practices of reason-giving, their outputs are best understood as sophisticated simulations rather than genuine exercises of linguistic understanding. In this respect, the paper makes an important contribution to ongoing debates in philosophy of language, philosophy of mind, and philosophy of AI, while providing a valuable framework for future research on the relationship between language, normativity, and machine intelligence. It shows that the question of whether LLMs understand language cannot be settled by appeal to performance alone, but requires a deeper examination of the normative conditions under which language use becomes possible.

This commentary has sought both to clarify and to extend the article’s central insights. In particular, it has emphasized the significance of the distinction between rule-following and mere conformity to patterns, explored competing accounts of AI-generated assertion, and developed the proposal that LLM outputs may be understood as instances of degenerate speech acts. This proposal helps explain why interactions with contemporary AI systems often resemble ordinary communicative exchanges while nevertheless lacking the normative commitments that characterize genuine assertion.

References

1. Brandom R. Making it explicit: Reasoning, representing, and discursive commitment. Harvard university press; 1994.
2. Brandom R. Articulating Reasons: An Introduction to Inferentialism. Cambridge, MA: Harvard University Press; 2000.
3. Wittgenstein L. Philosophical investigations (trans: Anscombe, GEM). Oxford: Blackwell. Retrieved October. 1953;21:2015.

4. Turing AM. Computing machinery and intelligence (1950). *Mind*. 1987;59(236):33–60.
5. Brandom R. Asserting. *Noûs*. 1983 Nov 1. pp. 637–50.
6. Nickel PJ. Artificial speech and its authors. *Minds and Machines*. 2013 Nov;23(4):489–502.
7. Green M, Michel JG. What might machines mean?. *Minds and Machines*. 2022 Jun;32(2):323–38.
8. Williams I, Bayne T. Chatting with bots: AI, speech acts, and the edge of assertion. *Inquiry*. 2024 Dec 19:1–24.
9. Mallory F. Fictionalism about chatbots. *Ergo an Open Access Journal of Philosophy*. 2023 Nov 17;10.
10. Walton KL. *Mimesis as make-believe: On the foundations of the representational arts*. Harvard University Press; 1993 Oct 15.
11. Butlin P, Viebahn E. AI assertion. *Ergo: An Open Access Journal of Philosophy*. 2023.
12. Austin JL. *How to Do Things with Words*. Oxford: Oxford University Press; 1962.
13. Searle JR. *Speech acts: An essay in the philosophy of language*. Cambridge University. 1969.
14. Coeckelbergh M. *AI ethics*. MIT press; 2020 Apr 7.
15. Vallor S. *The AI mirror: How to reclaim our humanity in an age of machine thinking*. Oxford University Press; 2024 May 6.
16. Gibson JJ. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin; 1979.
17. Simonelli R. Sapience without sentience: an inferentialist approach to LLMs. *Asian J Philos*. 2026 Apr 27;5(1):48.