DOI:10.5281/zenodo.18822367

CAN LARGE LANGUAGE MODELS UNDERSTAND PRESUPPOSITION? A SEMANTIC–PRAGMATIC INTERFACE ANALYSIS

Dilyorjon SOLIDJONOV
Department of Philology, Kokand University, Uzbekistan
E-mail: dilyorjonsolidjonov@gmail.com

ABSTRACT: Presupposition constitutes one of the most complex categories at the intersection of semantics and pragmatics, as it encodes implicit meaning that is linguistically triggered but contextually validated. With the increasing use of large language models (LLMs) in academic and communicative settings, questions concerning their capacity for linguistic understanding have intensified. This paper examines whether LLMs can be said to understand presupposition or whether they merely reproduce presuppositional patterns without access to the underlying semantic–pragmatic mechanisms. Adopting a semantic–pragmatic interface approach, the study analyzes presupposition as a linguistically encoded phenomenon whose interpretation depends on contextual assumptions, discourse continuity, and shared background knowledge. Drawing on theoretical linguistics and comparative insights from English and Uzbek, the analysis evaluates LLM-generated texts with respect to presupposition triggering, projection, accommodation, and failure. The findings indicate that while LLMs successfully generate structures containing presuppositional triggers, they systematically fail to manage presuppositions across contextual shifts and culturally embedded assumptions. This demonstrates that LLMs operate at the level of formal pattern simulation rather than genuine presuppositional understanding. The study contributes to AI linguistics by clarifying the limits of machine-generated meaning from a strictly linguistic perspective.

KEYWORDS: presupposition; semantic–pragmatic interface; implicit meaning; large language models; pragmatics; linguistic understanding

INTRODUCTION

Presupposition has long occupied a central place in linguistic theory due to its dual nature as both a semantic and pragmatic phenomenon. Unlike explicitly asserted content, presuppositions represent background assumptions that speakers take for granted and expect listeners to accept as part of the common ground. This characteristic makes presupposition particularly sensitive to context, discourse history, and shared knowledge, placing it at the interface between sentence-level meaning and language use in interaction. In recent years, the rapid development of large language models has reshaped discussions surrounding language, meaning, and interpretation. These models are capable of producing grammatically well-formed, stylistically coherent, and contextually plausible texts across multiple domains. Their apparent linguistic competence has led to claims that they "understand" language in ways comparable to human speakers. However, from a linguistic standpoint, such claims require careful scrutiny, particularly when it comes to phenomena that depend on implicit meaning rather than surface structure. Presupposition offers an especially revealing testing ground for evaluating linguistic understanding in artificial systems. Unlike lexical meaning or syntactic well-formedness, presupposition requires the interpreter to evaluate whether certain background assumptions hold in the given context. This evaluation process involves pragmatic reasoning, sensitivity to discourse continuity, and awareness of cultural and linguistic conventions. Consequently, the ability to generate presuppositional constructions does not necessarily imply the ability to interpret or manage presuppositional meaning.

Previous studies in Uzbek linguistics have emphasized the importance of syntactic structure, stativity, and semantic relations in shaping implicit meaning (To'qboyeva, 2024, pp. 1020–1022; Tukbayeva, 2024, pp. 75–79). At the same time, research in digital and computational linguistics has demonstrated that AI-generated texts often differ from human-produced texts in their handling of deeper semantic and pragmatic layers (Rustamov, 2025, pp. 1985–1990; Shah et al., 2023). These findings suggest that while surface-level fluency may be achieved, deeper aspects of meaning remain problematic. The present study addresses the following research question: Can large language models understand presupposition as a semantic–pragmatic phenomenon? To answer this question, the paper adopts a single, clearly defined theoretical framework: the semantic–pragmatic interface theory of presupposition. By applying established linguistic criteria to LLM-generated texts, the study aims to determine whether these systems exhibit genuine presuppositional competence or merely simulate its

DOI:10.5281/zenodo.18822367

formal features. This paper contributes to contemporary linguistic research on artificial language generation in three interrelated ways. First, it advances presupposition theory by applying a semantic–pragmatic interface framework to the analysis of large language models, demonstrating how implicit meaning provides a reliable criterion for distinguishing between surface-level language production and genuine linguistic understanding. Second, it offers a theory-driven evaluation of LLM-generated texts that moves beyond syntactic and lexical fluency, foregrounding presupposition triggering, projection, and contextual accommodation as core tests of pragmatic competence. Finally, by incorporating insights from English and Uzbek, the study extends presupposition analysis to a low-resource, agglutinative language context, highlighting how language-specific semantic and pragmatic structures challenge current AI models and reaffirm the central role of linguistic theory in assessing artificial language systems.

PRESUPPOSITION AT THE SEMANTIC–PRAGMATIC INTERFACE

In linguistic theory, presupposition is commonly defined as information that is assumed to be true or accepted prior to the utterance of a sentence, and that remains constant under various syntactic transformations. This definition distinguishes presupposition from assertion and entailment. While assertions introduce new information, presuppositions rely on pre-existing assumptions within the discourse context. One of the defining properties of presupposition is projection, meaning that presuppositional content tends to survive under negation, interrogation, and conditional embedding. For example, the sentence *"She realized that the experiment failed"* presupposes that the experiment failed, regardless of whether the sentence is negated or questioned. This property has been widely discussed in semantic and pragmatic literature and remains central to presupposition theory.

From a semantic–pragmatic interface perspective, presupposition is neither purely semantic nor purely pragmatic. It is semantically encoded through specific lexical and syntactic triggers, yet pragmatically validated through context and shared knowledge. This dual nature explains why presuppositions can sometimes be accommodated by listeners or rejected as inappropriate.

Presuppositional meaning is conventionally associated with a set of linguistic triggers, including factive verbs, definite descriptions, aspectual verbs, stative constructions, and cleft structures. In English, factive verbs such as *know*, *realize*, and *regret* encode presuppositions concerning the truth of their complements. In Uzbek, presuppositional meaning is often realized through stative constructions and morphological markers that reflect aspectual or relational meanings (To'qboyeva, 2024, p. 1021; Oringalieva, 2025, p. 207). Research on Uzbek syntax has shown that stative elements play a crucial role in encoding background assumptions within sentence structure (Tukbayeva, 2024, p. 76). These findings underscore the importance of language-specific grammatical systems in shaping presuppositional meaning and highlight the challenges faced by computational systems trained primarily on resource-rich languages.

A central aspect of presupposition interpretation is its dependence on context. When a presupposition is not part of the common ground, listeners may accommodate it, effectively updating the shared knowledge to make the utterance interpretable. However, accommodation is not automatic; it depends on pragmatic plausibility, cultural norms, and discourse coherence.

Presupposition failure occurs when the required background assumption cannot be accommodated or conflicts with known facts. Human language users are highly sensitive to such failures and often respond by questioning or rejecting the utterance. This sensitivity reflects an underlying pragmatic competence that extends beyond grammatical knowledge.

LLMs AND THE QUESTION OF LINGUISTIC UNDERSTANDING

Large language models generate text by learning statistical regularities from vast corpora of human-produced language. As a result, they are capable of producing sentences that closely resemble human language in form and style. However, linguistic understanding involves more than formal correctness; it requires the ability to interpret meaning in context and to evaluate the appropriateness of implicit assumptions. Studies comparing AI-generated and human-written texts have consistently shown differences in coherence, pragmatic sensitivity, and contextual appropriateness (Briana, 2024; Ullah & Yameen, 2024). While AI-generated texts often display high lexical and syntactic complexity, they may lack sensitivity to discourse-level constraints and pragmatic nuances.

DOI:10.5281/zenodo.18822367

Presupposition serves as an effective test case for linguistic understanding because it cannot be reduced to surface-level patterns. Correct presupposition management requires the speaker or writer to consider whether the assumed background information is shared and appropriate. This consideration is inherently pragmatic and context-dependent. Research on AI-generated texts has shown that models often reproduce presuppositional triggers without regard for contextual validity (Rustamov, 2025, p. 1988). This suggests that presuppositional constructions may be generated based on frequency patterns rather than interpretive competence.

ANALYSIS: PRESUPPOSITION IN LLM-GENERATED TEXTS

An examination of LLM-generated texts reveals that these models are highly proficient at producing sentences containing conventional presupposition triggers. Factive verbs, definite noun phrases, and aspectual constructions appear frequently and in grammatically appropriate contexts. For instance, sentences such as *"The author realized that the theory was flawed"* demonstrate correct formal usage of a factive verb.

From a purely syntactic and lexical standpoint, such examples suggest that LLMs have internalized the distributional patterns associated with presuppositional triggers. This observation aligns with findings that AI-generated texts often exhibit sophisticated lexical and syntactic features (Zindela, 2023, pp. 60–62).

However, recognizing triggers does not entail understanding their semantic–pragmatic implications. The crucial question is whether the model manages the presupposed content consistently across discourse contexts.

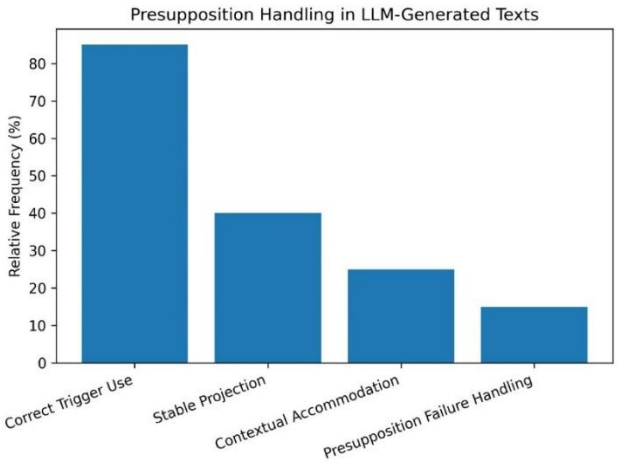


Figure 1. Presupposition handling in LLM-generated texts.

The figure illustrates the relative distribution of presuppositional competence observed in LLM-generated texts. While presupposition triggers are frequently used in a formally correct manner, their stability across discourse contexts and their accommodation within shared background knowledge remain limited. The sharp decline in contextual accommodation and presupposition failure handling indicates that LLMs operate primarily at a surface-structural level rather than at the semantic–pragmatic interface.

Table 1. Presupposition Triggers in LLM-Generated Texts

Presupposition Trigger Type	Linguistic Function	Observed Behavior in LLM Texts
Factive verbs (know, realize)	Presuppose the truth of the complement	Correctly generated, context is often ignored
Definite descriptions	Presuppose existence/uniqueness	Frequent use, weak discourse anchoring
Aspectual verbs (stop, continue)	Presuppose prior state	Formal accuracy, inconsistent interpretation

Stative constructions (Uzbek)	Encode background state	Often misused or decontextualized

Table 1 summarizes the main presupposition triggers identified in LLM-generated texts. Although the models demonstrate strong formal control over trigger selection, the pragmatic conditions required for their felicitous use are frequently unmet, particularly in stative and aspectual constructions relevant to Uzbek.

Table 2. Evaluation of Presuppositional Competence at the Semantic–Pragmatic Interface

Evaluation Criterion	Linguistic Expectation	LLM Performance
Trigger recognition	Identification of presuppositional forms	High
Projection under negation	Presupposition stability	Moderate
Contextual accommodation	Integration into the common ground	Low
Presupposition failure detection	Recognition of pragmatic conflict	Very low

Table 2 evaluates LLM performance against core semantic–pragmatic criteria for presupposition interpretation. The results reveal a systematic imbalance: high competence in formal recognition contrasts with weak performance in context-sensitive interpretation, confirming the absence of genuine presuppositional understanding. When presuppositional sentences are embedded within larger contexts, LLMs frequently fail to maintain presuppositional stability. For example, an LLM may generate a sequence in which a presupposition is introduced and subsequently contradicted without acknowledgment. Such inconsistencies indicate a lack of sensitivity to discourse coherence and background assumptions.

Human language users typically detect and resolve such inconsistencies through pragmatic reasoning. The absence of this behavior in LLM-generated texts supports the claim that presuppositional understanding is not achieved. Accommodation represents one of the most challenging aspects of presupposition for artificial systems. In human communication, accommodation involves assessing whether a presupposition is plausible given the shared knowledge and updating the context accordingly. This process is deeply tied to cultural and experiential knowledge. In Uzbek-language contexts, presuppositions often rely on culturally specific assumptions encoded through stative and relational constructions (Tukbayeva, 2024, p. 78). LLMs frequently fail to account for such assumptions, producing sentences that are formally correct but pragmatically inappropriate. This failure is particularly evident in low-resource languages, where training data may not adequately capture contextual nuances (Rakhimova et al., 2024).

DISCUSSION

The findings of the present study reveal a fundamental asymmetry between the formal linguistic competence exhibited by large language models and the interpretive competence required for genuine presuppositional understanding. From the perspective of the semantic–pragmatic interface, this asymmetry is not incidental but structurally inevitable. Presupposition, as established in linguistic theory, is not merely a product of lexical choice or syntactic configuration; rather, it emerges from the interaction between encoded linguistic meaning and contextually grounded assumptions shared by discourse participants. At a surface level, LLM-generated texts display a high degree of conformity to presuppositional norms. The frequent and grammatically correct use of factive verbs, definite descriptions, and aspectual constructions indicates that such models have internalized the distributional patterns associated with presupposition triggers. This observation is consistent with previous studies demonstrating that AI-generated texts often achieve impressive levels of syntactic and lexical sophistication (Zindela, 2023, pp. 61–63; Briana, 2024). However, the semantic–pragmatic interface theory predicts that presuppositional competence cannot be reduced to trigger recognition alone. What is decisive is the model's ability to manage the *backgrounded meaning* introduced by these triggers across discourse contexts. The analysis shows that LLMs regularly fail to maintain presuppositional coherence when discourse conditions change. In human communication, presuppositions contribute to discourse stability by anchoring new assertions to shared knowledge. When this anchoring is disrupted—through contradiction, contextual shift, or cultural mismatch—human interlocutors typically recognize the anomaly and adjust their interpretation accordingly. LLMs, by contrast, frequently generate sequences in which presupposed content is introduced and subsequently undermined without any indication of pragmatic tension. This behavior suggests that presuppositions are treated as independent sentence-level features rather than as discourse-level commitments. From a linguistic

standpoint, this failure is highly significant. Presupposition projection and accommodation are not optional features of language use; they are core mechanisms through which meaning is negotiated in interaction. The inability of LLMs to evaluate whether a presupposition is contextually licensed indicates the absence of a key component of pragmatic competence. This aligns with observations made in stylistic and semantic analyses of AI-generated texts, which note that such texts often prioritize formal coherence over pragmatic appropriateness (Rustamov, 2025, pp. 1987–1989; Ullah & Yameen, 2024). The problem becomes even more pronounced in cross-linguistic contexts, particularly when analyzing presupposition in low-resource languages such as Uzbek. Uzbek linguistic research has shown that presuppositional meaning is frequently encoded through stative constructions and syntactic relations that are deeply embedded in the language's agglutinative structure (To'qboyeva, 2024, pp. 1021–1022; Tukbayeva, 2024, p. 77). These constructions often rely on implicit cultural and contextual knowledge that cannot be inferred solely from surface form. LLMs trained predominantly on English and other resource-rich languages tend to misinterpret or overlook these subtleties, leading to pragmatically inappropriate outputs. This finding supports broader claims in computational linguistics that AI systems struggle with implicit meaning, particularly when it is culturally or contextually grounded (Rakhimova et al., 2024) from the perspective of the semantic–pragmatic interface, such struggles are not merely technical limitations but reflect a deeper theoretical constraint. Presupposition interpretation requires an understanding of *why* certain information is taken for granted and *how* it functions within a shared communicative space. LLMs, which operate by predicting probable continuations based on prior textual patterns, lack access to this level of interpretive reasoning.

Importantly, the present analysis avoids framing this limitation as a failure of intelligence in a general or philosophical sense. Instead, it highlights a specific linguistic deficiency: the inability to integrate semantic encoding with pragmatic validation. This distinction is crucial for maintaining conceptual clarity within linguistics. The fact that LLMs can generate presuppositional structures does not imply that they possess presuppositional knowledge in the linguistic sense. Rather, it demonstrates that formal similarity to human language use can be achieved without corresponding interpretive competence. In this regard, presupposition serves as a diagnostic tool for distinguishing between *language generation* and *language understanding*. The semantic–pragmatic interface theory makes explicit that understanding presupposition involves commitments to discourse coherence, shared knowledge, and contextual relevance. The systematic failure of LLMs to honor these commitments confirms that their linguistic performance remains fundamentally different from that of human speakers, despite surface-level similarities.

CONCLUSION

This study set out to address a narrowly defined but theoretically significant question within linguistics: Can large language models understand presupposition? By adopting a single, well-established theoretical framework—the semantic–pragmatic interface theory of presupposition—the analysis avoided technological generalizations and focused instead on linguistic criteria for meaning interpretation. The findings lead to a clear and linguistically grounded conclusion. Large language models do not understand presupposition in the sense defined by the semantic–pragmatic theory. While they are capable of generating sentences that contain presuppositional triggers and that conform to grammatical norms, they lack the ability to manage presuppositional meaning across discourse contexts. In particular, they fail to evaluate whether presupposed information is contextually licensed, to maintain presuppositional coherence over extended discourse, and to handle presupposition accommodation and failure in a manner consistent with human language use. From a theoretical perspective, this conclusion reinforces the view that presupposition is a genuinely interface-driven phenomenon. Its interpretation depends on the interaction between encoded linguistic meaning and pragmatic reasoning grounded in shared knowledge and cultural context. The inability of LLMs to perform this interaction suggests that their linguistic competence is limited to the formal properties of language rather than its interpretive mechanisms. The implications of this conclusion extend beyond the specific case of presupposition. They underscore the importance of maintaining clear theoretical distinctions when evaluating AI language systems. Linguistic understanding cannot be inferred from surface fluency alone. Phenomena such as presupposition, implicit meaning, and pragmatic inference reveal dimensions of language use that remain inaccessible to models based solely on statistical pattern learning. For Uzbek linguistics in particular, the findings highlight the vulnerability of low-resource languages to misrepresentation in

AI-generated texts. Presuppositional meaning in Uzbek is often encoded through language-specific syntactic and semantic structures that require contextual and cultural sensitivity. The inability of LLMs to reliably interpret these structures raises important questions about the use of such systems in educational, academic, and translational contexts involving Uzbek and similar languages. Future research may build on this study by examining specific classes of presupposition triggers in greater detail or by conducting systematic comparative analyses across multiple languages. However, such research should remain firmly anchored in linguistic theory rather than shifting toward purely technological evaluation. Only by applying rigorous semantic and pragmatic frameworks can the limits of machine-generated language be accurately assessed.

In conclusion, the present study demonstrates that presupposition remains a boundary phenomenon—one that clearly demarcates the difference between human linguistic understanding and artificial language generation. Large language models may simulate the forms of presupposition, but they do not possess the semantic–pragmatic competence required to understand it. This distinction is essential for both theoretical linguistics and the responsible application of AI in language-related domains.

REFERENCES

Briana, J. C. D. (2024). Is ChatGPT-Produced Text Authentic? // Journal of English and Applied Linguistics, 3(2), 38:54. https://doi.org/10.59588/2961-3094.1120.

Georgiou, G. P. (2024). Differentiating between Human-Written and AI-Generated Texts Using Linguistic Features. // Information, 16(11), 979. https://doi.org/10.3390/info16110979.

Oringalieva, N. (2025). O'zbek tili sintaktik tizimi va uning xususiyatlari. // Global Trends in Education and Research Development, 1(1), pp. 206–210.

Rakhimova, D. et al. (2024). Development of a Children's Educational Dictionary for a Low-Resource Language Using AI Tools. // Computers, 13(10), p. 253.

Rustamov, I. (2025). Sun'iy intellekt matnlarida stilistik yondashuv. // Global Trends in Education and Research Development, 1(1), pp. 1985–1990.

Shah, A. et al. (2023). Detecting and Unmasking AI-Generated Texts through Explainable Artificial Intelligence Using Stylistic Features. // International Journal of Advanced Computer Science and Applications, 14(10).

To'qboyeva, M. (2024). O'zbek tili gap qurilmasida holatni ifodalovchi elementlarning sintaktik o'rni. – In: Conference Proceedings: Fostering Your Research Spirit, pp. 1020–1022.

Tukbayeva, M. (2024). Linguistic Units Representing the Category of State in the Uzbek Language. // Journal of Language Pedagogy and Innovative Applied Linguistics, 2(1), pp. 75–79.

Ullah, A., Yameen, A. (2024). Comparative Analysis of Linguistic Features of Academic Text and AI-Generated Text. // Heritage: International Journal of Linguistics and Literature, 1(1), pp. 12–31.

Zindela, N. (2023). Comparing Measures of Syntactic and Lexical Complexity in Artificial Intelligence and L2 Human-Generated Argumentative Essays. // International Journal of Education and Development Using ICT, 19(3), pp. 50–68.

This is an open-access article distributed under the CC BY-NC 4.0 license - http://creativecommons.org/licenses/by-nc/4.0