

Literal Words, Hidden Worlds: Implicature and Context in Generative AI Communication

Minhas Rahim

Lecturer in English at Sarhad University of Science and Information Technology, Khyber Pakhtunkhwa (KPK), Pakistan, Peshawar, Email: minhas.english@suit.edu.pk

Nosheena Farwa Iqbal

MPhil English Linguistics Scholar, Department of English at Abdul Wali Khan University Mardan, (AWKUM) Khyber Pakhtunkhwa (KPK), Pakistan.
Email: nosheenafarwahiqbal@gmail.com

Fazal Ghufra

Lecturer in English at University of Shangla, Alpurai, Shangla, Khyber Pakhtunkhwa (KPK), Pakistan. Email: fazalghufra@gmail.com

Kainat*

Lecturer in English, Department of English and Oriental Languages, Kalam Bibi International Women Institute Bannu. Email: mughalkainat1997@gmail.com

Abstract

Generative artificial intelligence (GenAI) is increasingly used in education, professional communication, and everyday life. Although these systems can produce fluent responses, fluency does not always show that they have understood what a speaker intends. This paper examines conversational implicature: meaning that is suggested rather than directly expressed. Guided by Grice's (1975) theory of conversation, the study focuses on how a generative AI system distinguishes literal meaning from implied meaning and how its interpretations may change when linguistic, situational, and social context changes. A qualitative comparative design is proposed. Approximately 20 context-dependent utterances representing indirect requests, scalar implicatures, polite refusals, and implied criticism are to be presented to one selected language model under three contextual conditions, producing about 60 responses. The responses are to be examined using four categories: accurate literal interpretation, contextually supported implicature, ambiguous interpretation, and unsupported inference. Previous studies indicate that language models can make some pragmatic inferences, but performance differs across models, phenomena, contexts, and evaluation methods (Cho & Kim, 2024; Cong, 2024; Ma et al., 2025; Wagner et al., 2025). The paper addresses the need for a focused comparison that examines literal and implied meaning together while context is varied systematically. Since the actual response dataset is not included in this document, the discussion and outcomes provide an analytical framework rather than claims of completed experimental findings. The study may help researchers and educators evaluate AI-generated interpretations more carefully and recognise when additional context is needed.

Keywords: Generative AI, Conversational Implicature, Context, Pragmatics, Literal Meaning, Implied Meaning

Introduction

What happens when an AI system understands the words in a sentence but misses what the speaker actually means? Imagine that a student asks a chatbot, “Can I submit my assignment tomorrow?” and the chatbot replies, “The deadline is today.” The reply gives information about the deadline, but in this situation it also indirectly tells the student to submit the assignment today. This difference between what words directly say and what a speaker indirectly communicates is known as conversational implicature. Grice (1975) explains that people understand conversation by considering not only the words but also the context and the expectations that guide communication.

Generative artificial intelligence (GenAI), especially large language models (LLMs), is now used to answer questions, explain ideas, draft messages, and support learning. These systems often produce fluent and well-organised responses. However, successful communication requires more than correct grammar. It may also require understanding indirect requests, polite refusals, implied criticism, and meanings shaped by relationships or shared knowledge. Pragmatics—the study of meaning in context—therefore offers a useful way to examine human–AI communication.

For example, the sentence “Some students completed the assignment” directly states that at least some students completed it. In certain contexts, it may suggest that not all students completed it. Yet that suggestion is not logically guaranteed: the speaker might use “some” even if everyone completed the assignment. An AI system must therefore distinguish what is stated from what is reasonably inferred. This is important because a response can sound natural while still misunderstanding the speaker’s intention.

Recent studies have examined different aspects of this issue. Cho and Kim (2024) investigated scalar implicature and found that models differed in their responses to contextual cues. Cong (2024) studied manner implicatures and reported limits in the models’ handling of complex pragmatic distinctions. Andersson and McIntyre (2025) examined ChatGPT’s recognition of impoliteness and identified difficulties in context-sensitive cases. Wagner et al. (2025) also reported limitations in conversational pragmatics involving everyday reasoning and cultural expectations. Together, these studies suggest that AI systems can make some pragmatic inferences, but their performance should not be assumed to be consistent in every situation.

Further focused investigation is needed into how literal meaning, implied meaning, and context interact. In particular, a model may interpret an expression differently when information about the preceding conversation, the immediate situation, or the relationship between speakers changes. A useful evaluation must distinguish a correct interpretation from a merely possible interpretation and from an inference that is not supported by the context.

This issue is relevant to Pakistan, where educational and professional communication may involve English alongside Urdu and regional languages. Students and teachers may use AI to interpret feedback, write emails, or understand indirect instructions. However, Pakistan includes diverse language communities and communication practices; findings about one group cannot automatically be applied to everyone. The Pakistani setting is therefore a relevant area for further study, not a basis for assuming that all users communicate in the same way.

The present paper examines how generative AI interprets selected context-dependent utterances and how changes in context influence those interpretations. It is guided by Grice's (1975) theory of conversational implicature and focuses on two connected objectives: identifying whether the model distinguishes literal from implied meaning, and examining whether linguistic, situational, and social context affect the accuracy and consistency of its interpretations. In this way, the study connects established pragmatic theory with current research on AI-mediated communication.

Background of the Study

Literal Meaning, Implicature, and Context

Literal meaning is the direct meaning expressed by the words and sentence structure. Conversational implicature is an additional meaning that a speaker communicates indirectly. Context includes the information that helps a listener understand an utterance, such as the previous conversation, the situation, and the relationship between the speakers. Grice (1975) explains that people often infer unstated meanings because they assume that conversation is generally cooperative.

Grice describes four conversational maxims. Quantity concerns giving an appropriate amount of information; quality concerns saying what one has reason to believe is true; relation concerns being relevant; and manner concerns communicating clearly. These maxims are not strict rules that speakers always follow. Rather, they help explain how listeners infer meaning when an utterance seems incomplete, indirect, or unusual. For example, if a teacher asks whether an assignment is finished and a student replies, "I have written the introduction," the teacher may infer that the remaining sections are not finished. However, if the teacher asks specifically about the introduction, the same response may be a complete answer. Therefore, an utterance should not be interpreted through its words alone. The question being discussed and the surrounding circumstances also matter.

Generative AI and Pragmatic Communication

GenAI systems produce text in response to prompts and conversational input. Their grammatical fluency does not, by itself, prove that they have correctly interpreted indirect meaning. Pragmatic competence involves understanding what a person is trying to do with language in a particular situation. This may include recognising a request that is phrased as a statement, a refusal that is expressed politely, or criticism that is communicated indirectly.

Cho and Kim (2024) studied scalar implicature, including the possible inference that "some" means "some but not all." Cong (2024) examined manner implicatures and the relationship between linguistic form and inferred meaning. Andersson and McIntyre (2025) investigated whether ChatGPT could recognise impoliteness communicated indirectly. These studies examine different parts of pragmatic communication, so their findings should not be treated as if they measured exactly the same ability. They nevertheless show why AI interpretation needs to be tested beyond grammar and fluency.

Context, Culture, and Human–AI Interaction

Context can affect whether a sentence functions as a statement, request, criticism, refusal, or expression of politeness. For instance, “That is an interesting suggestion” may be sincere praise in one situation and indirect disagreement in another. The words alone do not establish which interpretation is intended. An AI system that chooses one meaning without considering the available evidence may produce a plausible but inaccurate interpretation.

Cultural and linguistic diversity makes the task more complex. People may rely on shared social practices or indirect expressions that are not fully explained by the words in a message. Wagner et al. (2025) examined ChatGPT’s conversational pragmatics in ethical, everyday, and culturally situated scenarios, including perspectives associated with Pakistani Muslim, Hindu, and Chinese participants. Their findings point to limitations in the handling of context and cultural expectations. However, the findings apply to the tasks studied and should not be generalised to every person or community.

In Pakistani education, AI may be used to interpret teachers’ feedback, understand instructions, or draft formal communication. Studying implied meaning may help identify when an AI response is useful and when the user needs to provide more context. Any claim about Pakistani learners specifically, however, requires data that directly includes those learners or locally grounded examples.

Theoretical Direction

This study uses Grice’s (1975) account of conversational implicature as its main theoretical framework. It considers literal meaning, the possible implied meaning, and the contextual evidence that supports an interpretation. The framework does not assume that every utterance has only one correct interpretation. Instead, it allows the analysis to identify reasonable inferences, alternative readings, and unsupported assumptions.

Research Objectives

The study has two objectives:

1. To examine how generative AI interprets conversational implicatures in selected context-dependent utterances, distinguishing literal meanings from implied meanings.
2. To investigate how changes in linguistic, situational, and social context influence the accuracy and consistency of generative AI’s interpretations of conversational implicatures.

Research Questions

The following questions correspond directly to the objectives:

1. How does generative AI interpret conversational implicatures in selected context-dependent utterances, and to what extent do its interpretations distinguish literal meanings from implied meanings?
2. How do changes in linguistic, situational, and social context influence the accuracy and consistency of generative AI’s interpretations of conversational implicatures?

Research Problem

Generative AI systems can produce fluent and coherent responses, but fluency does not necessarily show that they have understood implied meaning. Everyday communication often depends on indirect expressions, shared assumptions, conversational expectations, and situational information. As a result, a statement may communicate more than its literal wording.

Previous studies have investigated scalar implicature, manner implicature, impoliteness, and other aspects of pragmatic competence. However, further focused analysis is needed to examine how the same utterance is interpreted when contextual information changes. The central problem addressed by this study is how generative AI interprets conversational implicatures and how linguistic, situational, and social context influence the accuracy and consistency of those interpretations.

Significance of the Study

This study connects pragmatic theory with the growing use of GenAI in education, professional communication, and everyday interaction. It contributes to pragmatics by examining how literal and implied meanings relate to context. It also contributes to AI communication research by considering situations in which a fluent response may not reflect the speaker's intended meaning.

The study may help teachers and students use AI more critically, particularly when interpreting indirect instructions, feedback, requests, and polite responses. It may also provide a starting point for examining AI-mediated communication in Pakistan. However, conclusions about Pakistani classrooms or multilingual users will be justified only if the dataset includes relevant users or locally grounded examples. The study therefore supports careful evaluation without assuming that one set of examples represents all speakers.

Limitations of the Study

The study has several limitations. First, its findings will depend on the selected AI model, the prompts, and the types of implicature included. Second, an utterance may have more than one reasonable interpretation, so evaluation requires clear criteria and, where possible, independent human judgments. Third, written prompts cannot fully represent real conversation, which may include tone, facial expressions, gestures, and shared personal experience.

AI responses may also vary across model versions, sessions, or prompt wording. Finally, if the study uses only English examples, its results cannot automatically be generalised to Urdu, regional languages, or all Pakistani communication. These limitations will be stated when interpreting the results.

Originality and Novelty

The study does not claim to be the first investigation of implicature in AI. Its contribution lies in a focused comparison of the same utterances under systematically varied contextual conditions. For example, an expression may be presented in a friendly conversation, a formal academic exchange, and a disagreement. The model's interpretations can then be compared using the same

evaluation criteria. The analysis distinguishes three central questions: whether the model identifies the literal meaning, whether it identifies an implicature supported by the context, and whether it introduces an unsupported interpretation. This approach provides a clearer account of pragmatic performance than simply deciding whether a response sounds natural. If Pakistani English or multilingual examples are included, they may add a local contribution, but this claim must be supported by the data actually collected.

Literature Review

Theoretical Foundations of Conversational Implicature

Grice's (1975) theory provides the main explanation for how listeners derive meanings that extend beyond literal expressions. The cooperative principle and its maxims help listeners interpret indirect statements by considering whether the speaker's contribution is relevant, sufficiently informative, credible, and clear. For example, when someone responds to an invitation by saying, "I have an examination tomorrow," the listener may infer that the person cannot or does not wish to attend. Yet the inference depends on the situation; the sentence does not explicitly state a refusal.

Grice's framework does not mean that every indirect expression has one fixed interpretation. Instead, the meaning arises from the relationship between the utterance, conversational expectations, and context. This is important for AI research because a response can be grammatically correct and logically possible while still failing to capture the intended message. The present study therefore treats pragmatic interpretation as context-dependent rather than as a simple search for dictionary meanings.

Pragmatic Inference in Generative AI

Cho and Kim (2024) examined scalar implicature in BERT and GPT-2. Their study focused on expressions such as "some," which can suggest "some but not all." They found evidence of scalar inference in the tested conditions, but the models differed when contextual questions were introduced. This suggests that the presence of a plausible inference does not guarantee consistent interpretation across contexts.

Cong (2024) investigated manner implicatures using probability-based measures, embedding similarities, and text-generation approaches. The study found some sensitivity to pragmatic distinctions but also limits in capturing their full complexity. Its methods also highlight an important point: evaluations based on generated responses, token probabilities, or semantic similarity may measure different aspects of performance. Therefore, results from different methods should not automatically be treated as equivalent.

Together, these studies suggest that pragmatic inference in GenAI is neither completely absent nor consistently reliable. Performance may depend on the model, the linguistic phenomenon, the context, and the evaluation method. Scalar and manner implicatures are important but do not cover all forms of indirect requests, implied criticism, polite refusal, and socially situated language. The present study addresses this issue by analysing selected utterances and separating literal meaning from contextually supported implications.

Context, Conversation, and Culture

Context influences whether an utterance communicates agreement, refusal, criticism, or politeness. “That is an interesting suggestion,” for example, may express genuine appreciation or indirect disagreement. The preceding conversation, the communicative purpose, and the relationship between participants help determine the interpretation. For AI systems, the challenge is to recognise when a contextual change justifies a different reading.

Wagner et al. (2025) examined ChatGPT through interactive exchanges involving ethical questions, everyday reasoning, and culturally situated scenarios. Their qualitative analysis, informed by Grice’s maxims, identified limitations in the model’s handling of relevance, informativeness, contextual information, and cultural expectations. The study is useful because it considers interaction rather than only isolated sentences. However, its selected scenarios do not establish how every AI system will perform across all cultural groups. Ma et al. (2025) reviewed datasets and evaluation methods for pragmatic processing in LLMs. Their survey highlights continuing challenges in defining and assessing pragmatic ability. This complements individual experiments: a controlled test can reveal how one model handles one phenomenon, while a survey can identify broader methodological issues. A focused comparison of the same utterance under different contextual conditions is therefore a useful way to examine how context and implied meaning interact.

Hedging, Language Education, and Ghufraan’s Research

Hedging is the use of expressions such as “may,” “might,” “appears to,” and “suggests that” to qualify a claim. It is not the same as conversational implicature, but it shows how language choices work within a rhetorical and social context. Ghufraan, Durrani, and Khan (2025) examined hedging in academic writing by native and non-native English writers through a corpus-based comparison. Their work reported differences in the frequency, distribution, and rhetorical functions of hedging devices. It is relevant here because it highlights the role of context in understanding linguistic choices; however, it does not directly test AI’s interpretation of implicature.

Ghufraan, Ali, Khan, and Irshad (2026) examined task-based language teaching and speaking fluency among ESL learners using a quasi-experimental design. Speaking fluency and pragmatic interpretation are different abilities. A learner may speak fluently without understanding an indirect request, just as an AI system may produce fluent language without identifying an implied meaning. This work provides educational context for the importance of communicative ability but should not be used as direct evidence of AI pragmatic performance.

These studies support the broader importance of context-sensitive language analysis while keeping the present study’s focus clear. Research on human academic writing or learner fluency cannot replace direct evaluation of AI-generated interpretations. The central concern here remains whether an AI model distinguishes literal content from implied meaning and responds appropriately when context changes.

Methodological Issues in Existing Research

The studies reviewed use different methods. Cho and Kim (2024) used computational experiments, Cong (2024) compared probability-based, embedding, and generation methods, Wagner et al. (2025) analysed interactive conversations qualitatively, and Ma et al. (2025) surveyed datasets and evaluation approaches. These methods provide different kinds of evidence. Computational measures can support systematic comparison but may not capture every social nuance; qualitative analysis can reveal contextual detail but may be limited by the number and type of scenarios.

The present study uses a small, controlled qualitative comparison. It holds the target wording constant while varying contextual information and then examines the resulting interpretations using predefined categories. This design aligns the method with the two research objectives: to examine how the model interprets implicatures and to investigate how context affects accuracy and consistency.

Research Gap

The literature establishes three points. First, conversational implicature depends on conversational expectations and context (Grice, 1975). Second, LLMs can produce pragmatic inferences, but performance varies across phenomena, models, and evaluation conditions (Cho & Kim, 2024; Cong, 2024). Third, research on extended conversations and cultural sensitivity identifies limitations in contextually situated interpretation, while broader reviews identify continuing challenges in pragmatic evaluation (Ma et al., 2025; Wagner et al., 2025).

The gap is not a complete absence of research on AI pragmatics. Rather, the reviewed studies use different target phenomena, methods, contextual designs, and evaluation criteria. A focused comparison is needed to connect the distinction between literal and implied meaning with systematic variation in linguistic, situational, and social context. A model may recognise an implication in one setting but interpret the same expression differently when relevant circumstances change.

To address this gap, the present study examines selected context-dependent utterances and evaluates literal meaning, possible implied meaning, and contextual support together. The first objective examines whether the model distinguishes explicit content from implied content. The second compares interpretations across contextual conditions to assess accuracy and consistency. This alignment between theory, questions, and method allows the study to make a specific contribution without overstating what earlier research has or has not done.

Research Methodology

A qualitative comparative methodology is used to examine how GenAI interprets conversational implicatures and how contextual variation influences its interpretations. An interpretivist philosophy is adopted because meaning is understood in relation to linguistic, situational, and social context. An inductive approach is used to identify patterns in the responses, while Grice's (1975) theory guides the analysis.

A controlled, prompt-based design is employed. Approximately 20 utterances representing indirect requests, scalar implicatures, polite refusals, and implied criticism are selected purposively. Each utterance is presented under three contextual conditions, giving approximately 60 responses. The population is text-generating GenAI systems capable of responding to conversational prompts, while one accessible LLM is selected for the study. The same model version and prompt format are maintained wherever possible.

Data are collected through standardised written prompts, and responses are recorded in a structured dataset. The wording of each target utterance is kept constant while the linguistic, situational, or social information changes. This procedure makes it possible to compare how the model interprets the same expression under different conditions.

The responses are analysed thematically and comparatively. Each response is coded as (1) an accurate literal interpretation, (2) a contextually supported implicature, (3) an ambiguous interpretation, or (4) an unsupported inference. The analysis then compares these categories across contexts to identify patterns in accuracy and consistency. The coding criteria are applied throughout the dataset, and an audit trail is kept. Where feasible, a second coder reviews a subset of responses and disagreements are discussed. Independent human judgments may also be used to help assess interpretations.

Only public linguistic examples or researcher-created prompts are used, and no personal or confidential information is collected. AI-generated material is documented accurately, and limitations are acknowledged to avoid unsupported generalisations. The method is limited by the selected model, the small sample, and the text-only format. Therefore, the results will not be treated as representative of every GenAI system, language, or real-world conversation.

Discussion and Analysis

This chapter's analysis is organised around the two research objectives. It examines the distinction between literal and implied meaning and considers how changes in linguistic, situational, and social context may influence an AI model's interpretation. Because the actual response dataset is not included in the supplied paper, the examples below demonstrate how responses can be analysed; they are not presented as completed experimental findings.

Literal Meaning and Implied Meaning

Literal meaning is what an utterance directly states, while conversational implicature is what it may communicate indirectly in context. Grice's (1975) framework provides a basis for examining this difference. Consider the sentence, "Some students passed the examination." Its literal meaning is that at least some students passed. It may imply that not all students passed, but this is not a logical certainty. If an AI model presents "not all" as an unavoidable fact, it does not clearly distinguish a common inference from what the sentence strictly entails.

Cho and Kim (2024) found that the models they examined showed evidence of scalar inference but differed in how they handled contextual cues. This supports the need to examine both the interpretation and the explanation behind it. A response should be considered contextually

appropriate when it identifies a plausible implication and recognises the limits of the available evidence. It should not be judged only by whether it supplies a familiar interpretation.

Linguistic Context

Linguistic context includes the surrounding words and previous turns in a conversation. Consider, “The presentation was interesting.” If the speaker responds enthusiastically to a question about enjoyment, the sentence may express genuine interest. If the speaker responds hesitantly to a question about clarity, it may suggest indirect criticism. Still, the word “interesting” does not always signal criticism; the interpretation needs support from the surrounding discourse.

The study can compare a model’s response to the utterance alone with its response when the preceding conversation is supplied. The analysis should ask whether the literal meaning remains stable, whether the implied meaning changes appropriately, and whether the model explains its interpretation using contextual evidence. This comparison directly addresses the first research question by examining how the model separates literal meaning from implied meaning.

Situational Context

Situational context includes the circumstances and practical purpose of an exchange. For example, “The window is open” may simply describe a room. If the room is cold and someone is trying to make it warmer, the same statement may function as an indirect request to close the window. However, the request should not be assumed unless the situation supports it.

Cong (2024) examined manner implicatures, while Ma et al. (2025) discussed broader challenges in pragmatic evaluation. Their work supports the value of testing interpretation systematically rather than relying on a few examples. In the present design, the target sentence remains unchanged while the situation varies. An appropriate model response should adapt when the new circumstances support a different inference, while avoiding a definite claim when the context remains unclear.

Social Context

Social context includes speaker relationships, roles, familiarity, and expectations about politeness. Consider, “I will think about it.” Between friends, it may communicate genuine uncertainty. In a student–teacher exchange, it may indicate that a decision is being postponed, that the request will be considered, or that the speaker wishes to avoid an immediate commitment. It does not necessarily mean refusal.

Wagner et al. (2025) identified limitations in ChatGPT’s handling of contextually and culturally situated conversation. The present study can build on this concern by changing the social relationship while keeping the target utterance constant. It should assess whether the model acknowledges uncertainty when information is incomplete and avoids treating a social convention as a rule that applies to every person or situation.

Accuracy and Consistency

Accuracy concerns whether an interpretation is supported by the words and context. Consistency concerns whether the model gives similar interpretations when relevant information is the same and changes its interpretation when meaningful context changes. These ideas are related but not identical. If a model changes its interpretation because the situation changes, the change may show appropriate contextual sensitivity rather than inconsistency.

Four coding categories are used: accurate literal interpretation, contextually supported implicature, ambiguous interpretation, and unsupported inference. An accurate literal interpretation correctly explains what the words state. A contextually supported implicature identifies a plausible implied meaning that fits the situation. An ambiguous interpretation leaves uncertainty unresolved. An unsupported inference assumes an intention that the available evidence does not justify. These categories provide a transparent way to compare responses.

The categories should be applied using clear criteria. A response should not be labelled wrong simply because it differs from the researcher's preferred reading, particularly when several interpretations are plausible. Where feasible, independent human judgments and a second coder can help identify disagreements. The final analysis should report the actual patterns observed in the dataset and should not claim accuracy rates or consistent performance unless those claims are supported by coded responses.

Outcomes

The following are expected analytical points, not verified findings. They should be confirmed, revised, or removed after the responses have been collected and coded.

1. Literal meaning and implied meaning are related, but they are not always the same.
2. Context can affect whether an utterance functions as a request, criticism, refusal, or simple statement.
3. Surrounding discourse may clarify an implication that is unclear in an isolated sentence.
4. Situational information may support an indirect request, but the inference still requires evidence.
5. Social relationships may affect how politeness and non-committal responses are understood.
6. A plausible implication should not be treated as certain when other readings remain possible.
7. Accuracy depends on whether an interpretation is supported by the utterance and its context.
8. Consistency means maintaining interpretations under equivalent conditions and changing them when relevant context changes.
9. The coding framework separates supported interpretations, ambiguity, and unsupported inference.

10. The study's contribution is a systematic comparison of literal meaning, implied meaning, and contextual variation.

References

1. Andersson, M., & McIntyre, D. (2025). Can ChatGPT recognize impoliteness? An exploratory study of the pragmatic awareness of a large language model. *Journal of Pragmatics*, 239, 16–36. <https://doi.org/10.1016/j.pragma.2025.02.001>
2. Cho, Y.-E., & Kim, S. (2024). Pragmatic inference of scalar implicature by LLMs. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)* (pp. 10–20). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.acl-srw.2>
3. Cong, Y. (2024). Manner implicatures in large language models. *Scientific Reports*, 14, Article 29113. <https://doi.org/10.1038/s41598-024-80571-3>
4. Creswell, J. W., & Poth, C. N. (2018). *Qualitative inquiry and research design: Choosing among five approaches* (4th ed.). Thousand Oaks, CA: SAGE Publications.
5. Ghufuran, F., Ali, A., Khan, M. A., & Irshad, P. (2026). Evaluating the effect of task-based language teaching (TBLT) on speaking fluency among ESL learners: A quasi-experimental study. *The Critical Review of Social Sciences Studies*, 4(1), 41–53. <https://doi.org/10.59075/zdpypt08>
6. Ghufuran, F., Durrani, N., & Khan, J. (2025). Corpus-based comparative study of hedging in academic writing by native and non-native English writers. *Journal of Asian Development Studies*, 14(4), 350–364. <https://doi.org/10.62345/jads.2025.14.4.26>
7. Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: Vol. 3. Speech acts* (pp. 41–58). New York, NY: Academic Press. https://doi.org/10.1163/9789004368811_003
8. Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. Beverly Hills, CA: SAGE Publications.
9. Ma, B., Li, Y., Zhou, W., Gong, Z., Liu, Y. J., Jasinskaja, K., Friedrich, A., Hirschberg, J., Kreuter, F., & Plank, B. (2025). Pragmatics in the era of large language models: A survey on datasets, evaluation, opportunities and challenges. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 8679–8696). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.acl-long.425>
10. Wagner, W., Gaskell, G., Paraschou, E., Lyu, S., Michali, M., & Vakali, A. (2025). Limits of ChatGPT's conversational pragmatics in a Turing test on ethics, commonsense, and cultural sensitivity. *Computers in Human Behavior: Artificial Humans*, 5, Article 100191. <https://doi.org/10.1016/j.chbah.2025.100191>
11. Walsh, S. (2011). *Exploring classroom discourse: Language in action*. London, England: Routledge.