Rhetorical Effects of AI Paraphrasing on Argumentation and Authorial Stance in Arab EFL Literature Reviews

Dina Talal Alhajailan[1]

[1] Assistant Professor of Applied Linguistics at Princess Nourah bint Abdulrahman University, Saudi Arabia

Correspondence: Dina Talal Alhajailan, Assistant Professor of Applied Linguistics at Princess Nourah bint Abdulrahman University, Saudi Arabia. E-mail: dtalhajailan@pnu.edu.sa

Received: April 15, 2026 Accepted: July 11, 2026 Online Published: November 3, 2026

doi:10.5430/wjel.v17n1p490 URL: https://doi.org/10.5430/wjel.v17n1p490

Abstract

Recent research has identified systematic differences between AI-generated and human-authored academic writing. However, most studies focus on argumentative essays, leaving the rhetorical effects of AI use on literature reviews underexplored. This paper addresses this gap through a controlled analysis of how AI paraphrasing reshapes argumentation and authorial stance in Arab EFL postgraduate literature reviews under content-preserving conditions.

Using a paired corpus design, this study compares ten human-authored literature reviews produced in 2020, prior to the widespread use of large language models, with ChatGPT-paraphrased versions produced using a controlled prompt. This design isolates core AI transformation patterns. Analysis draws on Toulmin's model of argumentation and Hyland's stance framework, supported by corpus-based validation.

Results indicated systematic rhetorical transformation. While ChatGPT preserved structural coherence and improved surface fluency, the AI-generated texts showed reduced rhetorical complexity and weakened authorial stance. Self-mention and reader inclusion were sharply reduced, hedging became more formulaic, warrants shifted from situated reasoning to generic principles, and contrastive and causal connectives occurred less frequently. These findings clarify which rhetorical features are most vulnerable under controlled AI paraphrasing and position AI paraphrasing as a concern for EFL academic writing pedagogy.

Keywords: AI paraphrasing, authorial voice, argumentation, EFL academic writing, literature reviews

1. Introduction

Generative artificial intelligence (AI) has reshaped academic writing practices, raising new questions about authorship, cognition, and rhetorical representation in scholarly discourse. This is no longer a marginal development: Kobak et al. (2025) found that at least 13.5% of published research abstracts they examined from 2024 showed signs of LLM processing. At the same time, AI is now deeply embedded in student writing practices, with 86% of students around the world reporting regular AI use in their studies (Digital Education Council, 2024), including the use of models to paraphrase, polish, and improve academic writing. Large language models can produce fluent, coherent academic prose, but emerging research suggests that such texts may lack the interpretive reasoning, contextual awareness, and epistemic stance associated with human writing (Lund et al., 2023; Zaitsu & Jin, 2023). Although growing research has compared AI-generated and human-authored texts (Jiang & Hyland, 2024, 2025a, 2025b, 2025c), we still know little about how AI transforms existing human writing, particularly in demanding genres such as literature reviews.

Human academic writing is anchored in critical judgement, disciplinary awareness, and a sense of audience. Writers do not simply present information; they interpret evidence, negotiate claims, and signal evaluation in ways that shape both knowledge construction and authorial identity. By contrast, AI-generated writing can reproduce many of the visible features of scholarly discourse while failing to preserve the same depth of contextual understanding, interpretive reasoning, or rhetorical positioning (Tarchi et al., 2024; Doru et al., 2025).

Although AI can assemble coherent literature reviews, concerns remain about its ability to sustain argumentation, stance, and voice, as well as its tendency toward generic evaluation, redundant phrasing, and unreliable citation (Mugaanyi et al., 2024; Mohamed et al., 2024). Despite growing research on AI-assisted academic writing, few studies have examined how AI transforms human writing through paraphrasing in a controlled, genre-sensitive context. This gap matters especially in Arab EFL postgraduate writing, where literature reviews are a key site for developing disciplinary reasoning and authorial stance. Most existing work has focused on short essays, abstracts, or detection-based outputs, often without attending to the genre conventions and rhetorical functions that shape extended academic texts. Literature reviews, in particular, remain underexplored, even though they demand synthesis, critique, argumentation, and authorial positioning. This gap limits our understanding of how AI paraphrasing affects reasoning and stance in one of the more demanding forms of academic discourse. Because AI paraphrasing is increasingly part of digitally mediated writing support, this study approaches it not only as a question of text transformation, but also as a concern for EFL academic writing development.

The study addresses three research questions.

1.1 Research Questions

i. How does ChatGPT paraphrasing transform reasoning structures in literature reviews when analysed through Toulmin's (2003) model of argumentation (claim, data, warrant, backing, qualifier, and rebuttal)?

ii. How does ChatGPT paraphrasing transform authorial stance in literature reviews according to Hyland's (2005) stance model (hedges, boosters, evaluative language, self-mention, and reader inclusion)?

iii. To what extent does AI paraphrasing preserve or erode the interpretive reasoning and authorial stance in human-authored literature reviews?

2. Literature Review

2.1 AI, Academic Writing, and the Question of Rhetorical Depth

The growing presence of generative AI in academic writing has reopened questions of authorship, accountability, and the nature of scholarly production itself. As automated tools take on a more visible role in knowledge creation, it becomes harder to treat questions of credit, responsibility, and ethical accountability as settled (Hajibayova & Lee, 2025). The issue is not only procedural. It also affects how authenticity, expertise, and credibility are perceived when academic discourse is mediated by AI (Lamm et al., 2025; Gilardi et al., 2024).

The discussion, however, does not stop at ethics. AI-generated writing has sharpened an older distinction between surface fluency and meaningful academic thinking. Lund et al. (2023) argue that tools such as ChatGPT do not think in any genuinely cognitive sense; rather, they reproduce learned linguistic patterns with increasing sophistication. A growing body of comparative work suggests that this fluency can conceal important rhetorical and conceptual limits. Human-authored texts tend to show more lexical and syntactic variety, stronger conceptual layering, and greater contextual sensitivity. AI-generated texts more often rely on repetition, formulaic phrasing, and a kind of surface coherence that can look persuasive without carrying the same depth of thought (Durak, Eğin, & Onan, 2025; Agrahari et al., 2024). Even where AI-generated writing appears polished, that polish does not necessarily bring with it interpretive nuance, critical depth, or contextual grounding (Nathani et al., 2025; Danylchenko-Cherniak, 2024; Rad et al., 2024). The issue, then, is not simply whether AI can produce academic-looking prose, but whether it can preserve the interpretive and rhetorical work that gives academic writing much of its value.

2.2 Literature Reviews, Citation, and Epistemic Engagement

These concerns become especially important in literature reviews. A literature review is not a neutral summary of prior studies. It asks the writer to evaluate sources, integrate perspectives, and position themselves within an ongoing scholarly conversation. Citation, in this context, is not merely a technical convention. It is part of how writers build authority, organise knowledge, and take a deliberate authorial position (Badenhorst, 2019; Jaidka et al., 2013).

Research suggests that AI has difficulty sustaining this kind of work. Zaitsu and Jin (2023) identify clear differences in stylistic and rhythmic patterning between AI-generated and human-authored writing, while Tarchi et al. (2024) and Doru et al. (2025) show that AI-generated essays tend toward redundant phrasing and mechanically coherent sentences. More broadly, work on AI-assisted academic writing has raised concerns about weaker source grounding and thinner contextual relevance in genres that depend heavily on intertextual judgement (Mugaanyi et al., 2024; Kim et al., 2024). This matters because literature reviews depend not only on coherence, but on informed selection, positioning, and relation-building across texts. AI may reproduce some of the visible architecture of academic discourse while still remaining weaker at source-sensitive reasoning and intertextual judgement.

This becomes particularly consequential in literature review writing. Melliti (2024) argues that human-written literature reviews retain interpretive value and contextual knowledge that AI-generated reviews often lack, even when those reviews appear logically organised and clear. Similarly, Amirjalili, Neysani, and Nikbakht (2024) found that although ChatGPT can imitate the structural features of literature reviews, it struggles to create authentic authorial tone, evaluative subtlety, and intertextual interaction. Similar tendencies have been observed in adjacent academic genres, where human-written abstracts show greater rhetorical flexibility while AI-generated abstracts tend to settle into more formulaic organisational patterns (Kong & Liu, 2024). The crucial distinction, then, is not simply between human and machine text, but between structural imitation and epistemic engagement.

2.3 Student/EFL Writing and the Present Study

The contrast between human and AI-generated writing is also visible in student writing. In a series of comparative studies, Jiang and Hyland (2024, 2025a, 2025b, 2025c) show that student-authored arguments and essays tend to display stronger authorial presence, richer epistemic stance, and greater reader engagement than AI-generated counterparts, even when the AI texts appear more polished on the surface. For second-language academic writers, however, the attraction of AI assistance may be especially strong. AI tools can improve clarity, fluency, and conformity to academic conventions, which makes them practically appealing in ESL/EFL contexts. This broader pedagogical relevance is also reflected in Aljuaid's (2024) systematic review, which examines the impact of AI tools on academic writing instruction in higher education and highlights both their instructional potential and their challenges. Herbold et al. (2023), for example, found that ChatGPT-generated essays were rated more highly than non-native student essays by educators. Yet such advantages may come with rhetorical costs. Al Hosni's (2025) review similarly raises concerns that AI-assisted writing may erode the linguistic fingerprint through which student writers establish authorial voice and identity, a concern that is especially relevant in university ESL writing contexts.

For EFL postgraduate writers, literature review writing is a particularly important site of development because linguistic proficiency, disciplinary reasoning, and authorial positioning all intersect there. Despite growing research on AI-assisted writing, much of the existing literature has focused either on independently generated texts or on genres other than literature reviews. One important exception is Mostafapour et al. (2024), who compared human and GPT-4-generated literature reviews on the same topic and found that while AI demonstrated speed and breadth, human reviewers showed greater accuracy, depth, and contextual understanding. However, that study compared independently produced reviews rather than examining how AI transforms an existing human text through paraphrasing.

The present study addresses that gap. Rather than asking how AI generates new academic text, it asks what happens to the rhetorical features of academic writing when AI paraphrases human-authored literature reviews. That distinction matters because paraphrasing lets us observe AI's effect on existing human reasoning and stance, not just on newly generated output. To examine this, the study draws on Toulmin's model of argumentation and Hyland's stance framework in order to trace changes in reasoning structure, authorial voice, and writer-reader positioning. It also adopts a deliberately constrained design. This is not because naturalistic student-AI interaction is unimportant, but because a baseline understanding of AI-driven textual transformation is needed before more complex forms of AI-assisted revision can be interpreted with confidence.

3. Method

3.1 Research Design

This study employed a mixed-method, comparative design that combined corpus-based analysis with qualitative discourse analysis. Its central aim was to examine how AI paraphrasing reshapes the rhetorical features of student academic literature reviews under tightly controlled conditions. The design was deliberately constrained in order to establish a baseline account of AI-driven textual transformation. More naturalistic uses of AI, such as "improve this text" or "make this more critical," are also important and deserve study. However, before those more open forms can be interpreted clearly, it is important to identify what the model changes when it is restricted to paraphrasing. The findings should therefore be interpreted as evidence of rhetorical transformation under constrained paraphrasing conditions, not as a direct account of the full range of AI-assisted academic writing practices.

3.2 Corpus

The study was based on a parallel corpus of twenty literature reviews. Ten of these were human-authored texts written in 2020 by Arab postgraduate EFL students, prior to the widespread availability of generative AI tools. The remaining ten were AI-paraphrased versions of the same texts produced through a controlled ChatGPT prompt. All human-written literature reviews were original texts in applied linguistics.

Text length was used as the inclusion criterion. Literature reviews were included if they fell between 500 and 2,200 words, a range that reflects a typical span for postgraduate literature review writing. The human corpus contained 14,490 words across ten texts, with a mean length of 1,449 words. The AI-paraphrased texts were shorter than their human-authored counterparts, totalling 9,599 words (mean = 960).

3.3 Instruments

The analysis drew on two established frameworks: Toulmin's model of argumentation (2003) and Hyland's model of stance (2005). These frameworks were selected because together they allow the study to examine both how writers construct reasoning and how they position themselves rhetorically in relation to their claims and readers.

From Toulmin's model, the analysis examined claim, data, warrant, backing, qualifier, and rebuttal as components of reasoning structure. From Hyland's model, it examined hedges, boosters, evaluative language, self-mention, and reader inclusion as markers of stance and authorial presence.

An operational coding manual was developed from these two frameworks and used throughout the analysis. In addition to the qualitative coding scheme, two digital tools were used for quantitative validation and triangulation. Voyant Tools was used for lexical and structural profiling, including total word counts, type-token distributions, lexical diversity, sentence length, and readability. AntConc was used to examine transitional marker frequencies across five categories: additive, contrastive, causal, temporal, and exemplification.

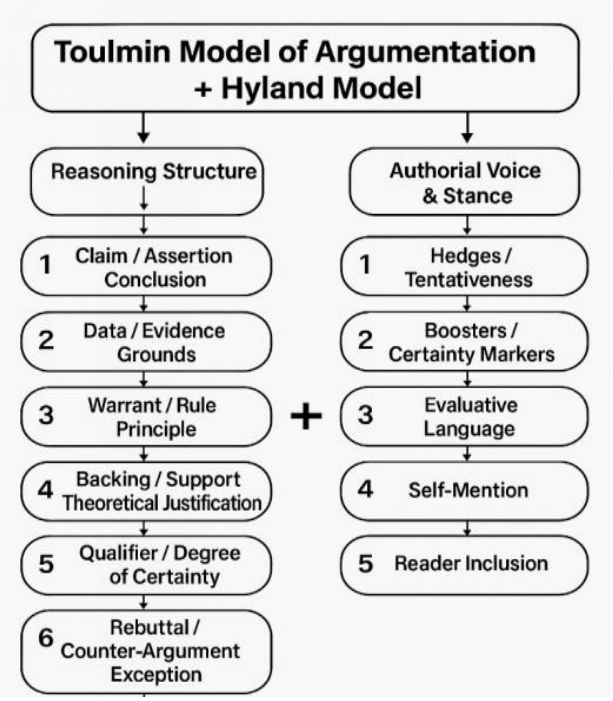


Figure 1. Hybrid Analytical Framework Integrating Toulmin's Model of Argumentation and Hyland's Model of Authorial Stance

3.4 Procedures

To isolate AI's rhetorical transformation effects, a content-preserving prompt was used to constrain the model while maintaining the meaning and citation order of the original human texts. All AI texts were generated in one controlled session using GPT-5.2 (OpenAI), accessed through the ChatGPT interface in January 2026, to minimise variation in model behaviour.

The prompt instructed the model to preserve the content, structure, citation placeholders, paragraphing, and approximate length of the source texts while rephrasing them in its own words. Each review was processed separately, with one human text entered and completed before the next was introduced. After generation, no manual editing or rephrasing was performed other than basic formatting. This helped ensure that the AI corpus remained free from human rewriting and that differences between the human and AI texts reflected the generative process itself rather than prompt variation or post-editing.

The prompt used was as follows:

> You are a writing assistant in academics. Rebuild the given literature review in your own words, ensuring that the meaning, structure, and placeholders of citations remain the same. No information to be added or removed, no reference numbers or years to be changed and the number of paragraphs and approximate length of words (within 5 percent) should remain the same. Write naturally and coherently, but do not strive to imitate human tone of rhetoric and his stylistic peculiarities. Restate the literature review below using your own words without altering the same concepts and flow. Keep the same focus on the topic.

For the qualitative analysis, clauses were used as the primary units of analysis because they represent meaningful segments in which claims, evidence, and stance can be realised. Since a clause may perform more than one rhetorical function, the operational manual based on the two analytical frameworks specified how such clauses were coded.

Two trained coders participated in the analysis: the author and a non-author collaborator, an assistant professor of applied linguistics specialising in corpus linguistics and discourse analysis, who has asked to remain anonymous. Both corpora were piloted in advance using the coding manual, and category definitions were discussed before full coding began. The coders then independently coded the full dataset, after which disagreements were reviewed and resolved through discussion. Inter-rater reliability was calculated using Cohen's kappa on the independent codes before disagreements were resolved, with a resulting score of $\kappa = 0.91$, indicating a high level of agreement. Word-frequency counts were normalised per 1,000 words to account for differences in text length (Tables 1 and 4 report raw and normalised frequencies), and descriptive statistics were calculated using means and percentages for the human and AI corpora.

Because rhetorical analysis necessarily involves interpretive judgement, the study also incorporated quantitative procedures to support the qualitative findings. Lexical and structural profiling was conducted in Voyant Tools to compare total word counts, type-token distributions, lexical diversity, sentence length, and readability across the human and AI corpora. In addition, AntConc was used to examine frequencies of transitional markers across five categories—additive, contrastive, causal, temporal, and exemplification—in order to validate patterns of cohesion and rhetorical organisation. These measures were used as triangulation tools to support the feature-based qualitative analysis and to reduce reliance on coder interpretation alone.

3.5 Ethical Considerations

No confidential data were used in this study. All human-authored literature reviews were written in 2020, before the widespread use of generative AI tools, and were drawn from a publicly accessible institutional repository at a public university. No human participants were involved, and no permissions were required. The AI corpus was generated solely for the purposes of this research, and ChatGPT was used only to produce controlled comparative versions of existing human texts. The corpus analysed in this study is available from the author upon request.

3.6 Limitations and Delimitations

This study examined AI transformation of human writing under constrained, content-preserving conditions. Although that design does not capture the full range of student or academic uses of AI, it does make it possible to identify systematic model tendencies under controlled paraphrasing conditions. All human texts were written in English by second-language users. While this limits generalisability, it also offers insight into how argumentation and stance are realised in EFL academic writing.

A further strength of the dataset is that the human corpus predates generative AI, making it possible to analyse texts that reflect unassisted human reasoning patterns. The analysis was restricted to selected rhetorical and linguistic features and did not examine content accuracy or the broader ethical implications of AI use. It was also limited to literature reviews and to the analytical categories drawn from Toulmin and Hyland. Importantly, the corpus comprised only ten pairs of literature reviews. While this small, controlled sample made close rhetorical analysis of each pair possible, the patterns identified here should be read as indicative rather than conclusive. These boundaries strengthen analytical precision, though they reduce the generalisability of the findings. Future research could address this limitation by drawing on larger and more diverse datasets—across proficiency levels, first-language backgrounds, and institutional contexts—as well as by comparing multiple AI models and examining additional academic genres such as abstracts or discussion sections.

4. Results

The analysis revealed consistent patterns of rhetorical transformation in AI-paraphrased literature reviews. Most notably, AI paraphrasing altered authorial presence, epistemic stance, and the balance between information presentation and argument development.

4.1 Comparative Frequency Summary

Table 1 summarises the frequencies of argumentation and stance features across the human and AI corpora.

Table 1. Comparative Summary of Human and AI Argumentation and Voice Features

Category	Raw Frequency				Per 1,000 Words		
	Human	AI	Difference	% Change	Human	AI	% Change
Claim / Assertion / Conclusion	92	141	+49	+53%	6.35	14.69	+131%
Data / Evidence / Grounds	163	297	+134	+82%	11.25	30.94	+175%
Warrant / Rule / Principle	64	77	+13	+20%	4.42	8.02	+82%
Backing / Support / Theoretical Justification	85	91	+6	+7%	5.87	9.48	+62%
Qualifier / Degree of Certainty	47	43	−4	−9%	3.24	4.48	+38%
Rebuttal / Counterargument / Exception	30	30	0	0%	2.07	3.13	+51%
Hedges / Tentativeness	52	32	−20	−38%	3.59	3.33	−7%
Boosters / Certainty Markers	54	56	+2	+4%	3.73	5.83	+57%
Evaluative Language	36	35	−1	−3%	2.48	3.65	+47%
Self-Mention	25	5	−20	−80%	1.73	0.52	−70%
Reader Inclusion	16	1	−15	−94%	1.10	0.10	−91%

Note. Normalised rates were calculated as raw frequency ÷ corpus word count × 1,000, using the corpus totals in Table 3 (human = 14,490 words; AI = 9,599 words). Normalised percentage changes were calculated from unrounded rates.

The frequency summary reveals two contrasting rhetorical profiles. Human-authored reviews show greater variation across categories, more flexible argumentation, and greater sensitivity to academic voice. By contrast, the AI-paraphrased reviews show a more regularised and impersonal rhetorical profile, with increased claims and data, more formulaic hedging, and reduced self-reference and reader engagement.

To make these shifts more concrete, Table 2 presents representative human-AI example pairs drawn from the corpus across the three major

transformation patterns.

Table 2. Representative Human-AI Transformations Across Key Rhetorical Features

Pattern	Feature	Human Example	AI Example	Main Shift
Informational densification	Claim / Assertion / Conclusion	"These results attract more attention to the importance of early language development in children with autism."	"These findings highlight the significance of both child-based and parental language factors."	More generalised, declarative claim; less context-grounded framing
	Data / Evidence / Grounds	"They employed 32 preschoolers with ASD and 35 matched peers aged 2 to 5 years."	"Fusaroli et al. (2019) analysed 32 preschoolers with ASD and 35 matched peers aged two to five years."	Evidence becomes citation-centred and compressed into summary form
Epistemic flattening	Warrant / Rule / Principle	"…it has been suggested that effective government communication of COVID-19 information with the public is critical…"	"…highlighting the necessity for effective government communication with the public…"	Qualified framing shifts toward a more categorical statement of necessity
	Qualifier / Degree of Certainty	"We can say that…"	"These findings suggest…"	Human qualification is dialogic; AI qualification is more formulaic
Depersonalisation	Self-Mention	"We can say that autism affects children's language abilities."	"The present research aims to explore…"	First-person presence replaced by impersonal academic framing
	Reader Inclusion	"To do so, linguists in Arab countries should search more…"	"This study is relevant as it…"	Direct reader engagement replaced by detached informational framing

Note. Citations appearing within the example excerpts (e.g., Fusaroli et al., 2019) belong to the source corpus texts and are not references of the present study.

These examples help illustrate the movement from more situated, dialogic, and writer-visible discourse toward a more compressed, impersonal, and pattern-driven rhetorical style. Appendix A provides two complete paired passages that illustrate these transformations at greater length.

4.2 Pattern 1: Depersonalisation

The clearest transformation involved the reduction of authorial presence. Self-mention fell from 25 to 5 instances, a 70% decrease per 1,000 words, while reader inclusion dropped from 16 to 1, a 91% reduction. Texts that were originally more dialogic in orientation became markedly more monologic in form. In concrete terms, language that had made the writer visible and brought the reader into the reasoning process was often replaced by more detached and impersonal framing.

4.3 Pattern 2: Epistemic Flattening

A second pattern appears in the treatment of epistemic stance. AI paraphrasing reduced hedges from 52 to 32 instances, a decline of only 7% per 1,000 words, while increasing boosters from 54 to 56, a 57% rise per 1,000 words. Qualifiers fell from 47 to 43 but rose 38% when normalised. Taken together, these shifts point toward a flatter and more uniformly assertive rhetorical stance. Even where caution remained, it was more often expressed through a narrower and more formulaic set of phrases, giving the texts a more even but less nuanced stance profile.

4.4 Pattern 3: Informational Densification

Claims increased by 131% and data by 175% per 1,000 words, while warrants and backing rose more modestly. On the surface this might suggest stronger argumentation, but the qualitative analysis indicates a different shift: the AI-paraphrased texts became denser in information while thinner in interpretation. Evidence was more often compressed into citation-focused summary, while the links between evidence and interpretation became less visible and less context-sensitive.

4.5 Lexical Validation and Triangulation Through Voyant Tools

To support the qualitative findings, lexical and structural profiling was conducted using Voyant Tools. Table 3 summarises the aggregate differences between the human and AI corpora.

Table 3. Summary of Aggregate Lexical and Structural Averages for Human and AI Corpora

Group	Total Words	Total Types	Total Ratio (%)	Total Words/Sentence	Readability (Coleman-Liau Index)
10 Human	14,490	2,711	19	27.1	14.950
10 AI	9,599	2,203	23	25.0	18.279

The Voyant results show that the human corpus contained more words and word types than the AI-paraphrased corpus. The AI corpus showed a higher type-token ratio and shorter sentences on average.

4.6 Transitional Marker Analysis Through AntConc

A further form of validation was carried out through AntConc-based analysis of transitional markers. Table 4 summarises the frequencies and proportions of five categories of transitions across the two corpora.

Table 4. Summary of Aggregate Transitional Marker Frequencies, Normalised Rates, and Percentages for Human and AI Corpora

Group	Category	Additive	Contrastive	Causal	Temporal	Exemplification	Total
Human	Frequency	573	122	95	66	70	926
Human	Per 1,000 Words	39.54	8.42	6.56	4.55	4.83	63.91
Human%	% of Total	61.90%	13.20%	10.30%	7.10%	7.60%	
AI	Frequency	305	57	33	29	43	467
AI	Per 1,000 Words	31.77	5.94	3.44	3.02	4.48	48.65
AI%	% of Total	65.30%	12.20%	7.10%	6.20%	9.20%	

Note. Per 1,000 words = raw frequency ÷ corpus word count (Table 3) × 1,000. Percentages of total are within-corpus proportions.

The human-authored texts used a substantially higher overall number of transitional markers than the AI-paraphrased texts, and still contained 31% more per 1,000 words once text length was taken into account (Table 4). Additive markers were the most frequent category in both corpora, which reflects shared academic conventions of cohesion. More revealing, however, is the relative reduction in contrastive and causal markers in the AI corpus. Human-authored reviews relied more heavily on these transitions to develop interpretive contrast and causal reasoning, both of which are central to argumentative literature review writing.

5. Discussion

5.1 Depersonalisation

The shift to depersonalisation matters because self-mention and reader inclusion do more than personalise the prose. They locate the writer within the discourse, signal responsibility for interpretation, and bring the reader into the reasoning process. In the human-authored corpus, first-person references and inclusive forms contributed to a sense of epistemic ownership and disciplinary participation. In the AI corpus, these features were largely removed or replaced with impersonal framing. The effect is not simply stylistic neutrality. It is a movement toward depersonalised academic discourse.

Taken together, the patterns identified in this study suggest that AI paraphrasing reduces the visible presence of the writer even when it preserves topic and content. This tendency aligns with Moussa et al. (2026), whose analysis of EFL undergraduate writing found that AI-assisted texts often appeared more polished and impersonal while reducing visible markers of writer presence. These findings extend that observation to a controlled paraphrasing context, suggesting that depersonalisation may be not only an artefact of text generation, but also a recurring feature of AI-mediated rewording.

5.2 Epistemic Flattening

In the human-authored reviews, hedging was used more flexibly to signal caution, openness to interpretation, and the limits of generalisation. In the AI corpus, hedging was not absent, but it became narrower and more formulaic, often relying on a restricted set of expressions such as *suggest* and *indicate*. This gives the appearance of caution without preserving the same degree of context-sensitive rhetorical judgement.

At the same time, the increase in boosters matters because it occurs alongside the modest reduction in hedging. This does not make the AI texts aggressively assertive, but it does make them more even in tone and less responsive to the subtle negotiations of certainty that characterise stronger academic writing. Rebuttal structures, like qualifiers and evaluative language, rose once normalised for length, but they too became more linear and less dialogic in expression. Overall, the result is not the disappearance of stance, but its flattening: a movement away from situated rhetorical judgement and toward a more standardised, pattern-based presentation of claims.

5.3 Informational Densification

In the human-authored reviews, claims were more often embedded in narrative or context-sensitive reasoning, and data were used not only to report prior studies but to support a developing line of argument. In the AI corpus, by contrast, evidence was more often compressed into citation-heavy summaries. This makes the prose appear efficient and well-supported, but it also tends to decontextualise the relationship between evidence and interpretation.

A similar change appears in warrants. Although warrants increased from 64 to 77, they were more likely in AI-paraphrased texts to take the form of generalised principles rather than situated reasoning tied to a specific pedagogical, psychological, or disciplinary context. Backing also shifted in function. Human-authored texts more often used support in ways grounded in empirical or programmatic context, whereas AI texts leaned more heavily toward canonical attribution and abstract scholarly authority.

For that reason, the increase in argumentative elements should not be read simply as argumentative improvement. The AI texts become richer in propositional content, but less anchored in the kinds of contextual and interpersonal reasoning through which academic arguments

are usually made persuasive. Informational density increases, but rhetorical depth does not increase with it.

These three patterns—depersonalisation, epistemic flattening, and informational densification—were observed consistently across all ten text pairs. The regularity of these changes suggests systematic algorithmic tendencies rather than incidental variation between individual texts.

In one sense, the shorter sentences in the AI texts support the common perception that AI improves fluency. At the same time, the qualitative analysis suggests that such simplification can come at the cost of rhetorical sophistication. The AI texts read more smoothly, but are less layered in structure.

A similar pattern emerged in the use of discourse markers: the AI texts relied proportionally more on additive and exemplification markers. This produces a discourse that remains coherent, but is more sequential and descriptive than analytically developed. The reduction in causal markers is especially important, because causal reasoning is one of the ways literature reviews move beyond summary and toward explanation. These quantitative findings therefore reinforce the qualitative analysis: AI paraphrasing tends to shift discourse away from interpretive argumentation and toward information presentation.

6. Conclusion

This study examined how AI paraphrasing, specifically through ChatGPT, reshapes the rhetorical features of student academic literature reviews. Even under content-preserving constraints, three interrelated transformation patterns emerged consistently across the dataset: depersonalisation, epistemic flattening, and informational densification.

Depersonalisation was reflected in the substantial reduction of self-mention (−70%) and the near elimination of reader inclusion (−91%), two features through which student writers signal authorial presence and engage readers in disciplinary dialogue. Epistemic flattening appeared in more formulaic use of hedging and in a shift from more contextually situated warrants toward more generalised principles. Informational densification was evident in the increased frequency of claims and data references, alongside a reduction in contrastive and causal connectives, producing texts that were richer in information but less interpretively connected. Taken together, these findings suggest that AI paraphrasing does not simply restate academic content. It alters the rhetorical character of the text.

These patterns matter for academic writing pedagogy, particularly for student writers who are still developing an authorial voice. Almousa and AbuSa'aleek (2025) report generally positive postgraduate attitudes toward using ChatGPT in thesis writing, which makes it especially important for pedagogy to address not only what AI may improve at the surface level, but also what it may reduce in terms of stance, voice, and rhetorical presence. Features such as self-mention, hedging, and reader engagement are not ornamental. They support emerging confidence, epistemic positioning, and participation in scholarly communities. When AI paraphrasing systematically reduces these features, it risks narrowing the rhetorical resources through which students learn to position themselves as academic writers. This concern may be particularly relevant in EFL contexts, where students may turn to AI tools to address surface-level linguistic difficulties without being fully aware of the rhetorical trade-offs involved. From an educational perspective, the issue is not only that individual students may lose aspects of voice and stance, but that widespread reliance on such transformations could contribute to a collective narrowing of rhetorical range and academic thought.

In classroom terms, several strategies follow from these findings. Instructors can build paired-text comparison tasks into writing courses, asking students to place their own drafts alongside AI-paraphrased versions and to trace what has happened to self-mention, hedging, and reader engagement; the paired excerpts in Appendix A illustrate the kind of material such tasks might draw on. Stance-focused revision checklists can serve a similar purpose, prompting students to check, after any AI-assisted revision, whether their claims are still hedged where caution is warranted and whether the writer is still visible in the text. Alongside these tasks, instruction can make explicit why the features at stake are worth protecting. Self-mention and reader engagement are the means by which a text enacts a relationship between writer and reader, and the situated reasoning observed in the human corpus reflects an individual writer thinking through a problem rather than a pattern being reproduced. When students understand these features as the human dimension of academic writing rather than as stylistic preferences, they have a principled reason to preserve them. Students can also be taught voice-preserving prompting, in which the tool is explicitly instructed to retain first-person positioning, hedging, and evaluative commentary rather than being left to its defaults. Finally, where AI paraphrasing is used, it is best treated as a draft to be revised rather than as a final version. The present findings indicate that paraphrasing even a complete, well-developed text can alter its rhetorical character, so students should be encouraged to return to the AI output and restore hedges, self-mention, and reader engagement where the tool has removed them. None of these strategies requires abandoning AI tools. They require treating AI output as an object of rhetorical analysis rather than as a finished product.

Given the speed with which generative AI has been adopted across educational and academic settings (Bick et al., 2026), these pedagogical issues are unlikely to remain peripheral. In practical terms, this suggests that EFL writing instruction should not treat AI paraphrasing simply as a fluency aid. Students may need explicit support to notice what AI adds in fluency and what it may reduce in self-mention, hedging, reader engagement, and context-sensitive reasoning. One implication for EFL writing pedagogy is that AI output should be discussed comparatively, so that students can evaluate not only correctness and clarity, but also rhetorical presence and argumentative control. Within the methodological limits of this study, the findings provide a baseline for understanding the rhetorical effects of AI paraphrasing on human academic writing. Future research could examine how these patterns vary across disciplines, writer populations, and more naturalistic forms of AI-assisted revision.

Acknowledgments

The author thanks an anonymous trained coder for contributing to independent coding and the discussion of coding disagreements.

Author contributions

The author conceived and designed the study, collected and analysed the data, interpreted the findings, and wrote and approved the final manuscript.

Funding

This research received no external funding.

Competing interests

The author declares no conflict of interest.

Informed consent

Not applicable.

Ethics approval

Not applicable.

Provenance and peer review

Not commissioned; externally double-blind peer reviewed.

Data availability statement

The human-authored texts were obtained from a publicly accessible institutional repository. The corpus analysed in this study is available from the author upon request.

Data sharing statement

No additional data are available.

Open access

This is an open-access article distributed under the terms and conditions of the Creative Commons Attribution license (http://creativecommons.org/licenses/by/4.0/).

Copyrights

Copyright for this article is retained by the author(s), with first publication rights granted to the journal.

AI Disclosure Statement

The author used generative AI tools during the drafting and language-editing process to improve readability and expression. ChatGPT was also used as the object of analysis in the study's controlled paraphrasing procedure. The author takes responsibility for the analysis, interpretation, and final content of the manuscript.

References

Agrahari, S., Bisht, S., & Sanasam, R. (2024). Text authorship attribution: Stylometric insights into human and LLM-generated text. In *Proceedings of the 8th International Conference on Data Science and Management of Data (CODS-COMAD '24)* (pp. 344-346). ACM. https://doi.org/10.1145/3703323.3703712

Al Hosni, J. (2025). Preserving authorial voice in academic texts in the age of generative AI: A thematic literature review. *Arab World English Journal, 16*(3), 244-258. https://doi.org/10.24093/awej/vol16no3.14

Aljuaid, H. (2024). The impact of artificial intelligence tools on academic writing instruction in higher education: A systematic review. *Arab World English Journal (AWEJ)*, Special Issue on ChatGPT, 26-55. https://doi.org/10.24093/awej/ChatGPT.2

Almousa, M. M., & AbuSa'aleek, A. O. (2025). Postgraduate students' attitude toward using ChatGPT in enhancing their master's thesis: A mixed method study. *Arab World English Journal*, (Special Issue on Artificial Intelligence), 157-175. https://doi.org/10.24093/awej/AI.9

Amirjalili, F., Neysani, M., & Nikbakht, A. (2024). Exploring the boundaries of authorship: A comparative analysis of AI-generated text and human academic writing in English literature. *Frontiers in Education, 9*, Article 1347421. https://doi.org/10.3389/feduc.2024.1347421

Badenhorst, C. (2019). Literature reviews, citations and intertextuality in graduate student writing. *Journal of Further and Higher Education, 43*(2), 263-275. https://doi.org/10.1080/0309877X.2017.1359504

Bick, A., Blandin, A., & Deming, D. J. (2026). The rapid adoption of generative AI. *Management Science*. Advance online publication. https://doi.org/10.1287/mnsc.2025.02523

Danylchenko-Cherniak, O. (2024). Philological perspectives of design thinking and artificial intelligence in the Education 5.0-7.0 era.

Filologichni Traktati, 16(1), 3-15. https://doi.org/10.21272/Ftrk.2024.16(1)-3

Digital Education Council. (2024, August 7). *What students want: Key results from DEC Global AI Student Survey 2024*. https://www.digitaleducationcouncil.com/post/what-students-want-key-results-from-dec-global-ai-student-survey-2024

Doru, B., Maier, C., Busse, J. S., Lücke, T., Schönhoff, J., Enax-Krumova, E., Hessler, S., Berger, M., & Tokic, M. (2025). Detecting artificial intelligence-generated versus human-written medical student essays: Semirandomized controlled study. *JMIR Medical Education, 11*, Article e62779. https://doi.org/10.2196/62779

Durak, H. Y., Eğin, F., & Onan, A. (2025). A comparison of human-written versus AI-generated text in discussions at educational settings: Investigating features for ChatGPT, Gemini and BingAI. *European Journal of Education, 60*(1), Article e70014. https://doi.org/10.1111/ejed.70014

Gilardi, F., Di Lorenzo, S., Ezzaini, J., Santa, B., Streiff, B., Zurfluh, E., & Hoes, E. (2024). *Quality perceptions and intended engagement in response to AI-generated and AI-assisted news*. *arXiv*. https://doi.org/10.48550/arXiv.2409.03500

Hajibayova, L., & Lee, S. (2025). Authorship in the age of artificial intelligence. *Journal of Documentation, 81*(5-6), 1401-1414. https://doi.org/10.1108/JD-04-2025-0109

Herbold, S., Hautli-Janisz, A., Heuer, U., Kikteva, Z., & Trautsch, A. (2023). A large-scale comparison of human-written versus ChatGPT-generated essays. *Scientific Reports, 13*(1), Article 18617. https://doi.org/10.1038/s41598-023-45644-9

Hyland, K. (2005). *Metadiscourse: Exploring interaction in writing*. London, England: Continuum.

Jaidka, K., Khoo, C. S. G., & Na, J.-C. (2013). Literature review writing: How information is selected and transformed. *ASLIB Proceedings, 65*(3), 303-324. https://doi.org/10.1108/00012531311330665

Jiang, F. K., & Hyland, K. (2024). Does ChatGPT argue like students? Bundles in argumentative essays. *Applied Linguistics, 46*(3), 375-391. https://doi.org/10.1093/applin/amae052

Jiang, F. K., & Hyland, K. (2025a). Does ChatGPT write like a student? Engagement markers in argumentative essays. *Written Communication, 42*(3), 463-492. https://doi.org/10.1177/07410883251328311

Jiang, F. K., & Hyland, K. (2025b). Metadiscursive nouns in academic argument: ChatGPT vs student practices. *Journal of English for Academic Purposes, 75*, Article 101514. https://doi.org/10.1016/j.jeap.2025.101514

Jiang, F. K., & Hyland, K. (2025c). Rhetorical distinctions: Comparing metadiscourse in essays by students and ChatGPT. *English for Specific Purposes, 79*, 17-29. https://doi.org/10.1016/j.esp.2025.03.001

Kim, H., Baghestani, S., Yin, S., Karatay, Y., Kurt, S., Beck, J., & Karatay, L. (2024). ChatGPT for writing evaluation: Examining the accuracy and reliability of AI-generated scores compared to human raters. In C. A. Chapelle, G. H. Beckett, & J. Ranalli (Eds.), *Exploring artificial intelligence in applied linguistics* (pp. 73-95). Iowa State University Digital Press. https://doi.org/10.31274/isudp.2024.154.06

Kobak, D., González-Márquez, R., Horvát, E.-Á., & Lause, J. (2025). Delving into LLM-assisted writing in biomedical publications through excess vocabulary. *Science Advances, 11*(27), Article eadt3813. https://doi.org/10.1126/sciadv.adt3813

Kong, X., & Liu, C. (2024). A comparative genre analysis of AI-generated and scholar-written abstracts for English review articles in international journals. *Journal of English for Academic Purposes, 71*, Article 101432. https://doi.org/10.1016/j.jeap.2024.101432

Lamm, A. J., Lamm, K., Byrd, A. R., Yazdanpanah, M., Gabler, N. K., Johnson, A. K., Sanders, C. E., Masambuka-Kanchewa, F., & Retallick, M. S. (2025). When machines speak science: Testing consumers' perceptions of AI-generated communication messages. *Journal of Applied Communications, 109*(3). https://doi.org/10.4148/1051-0834.2594

Lund, B., Wang, T., Mannuru, N. R., Nie, B., Shimray, S. R., & Wang, Z. (2023). ChatGPT and a new academic reality: Artificial intelligence-written research papers and the ethics of large language models in scholarly publishing. *Journal of the Association for Information Science and Technology, 74*(9), 1125-1135. https://doi.org/10.1002/asi.24750

Melliti, M. (2024). Using genre analysis to detect AI-generated academic texts. *Diá-logos, 16*(29), 9-27. https://doi.org/10.61604/dl.v16i29.377

Mohamed, Y. A., Mohamed, A. H., Kannan, A., Bashir, M., Adiel, M. A. E., & Elsadig, M. A. (2024). Navigating the ethical terrain of AI-generated text tools: A review. *IEEE Access, 12*, 197061-197120. https://doi.org/10.1109/ACCESS.2024.3521945

Mostafapour, M., Fortier, J. H., Pacheco, K., Murray, H., & Garber, G. (2024). Evaluating literature reviews conducted by humans versus ChatGPT: Comparative study. *JMIR AI, 3*, Article e56537. https://doi.org/10.2196/56537

Moussa, A., Elkhetabi, S., & Eddahmani, H. (2026). Tracing artificial intelligence in EFL student writing: A linguistic approach to identifying machine-assisted text in undergraduate writing. *Arab World English Journal*, Special Issue on Artificial Intelligence, (3), 73-87. https://doi.org/10.24093/awej/AI3.5

Mugaanyi, J., Cai, L., Cheng, S., Lu, C., & Huang, J. (2024). Evaluation of large language model performance and reliability for citations and references in scholarly writing: Cross-disciplinary study. *Journal of Medical Internet Research, 26*, Article e52935.

https://doi.org/10.2196/52935

Nathani, K. R., Nathani, A. M., Delawan, M., Safdar, A., & Bydon, M. (2025). Can artificial intelligence write science? A comparative analysis of human-written and artificial intelligence-generated scientific writings. *Journal of Neurosurgery: Spine, 43*(6), 767-772. https://doi.org/10.3171/2025.4.SPINE25519

Rad, M. H., Farsi, F., Bali, S., Etezadi, R., & Shamsfard, M. (2024). RFBES at SemEval-2024 Task 8: Investigating syntactic and semantic features for distinguishing AI-generated and human-written texts. In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)* (pp. 1912-1917). Association for Computational Linguistics. https://doi.org/10.48550/arXiv.2402.14838

Tarchi, C., Zappoli, A., Casado Ledesma, L., & Brante, E. W. (2024). The use of ChatGPT in source-based writing tasks. *International Journal of Artificial Intelligence in Education, 35*, 858-878. https://doi.org/10.1007/s40593-024-00413-1

Toulmin, S. E. (2003). *The uses of argument (Updated ed.)*. Cambridge, England: Cambridge University Press. https://doi.org/10.1017/CBO9780511840005

Zaitsu, W., & Jin, M. (2023). Distinguishing ChatGPT (-3.5, -4)-generated and human-written papers through Japanese stylometric analysis. *PLOS ONE, 18*(9), Article e0288453. https://doi.org/10.1371/journal.pone.0288453

Appendix A. Representative Human-Authored Literature Review Excerpts and Their AI-Paraphrased Versions

The following paired excerpts provide fuller examples of the rhetorical transformations identified in the main analysis and complement the shorter examples presented in Table 2. The human-authored excerpts are reproduced from the corpus with original language errors retained.

Table A1. Example 1: Informational Densification and Epistemic Flattening

Human-Authored Original	**AI-Paraphrased Version**
"It had six major results: 1. 2. The children produced more kinds of words, symbols. Parents' production also increased and was predicted by their children's nonverbal perception (positive) and the prognosis (negative). 3. The children and parents were matching with each other in lexical and grammatical production. 4. 5. The parents had to produce more types of words and symbols. When both the child and the parent were included, they both contributed greatly to the child's future language. 6. child speech significantly anticipates parental speech. These results attract more attention to the importance of child-based factors in the early language development of children with autism the importance of parental language factors in the early language development of children with autism."	"The study yielded six main findings: children produced a greater variety of words and symbols; parental language production increased and was positively predicted by children's nonverbal perception and negatively by prognosis; there was lexical and grammatical alignment between children and parents; parents produced more types of words and symbols; both child and parent contributions were significant for the child's future language development; and child speech was a strong predictor of parental speech. These findings highlight the significance of both child-based and parental language factors in the early language development of children with autism."

Note. The AI-paraphrased version presents the information in a more condensed and generalised form. Evidence and findings are integrated into a more fluent academic summary, while the links between evidence and interpretation become less explicit. This illustrates both informational densification and epistemic flattening.

Table A2. Example 2: Depersonalisation and Epistemic Flattening

Human-Authored Original	**AI-Paraphrased Version**
"In the study by James, K. & Tanian, S (2002), aim to envisions a desired future scenario for incorporation of oral communication skills into an ideal online learning environment. And the findings were smart to step forward, online learning has been seen by many educational institutions and governments as the panacea to improve the delivery and quality of student learning also the results of this simulated online trial will be compared to the existing face-to-face and print-based distance education courses, for learning outcomes skills, reduction of communication apprehension, willingness to communicate and student satisfaction. The paper outlines a simulated scenario that will be trailed and modified for future use. This study mentioned an important skill like oral learning via online and this subject related to our research and that what we are seeking for."	"James, K. & Tanian, S. (2002) envisioned the integration of oral communication skills into ideal online learning environments. Their findings suggested that online learning is increasingly viewed as a solution for improving educational delivery and quality. The study compared simulated online trials with face-to-face and print-based distance education, focusing on learning outcomes, communication apprehension, willingness to communicate, and student satisfaction. This research highlights the importance of oral skills in online learning, which is central to the current study."

Note. The AI-paraphrased version reduces writer visibility by replacing more situated and personally inflected commentary with detached academic framing. At the same time, evaluative and interpretive language becomes more standardised and formulaic, illustrating both depersonalisation and epistemic flattening.

These excerpts are illustrative rather than exhaustive and were selected because they clearly represent the principal rhetorical transformation patterns identified across the corpus.