English Language Teaching and Linguistics Studies
ISSN 2640-9836 (Print) ISSN 2640-9844 (Online)
Vol. 7, No. 6, 2025
www.scholink.org/ojs/index.php/eltls

Original Paper

How AI Constructs Stance: A Corpus-Based Comparison of Interactional Metadiscourse in Student and AI Essays

Qianyu He[1]

[1] Chongqing University of Education, Chongqing, China

Received: October 13, 2025 Accepted: November 27, 2025 Online Published: December 06, 2025
doi:10.22158/eltls.v7n6p103 URL: http://dx.doi.org/10.22158/eltls.v7n6p103

Abstract

The application of Large Language Models (LLMs) in the writing has prompted a need to scrutinize how those models establish authorship and stance. This study adopts a corpus-based approach to investigate the interactional metadiscourse in AI-generated essays versus human writing, under Hyland's framework. From a parallel corpus of approximately 1.3 million words from the DAIGT V2 dataset, the research combines quantitative frequency analysis with qualitative concordance investigation. The findings reveal that while AI models demonstrate structural proficiency, they exhibit a deficit in rhetorical conviction and reader engagement characterized by a detached stance. Compared to the dialogic nature of human writing, AI models show underuse of boosters and rhetorical questions, and a reliance on vague attitude markers, and hence lack the capability of building writer-reader relationship. These results highlight the deficiency of AI writing's stance-taking and offer corresponding pedagogical implications.

Keywords

AI writing, interactional metadiscourse, stance, corpus analysis

1. Introduction

In current writing environment, the integration of Large Language Models (LLMs) has witnesses a profound increase. Scholars have pointed out that tools LLMs like Chat GPT offer capabilities in generating grammatically sophisticated and well organized passages, pushing educators and learners to reconsider the assessment. (Kasneci et al., 2023; Cotton et al,. 2024). Hence, a growing interest on detecting AI-generated essays, comparing AI-generated essays with human writing is appearing. For instance, Desaire et al. (2023) analyzed scientific writing and argued that AI-generated text is distinct mainly because of its reduced lexical diversity and simpler sentence structures, appealing to develop better detection algorithms. Similarly, Gao et al. (2023) discussed AI's ability to bypass plagiarism detectors. Though existing studies confirm that AI can imitate the pattern of academic writing, their insight into the functional aspects of how AI model establish a voice and create a stance to show authorship is comparatively limited. Some recent pedagogical researches noticed this limitation and started to critique the quality of AI-generated argumentation. Imran and Almusharraf (2023) claimed that while AI-generated texts are structurally organized, they often lack the critical depth and emotional engagement as could be easily found in human writing. Similar researches (e.g., Jiang & Hyland, 2025; Zhang & Zhang, 2025) have been conducted but there is still a scarcity of large-scale, empirical research that quantifies exactly how AI-generated texts fails to engage readers.

Writing proficiency considers more than just linguistic delicacy; instead, it is a social act involving the negotiation and communication of meaning between writer and reader, which is gained through the establishment of writers' stance (Hyland, 2005). Hyland then proposed interactional metadiscourse as a complete set of rhetorical resources which help writers to engage with their audience and create a space for building writer-reader relationship. This study aims to adopt Hyland's classic interactional metadiscourse as the fundamental framework in detecting the establishment of authorship and engagement of readers. Utilizing a parallel corpus derived from the DAIGT V2 dataset, this research tries to investigate whether AI models possess interactional competence and how its writing strategies are different from that adopted by human writing.

2. Methodology

A corpus-based approach will be adopted to compare interactional metadiscourse devices in human and AI-generated writing, through the construction of the parallel corpus and the subsequent combination of quantitative and qualitative examination of stance markers based on Hyland's (2005) framework.

2.1 Analytical Framework: Interactional Metadiscourse

Metadiscourse refers to the linguistic devices that writers use to evaluate their discourse, engage readers, and shape the writer-reader relationship within a text. It includes two dimensions: interactive and interactional (Hyland, 2005). The Interactive dimension (e.g., transitions, frame markers) create the flow of information, which is the exact strength of AI tools; the interactional dimension is the one that demonstrates how writers project their stance and address readers. Hence, this study puts a particular focus on the interactional dimension.

Interactional metadiscourse consists of five subcategories: hedges, boosters, attitude markers, self-mentions, and engagement markers (see Table 1 for detailed information). These devices are used strategically by writers to express attitudes, signal their degree of commitment, and involve readers in the argument. Examining these devices allows us to compare how human writers and AI systems establish stance, position themselves in relation to readers, and attempt to engage readers in the reasoning process.

Table 1. Interactional Metadiscourse Resources (Hyland, 2005)

Category	Function	Examples
Interactional	**Involve the reader in the text**	**Resources**
Hedges	Withhold commitment and open dialogue	*Might; perhaps; possible; about*
Boosters	Emphasize certainty or close dialogue	*In fact; definitely; it is clear that*
Attitude markers	Express writer's attitude to proposition	*Unfortunately; I agree; surprisingly*
Self-mentions	Explicit reference to author(s)	*I; we; me; our*
Engagement markers	Explicitly build relationship with reader	*Consider; note; you can see that*

2.2 Corpus Construction and Data Analysis

The corpus is sourced from the DAIGT V2 dataset, which compares argumentative essays written by U.S. students with AI-generated essays in parallel writing topics. To ensure a balanced comparison and sufficient tokens, human-authored and AI-generated texts are extracted and aligned. The final corpus comprises approximately 1.3 million words, with 759,540 words in the Human sub-corpus and 555,072 words in the AI sub-corpus. Then, a Python script is developed to tokenize the texts and calculate the frequency of interactional metadiscourse markers (hedges, boosters, attitude markers, self-mentions, and engagement markers). All raw frequencies are normalized to per 1,000 words and a corresponding comparison between the two groups will be made. The analysis is conducted by a sequential order: a quantitative tendency will be identified, followed by a qualitative concordance analysis. This involved the examination of specific sentence structures containing high-frequency markers to determine their rhetorical functions within the essays.

3. Results and Discussion

3.1 Overview

To gain a general understanding of how human students and AI models construct interactional stance respectively, we analyzed the frequency of five key metadiscourse categories. Table 2 presents the normalized frequencies (per 1,000 words) of the interactional dimension's usage.

Table 2. Frequency of Interactional Metadiscourse Ctegories (per 1,000 words)

Interactional Dimension	Human Writers (Freq.)	AI Writers (Freq.)	Difference
Hedges	9.06	9.63	+6.3%
Boosters	4.01	1.75	-56.4%
Attitude Markers	1.38	2.71	+96.4%
Self-mentions	15.37	17.35	+12.9%
Engagement Markers	15.28	10.12	-33.8%

Total Interactional	45.10	41.56	-7.8%

The table generally conveys a complex rhetorical landscape. AI models have shown the features of academic stance demonstrated by high frequencies of almost each interactional metadiscourse devices, but they also show sharp disparity in conviction (boosters) and reader engagement (engagement markers). To be more specific, some obvious contrasts are as follow.

While both groups employ hedges at similar rates, a striking difference exists in boosters, where human writers use boosters (e.g., *always, know, fact*) at a rate more than double that of AI-generated essays. This suggests that human authors write with a stronger awareness of conviction, often framing their arguments as absolute truths, while AI models prefer to soften claims to avoid being too definite.

Unlike the conventional expectation of being neutral, AI-generated essays employ attitude markers nearly twice as frequently as humans. However, this high frequency is not a sign for emotional stance, but results from the repetition of vague or general adjectives for evaluation (mainly *important* and *essential*).

Similarly, AI models show a somewhat unexpected by applying more self-mentions than human students. The result suggests that AI models are capable of imitating the explicit structure of opinion (e.g., *I believe*), and even brings it to a higher frequency than the actual human authors.

As for the engagement markers, human writers significantly outperform AI in engaging potential readers. A critical example in this gap lies in the use of rhetorical questions. The corpus shows that human writing texts contain 1,184 question marks, whereas the entire AI corpus contains only 265, which implies a profound difference of how the two groups provoke thought and bridge the writer-reader distance.

In summary, the quantitative overview indicates that the disparity between the two groups of writing mainly lies in the extent of assertiveness and direct writer-reader dialogue. The following sections will provide a detailed analysis on those dimensions.

3.2 Hedges and Boosters

Hedges and boosters function as the primary tools for writers to modify their commitment to a claim. The comparative frequency of key hedges reveals a distinct difference in how human writing and AI-generated writing manage this commitment in their own way, as demonstrated in Table 3 and Table 4. While the overall frequency of hedges is comparable between the two groups, the use of boosters shows a critical difference in rhetorical assertiveness, with human writers using them at more than double the rate.

Table 3. Comparative Frequency of Key Hedges (per 1,000 words)

Lexical Item	Human Freq.	Human Rank	AI Freq.	AI Rank
could	2.30	1	1.44	3

may	1.78	3	2.62	1
believe	0.81	5	1.88	2

In the above table, it can be seen that human writers heavily favor the modal verb *could* (Rank 1). Sample sentences show that those writers use *could* to construct hypothetical narratives or describe potential capabilities in concrete scenarios, like "Lowering car usage *could* be a way for the government to create more bicycle paths." On the contrary, AI-generated texts prefer *may* and the verb *believe*, which reflects an academic tone regarding probability in a more modified way. This modification may create a wider space for writer-reader communication, or at least acts as a device to explicitly label generated text as opinion rather than fact, likely a result of safety alignment training.

Table 4. Comparative Frequency of Key Boosters (per 1,000 words)

Lexical Item	Human Freq.	Human Rank	AI Freq.	AI Rank
know	1.17	1	0.58	1
always	0.91	2	0.26	3
actually	0.54	3	0.15	-
fact	0.38	5	0.26	2

The divergence is even more profound in the use of boosters, where the lexical choices suggest different approaches to certainty. As shown in Table 4, the most telling difference is the prevalence of the absolute word *always* in human writing that is used to emphasize their points through hyperbole. Human writers frequently use *always* to convey the intensity of the experience (e.g., "When I had a car I was *always* tense."). Similarly, human writers tend to use the conversational booster *actually,* a dynamic strategy almost absent in AI text, to correct expectations, as in the sample sentence "It *actually* makes a difference".

Though the word *know* is the top boosters for both groups, human writers usually use it with the first and second person pronouns to strengthen personal authority (e.g., "I *know*" "You *know*"), AI models tend to use the it to reach general consensus (e.g., "We *know* that…"). Also, when the AI models need to show their certainty, they prefer evidential nouns like *fact* (e.g., "This ignores the *fact* that...") or formal verbs like *demonstrate*. This highlights a more evidence-based certainty in AI-generated writing.

3.3 Attitude Markers

Attitude markers indicate the writer's affective, rather than epistemic, attitude to propositions, encoding an explicit positive or negative value that is gradable (Hyland, 2005, p. 149). The overall investigation on attitude markers reveals a trend in which AI models apply this device at a frequency nearly double that of human writers; however, a detailed lexical analysis presents that this high frequency is not driven by a rich emotional range, but by the repetitive utilization of a set of evaluative adjectives. See Table 5.

Table 5. Comparative Frequency of key Attitude Markers (per 1,000 words)

Lexical Item	Human Freq.	Human Rank	AI Freq.	AI Rank
important	0.59	1	1.84	1
essential	0.06	-	0.54	2
agree	0.24	2	0.06	-
hope	0.22	3	0.15	5

The most dominant feature in the AI models' attitude marker usage is the adjective *important*. Though this word ranks as the top one frequently used attitude marker in both groups, AI models use it at a rate more than three times that of humans, often considering it as a structural device to organize text rather than to express genuine emotion. Using this item, common phrases in AI writing are like "It is *important* to note" or "It is *important* for individuals to be mindful". Furthermore, this word can be replaced or be treated equally with its synonym *essential*, which appears shockingly nine times more frequently in AI text than in human text. The two words may indicate a formulaic approach to evaluation instead of real attitude expression.

Conversely, AI-generated texts show obvious limitation in markers of human desire and alignment, such as the limited usage of *hope* and *agree*. Human writers frequently use *agree* to align themselves with a side of the argument, stating explicitly, "I *agree* with the idea that...". Likewise, the verb *hope* is applied by human writers to build their arguments or to appeal to the readers' conscience, such as in the sample sentence "I *hope* you take this information to good use". When AI models use *hope*, the context is typically confined to polite closings like "I *hope* this helps", lacking the actual attitude expression found in human essays.

3.4 Self-mentions

Self-mention refers to the degree of explicit author presence in the text measured by the frequency of first-person pronouns and possessive adjectives (me, mine, exclusive we, our, ours) (Hyland, 2005, p. 148). They play a crucial role in establishing authorial identity and stance. Contrary to the assumption that AI writing is impersonal, the overall data shows that AI models in fact employ self-mentions at a higher frequency than human writers (Table 2); however, when we break down the detailed usage, it can be found that human writers and AI models show distinct ways to construct the writer identity. See Table 6.

Table 6. Comparative Frequency of Dominant Self-mentions (per 1,000 words)

Lexical Item	Human Freq.	Human Rank	AI Freq.	AI Rank
I	5.07	1	5.68	1
We	4.56	2	4.17	2
My	1.13	4	1.75	4
Our	2.56	3	4.64	3

AI writers show a slightly more frequent use of the singular pronoun *I* than human writers, but sample sentences suggest that AI models' use of *I* is mainly structural. They typically use this word within phrases such as "*I* believe" "*I* think" or "*I* urge", like in the sample "*I believe* that this system is outdated". In this context, the word *I* functions as a stance marker rather than an identity marker. In contrast, human students use *I* to establish a narrative identity, frequently using this word to recall specific personal experiences, like in sentence patterns "When *I* was in school..." or "*I* remember a time when...". Hence, for human writers, this word helps with building argument in reality, while for AI models, this word serves to fulfill the structural requirement of stating a position.

The similar disparity can be observed in the use of the plural pronoun *we*. Despite that human writers use this pronoun just slightly more than AI models, the pronoun's referent differs. In human writing, phrases like "*we* students" or "*we* teenagers" imply that they typically use *we* to refer to the specific community to which they belong; however, AI models prefer to use *we* (and its possessive *our*) to refer to a generic, universal humanity, as found in sentences like "*We* need to protect *our* planet" or "*We* can facilitate communication". This helps to create an effect of inclusivity, allowing the non-human writers to simulate a collective voice, while lacking the specific group membership that characterizes human writing.

3.5 Engagement Markers

Engagement markers reefer to the various ways writers bring readers into the discourse to relate to them and anticipate their possible objections (Hyland, 2005, p. 151). Unlike other categories in interactional metadiscourse where the AI models' use was higher, engagement markers stand for the area of the most profound deficit in AI-generated writing in addressing readers, since AI models employ these markers at a much lower frequency compared with human writing. See Table 7 for the detailed observation.

Table 7. Comparative Frequency of Engagement Markers (per 1,000 words)

Lexical Item	Human Freq.	Human Rank	AI Freq.	AI Rank
you	6.30	1	4.25	1
your	2.52	2	3.02	2
? (Questions)	2.60	3	0.52	-

think	1.84	4	0.54	-
consider	0.15	-	0.76	3

The most profound disparity in reader engagement lies in the use of rhetorical questions. As can be seen in Table 7, this type of question, indicated by question marks, are frequently used in human writing, appearing 2.60 times per 1,000 words; in contrast, they are rare in AI-generated texts. Human writers show a strong inclination to raise questions to provoke thought, come up with the texts' central claim, and bridge the gap between writers and readers, as in the sample sentence "Have you ever thought about...*?*". The absence of the equivalent strategy in AI writing leads to a comparative monologue, where the information is just delivered to readers rather than negotiated or communicated with them.

Additionally, while both groups rely on the second-person pronouns *you* and *your*, the rhetorical function of these words differs. Human writers often use *you* to place readers in a hypothetical scenario or to appeal to shared experiences (e.g. "*You* might think that..."); AI models tend to use it in generic directives, as in the sample sentence "Consider the impact on *your* community".

Furthermore, AI models prefers formal directive verbs like *consider* over the more cognitive verbs preferred by humans, such as *think*. This implies that human writing is more conversational and cognitively engaging, whereas AI writing is more instructional and directive.

4. Conclusion

This study sets out to investigate the interactional metadiscourse application of AI-generated essays compared to those written by human students. Through the analysis of a parallel corpus of approximately 1.3 million words, the findings reveal a closer understanding of AI writing capabilities. Though Large Language Models have to some extents mastered the macro-level mechanics of academic interaction by employing hedges and self-mentions at frequencies comparable to or exceeding human writers, they still show profound deficits in rhetorical conviction and reader engagement.

The primary finding is that AI-generated writing is featured by structural proficiency but rhetorical detachment. The significant underuse of boosters (e.g., *always, never*) and the formulaic overuse of general attitude markers (e.g., *important*) indicate that AI models prioritize neutrality and safety over the persuasive force required in argumentation, even trying to minimize their persuasion to engage readers. Different from human writers who ground their stance in personal narrative and absolute conviction, AI models establish a stance that is detached and risk-averse. Also, the utilization of pronouns and engagement markers reflects a disembodied voice in AI texts. Even though they frequently use *I* and *we*, those pronouns act more like structural signposts for instruction rather than symbols for writer identity or community membership. On the other hand, the lack of rhetorical questions and the rarity of dialogic markers like *agree* further enhance AI writing's alienation from potential readers, through avoiding the writer-reader conversation.

These findings may provide implications for writing pedagogy. First, educators should notice that AI-generated writing often reflect a distinct register that lacks writer's voice compared to human writing, no matter how grammatically flawless it is. If students turn to AI models in writing without proper revision, this detached writing style may be internalized, and the importance of creating writer-reader relationship may be ignored. Second, in the field of AI writing detection, we suggest that the linguistic examination should move beyond syntax to investigate the process of establishing stance and engagement, paying special attention to those interactional metadiscourse devices presented before. Specific overuse or underuse of certain markers should be avoided. Future research could expand this analysis to more diverse genres beyond argumentative essays to see whether these patterns AI-human writing disparity persists across different types of discourse.

References

Cotton, D. R., Cotton, P. A., & Shipway, J. R. (2023). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, *61*(2), 1-12.

Desaire, H., Chua, A. E., Isom, M., Jarosova, R., & Hua, D. C. (2023). Distinguishing academic science writing from humans or ChatGPT with over 99% accuracy using off-the-shelf machine learning tools. *Cell reports Physical science*, *4*(6), 101426.

Gao, C. A., Howard, F. M., Markov, N. S. et al. (2023). Comparing scientific abstracts generated by ChatGPT to real abstracts with detectors and blinded human reviewers. *npj Digital Medicine*, *2023*(6), 75.

Hyland, K. (2005). *Metadiscourse: Exploring Interaction in Writing*. London: Continuum.

Imran, M., & Almusharraf, N. (2023). Analyzing the role of ChatGPT as a writing assistant at higher education level: A systematic review of the literature. *Contemporary Educational Technology*, *15*(4), ep464.

Jiang, F. K., & Hyland, K. (2025). Rhetorical distinctions: Comparing metadiscourse in essays by ChatGPT and students. *English for Specific Purposes*, *2025*(79), 17-29.

Kasneci, E., Seßler, K., Küchemann, S. et al. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, *2023*(103), 102274.

Zhang, M., & Zhang, J. (2025). Disciplinary variation of metadiscourse: A comparison of human-written and ChatGPT-generated English research article abstracts. *Journal of English for Academic Purposes*, *2025*(76), 101540.