MEGA INSTITUTE FOR ADVANCE RESEARCH AND DEVELOPMENT

Journal of Arts and Linguistics Studies
ISSN Print 2959-0957 ISSN Online 2959-0965
Volume 4 Issue 1 2026
Pages 387-402
Open Access

A Comparative Discourse Analysis of AI Generated and Human-Written Texts: Coherence, Framing, and Persuasion

Publication Details
Paper Received:
December 24, 2025
Paper Accepted:
Janurary 28, 2026
Paper Published:
Janurary 30, 2026

Masooma Amjad
M.Phil Scholar, Department of English, University of Azad Jammu and Kashmir. Email: masoomaqureshi09@gmail.com
Hamna Tanveer
M.Phil Scholar, Department of English, National University of Modern Languages, (NUML), Islamabad.
Emal: hamna.tanveer000@gmail.com

Article DOI: https://doi.org/10.71281/jals.v4i1.627

Abstract
The rapid growth of artificial intelligence (AI) in natural language generation has reshaped how written discourse is produced, assessed, and consumed. Even though AI-generated texts are becoming similar to human-written texts, there are still concerns regarding their discursive richness and communicative efficiency. The current study is the critical evaluation between AI-generated and human-written texts, which is based on three dimensions: coherence, framing, and persuasion. Moreover, this study analyzes how texts are framed to construct meaning, leading to certain interpretations, and their overall impact on audiences. It is in contrast with other studies that mostly evaluate linguistic accuracy on superficial level. A mixed-methods research methodology, which integrated qualitative discourse analysis with systematic and structured evaluation measures, was employed to analyze a corpus of 60 texts in the academic, journalistic, and argumentative genres. The findings illustrate that texts generated through AI possess a significant amount of grammatical and local coherence, yet they fail to effectively construct compelling arguments or cohesively connect various ideas. Humanly composed texts, on the contrary, tend to be logical and coherent as a whole due to the way they show how thoughts are arranged and the way underlying themes progress ahead. Regarding framing, AI-written texts are more prone to delivering information in general terms, whereas human writings can build interpretive frames, relying on context and communication objectives. The AI-generated texts are convincing yet lack adaptive variation and, therefore, unable to capture the attention of the audience. This paper explores that AI-generated discourse is well-knitted through discourse markers at linguistic level, yet it is limited in its capacity to act as a dynamic communicative process and engage the learners. The results of this research study can be used in the application of AI in education, media, and knowledge production because it is evident that one must approach machine-generated text with critical thinking.

Keywords: AI, discourse, coherence, framing, persuasion, rhetorical analysis, human writing.

Published by MIARD. Copyright: © the author(s). This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY NC) license (https:// creativecommons.org/licenses/by/ 4.0/).

1. Introduction

The rapid integration of artificial intelligence in textual production has transformed contemporary communication that transcends the technological progress. The use of AI-generated texts has gained popularity in academic writing, journalism, business communication, and digital content creation. Massively-driven language models can produce grammatically correct, stylistically coherent, and contextually suitable text. This has entrenched them in everyday communicative activities.

Still, knowing that more and more people are using AI-generated text raises serious questions about the nature of discourse. Writing isn't just putting words together in the right way. Instead, writing is the process of putting ideas together, making sense of them, and getting them across to other people. How people write is affected by their thoughts, the social setting, and their communicative goals. Authors determine the arrangement of their arguments, the focal points of a subject, and the interpretation of information. These choices are based on what they know, how much experience they have, and how well they understand the needs of the audience.

Machine-generated text, on the other hand, comes from computer systems that copy language patterns instead of using mental processes to make meaning. Language models generate text by forecasting word sequences as per statistical possibilities using large collections of data. Although this enables them to generate fluent and coherent sentences, it does not furnish them with the ability to make

intentional meaning. This difference between meaning construction and pattern generation is the focus of the current research.

This study aims to investigate the differences between AI-generated and human-written texts in terms of discourse. The study explores the ways in which texts are communicative acts as opposed to grammatical correctness or stylistic similarity. In particular, it breaks down three dimensions, namely, coherence, framing, and persuasion.

Coherence is the logical organization of ideas in a text. It is not only the application of cohesive devices but also the logical development of arguments and the connection of various parts into one organic whole. Coherence in human writing is obtained by planning and revision. Writers make sure that every section of a text has its contribution to a greater goal and that ideas are interrelated in significant ways. Coherence in AI-generated text is created by learned language patterns, and can lead to structurally coherent but conceptually discontinuous discourse.

Framing involves the selection and presentation of information in ways that shape interpretation. It is not merely what is said but how it is said. Human writers use framing to bring readers to certain interpretations, which are, in many cases, based on their perceptions and communicative purposes. Framing has a direct relation with the context since it depends on the cultural knowledge, situational analysis, and the expectation of the audience. The texts generated by AI, without possessing their own point of view, are grounded on the patterns present in the training data, which can result in generalized or neutral representations.

Persuasion is the ability of writing to influence faith, attitudes or understandings. It involves the use of rhetorical strategies including positioning and engagement strategies as well as reasoning. Human authors are adaptable, receptive and able to alter their persuasion strategies according to the audience and

situation. Despite the ability of AI-generated texts to produce persuasive language, it is possible that they do not offer this flexibility.

There are far-reaching implications of increasing the use of AI-generated text across different disciplines. In education, it raises issues about the development of writing skills and the significance of technology in education. Within journalism, it can influence how information is produced and reported, which can influence how information is perceived by general population. It changes the process of creation and consumption of content in professional communication. These implications are not mere practical implications, but also epistemological, in the sense that they concern the process of knowledge building and verification.

Determining whether AI-generated texts are better or worse than human-written texts is not the main objective of the research. Rather, it seeks to identify the distinctions between their discursive characteristics. The study highlights coherence, framing, and persuasion in order to outline the advantages and disadvantages of AI in textual production.

It is noteworthy that the study is comparison-based, focusing on functional and structural distinctions rather than superficial similarities. This method will enable a thorough examination of discourse produced by AI, emphasizing both its parallels with and differences from human writing.

The results of this paper are meant to contribute to the current discussion about the importance of AI in communication. As AI systems continue to evolve, it is essential to develop a precise understanding of their capabilities and constraints. In situations where the quality of the discourse is crucial, this data can be used to ensure that AI is used appropriately.

2. Literature Review

Coherence in discourse analysis underwent a crucial change within its structural and cognitive frameworks. Early linguists regarded it as a measurable trait of the

text, focusing on syntactic constructions i.e. grammatical and lexical linkages between words to create sentences. Later studies asserted that coherence is not an intrinsic feature of a text; rather it arises from the readers' interpretation by their engagement with the text. This approach affirms the notion that coherence comes from the connection between textual cues and cognitive operations operating in readers' mind such as inference and schema activation. This is particularly essential when encountering AI-generated text, where such means of evaluating coherence are not sufficient. While AI-generated texts may exhibit notable coherence, studies suggest they fail to uphold interpretive coherence, especially in more multifaceted argumentative contexts (McNamara et al., 2010; Crossley et al., 2016).

Latest research mostly focuses on how coherence is perceived rather than its measurement in fixed statistics. The research shows that readers detect slight inconsistencies in AI-generated text, despite its superior performance on computational metrics. These discrepancies appear due to abrupt changes in the topics, poorly constructed arguments, or ideas that aren't well crafted. These results depict that coherence is a qualitative phenomenon instead of quantitative. It has brought about the necessity to introduce more unified approaches that entail using computational tools and human judgment in evaluating the quality of discourse (Ippolito et al., 2020; Clark et al., 2021).

Framing is the concept that has been discussed extensively in the sphere of communication as the process by which texts shape the interpretation. Rather than just presenting information, texts are also dynamic in the process of constructing meaning, through organizing and emphasizing certain areas of a topic. Research has found out that framing does not simply influence what the readers understand but also their evaluation and response to information. The concept of framing is particularly important when it comes to AI-generated text because AI systems do

not possess their own interpretative frameworks. Rather, they mimic training data patterns, which can be biased vis-a-vis dominant opinions but removes all others. This brings up doubts about how varied and inclusive AI-generated discourse is (Entman, 1993; Bender et al., 2021).

Research on biased algorithms also backs up the idea that framing is important in AI-generated text. Language models may replicate social biases in training data, which affects how topics are represented. It indicates that the organization of AI-generated text is not neutral; it is affected by the data structures that are behind it. Unlike human authors, who can intentionally alter their framing strategies, AI systems do not have the ability for critical reflection on these strategies, making it difficult to address such biases solely through text generation (Crawford, 2021; Bender et al., 2021).

Persuasion in discourse has grown to encompass not only rhetorical structures but also audience-centered perspectives. Recent research demonstrates that persuasion is dependent on the ability to adjust arguments to the particular context and audience. Experimental research comparing human-produced and AI-generated persuasive texts has found that AI can create solid arguments, but it rarely exhibits variability and responsiveness. It leads to persuasive content that is less interesting but more consistent. The lack of adaptive strategies constrains the effectiveness of AI-generated persuasion, especially in contexts necessitating nuanced communication (Tan et al., 2016; Floridi and Chiriatti, 2020).

The other major area of research is the concern of stylistic variation. Human writing is heterogeneous in terms of tone, structure and expression and it is a manifestation of personal opinion and imagination. In turn, AI-generated text can be more predictable and consistent. This consistency may render it easier, but it may also reduce originality and engagement. The stylometric analyses have also shown that the texts generated by the AI are more likely to be placed in the

narrower range of styles, and this can affect how the readers perceive them (Gehrmann et al., 2019; Ippolito et al., 2020).

The research on the perception of readers can also help shed light on the rating of the texts created by AI. It has been demonstrated that readers are sensitive to differences in depth, engagement and authenticity even in linguistically similar texts. These findings show that the quality of discourse cannot be determined only by the grammatical correctness but also by the performance of a text in reference to the interaction with the audience. This justifies the relevance of the discourse analysis on different levels, including structure, meaning, and reception (Clark et al., 2021).

3. Methodology

The current study is grounded in a mixed-methods approach of research to compare AI-generated and human-written texts regarding three analytical dimensions of coherence, framing, and persuasion. The study analyzed a corpus of 60 texts, comprising 30 AI-generated and 30 human-generated. To ensure variety in discourse structure and communicative intentions, the texts were of three genres, including academic essays, opinion articles, and informational reports.

Standardized prompts were used to create the AI-generated texts to make them consistent in terms of scope. The sources of human-written texts were peer-reviewed journals, existing media sources, and writing samples of academic literature. To ensure that all texts were comparable, all the texts were standardized to a similar length (approximately 900-1000 words) and were scaled to remove extraneous formatting variations.

The analytical framework was structured in such a way that it can represent both structural and functional elements of discourse. Coherence was assessed by structural mapping, which was concerned with the introduction of ideas, their development, and integration throughout the text. The analysis of framing was

conducted through the detection of emphasis patterns, thematic organization, and interpretive positioning. Persuasion was evaluated by analyzing the argumentative strategies, such as reasoning, positioning, and engagement strategies.

The qualitative analysis was done by two independent coders to increase reliability. The analysis was preceded by the development of coding guidelines, and the inconsistencies were discussed. Cohen's kappa (k) was used to calculate inter-coder agreement and the result was 0.82 which is very high in terms of reliability.

The quantitative analysis was performed by giving marks to the texts according to the predetermined standards of coherence, framing, and persuasion. Patterns in AI-generated and human-written samples were compared using these scores.

The ethical issues were addressed by anonymizing all man-made texts and making the use of AI-generated material transparent. All the resources used were publicly available or created to be used in the research.

4. Discussion and Analysis

The relative comparison presented in this paper shows that the difference between AI-generated and human-written texts cannot be appropriately captured using traditional linguistic metrics like grammatical correctness or lexical cohesion. Rather, the most significant divergences are found at the discourse functionality level, that is, in the way texts create and formulate meaning, set the interpretive direction, and address the readers throughout the process of argumentation. Such distinctions are not just technical but indicate more profound differences in processes of text production.

4.1. Structural Coherence as Process vs. Output

The results show that AI-generated texts attain coherence mainly due to the regularity of patterns but not due to a dynamic compositional process. This is a

crucial difference. Coherence is not static in human writing; rather, it is a fluid element that makes up the process of making decisions through repetition. Writers keep changing their arguments, making the links between ideas stronger, and thinking about how readers will understand transitions that occur. This process-based coherence not only results in internally consistent texts, but also provides a directional momentum.

The AI-generated texts exhibit what can be called "sequential coherence," which implies that every component is correct in its own respect, but the entire sequence does not always show a distinct progression. This is especially noticeable in argumentative writing, where AI-generated texts are more likely to give multiple valid arguments, but they do not clarify how these arguments relate to each other in terms of their significance. This order of preference triggers a kind of discursive flattening, where all ideas are interpreted in the same way, thereby rendering the argument ambiguous.

Texts produced by AI are also inclined to employ transitional markers (such as "however," "moreover," and "in conclusion") as signs of coherence rather than establishing a logical order.

Although these cues produce the illusion of organized argumentation, they may not always be associated with substantive changes in thought. Transitions are more selective and meaningful in human-written texts, which tend to entrench logical relationships in the structure of the argument itself rather than indicate them.

This difference implies that AI-generated discourse coherence is mostly syntactic and positional, and human discourse is conceptual and intentional. The consequences are significant: a text can be coherent but without the structure needed for effective communication.

4.2. Framing as Interpretive Control vs. Representational Distribution

The framing analysis shows that texts written by human beings have a consistent ability to produce and sustain interpretive control. This is done by positioning the topic early, selective focus on key aspects, and maintaining consistency between introduction, development and conclusion. The framing in most human-written texts is not only a characteristic of individual passages but a textual property, which determines the manner in which information is presented, developed, and resolved.

AI-produced texts, however, are more likely to spread information throughout the text more evenly without creating a dominant interpretive viewpoint. The outcome of this is what can be termed as representational distribution where various facets of a subject are brought out without any apparent prioritization or judgmental position. Although this method can produce a sense of balance, it tends to decrease interpretive clarity. Information is given to the readers but there is little indication on how to incorporate or assess it.

Another important difference lies in the temporal stability of framing. Human authors tend to keep the frame of the text consistent, even in cases where they present counterarguments or opposing points of view. This uniformity allows readers to understand new information in accordance with what they already knew. However, AI-generated texts may have slightly different frames in various parts, depending on underlying information patterns rather than purpose. This type of change might make it more difficult to understand the text because the reader has to keep adjusting their interpretation of where the text is headed.

In addition, human framing typically includes contextual information that places the subject matter in a larger social, cultural, or disciplinary context. These clues make the text deeper and more relevant, enabling the reader to connect the text with other knowledge. Though AI-written texts can refer to such contexts,

they are generic, not specific and integrated. This constrains their capacity to create significant interpretive spaces.

The outcomes suggest that framing in human discourse is an interpretive control mechanism, whereas framing in AI-generated discourse is an information arrangement mechanism. This reduces the capacity of AI-generated texts to guide readers' interpretation in a consistent and long-term manner because of the absence of intentional framing.

4.3. Persuasion as Adaptation vs. Repetition

Persuasion is the area where the differences between texts created by AI and humans are most noticeable. Human-written texts can practice adaptive persuasion, adjusting their strategies in accordance with the dynamically evolving structure of the argument and the anticipated reactions of an audience. Such flexibility may be observed in the use of different tones, the changes of emphasis, and the sensible usage of examples or analogies.

Conversely, AI-written texts are repetitively convincing i.e. they contain the same generalized argument structures. They are more likely to present balanced arguments, listing strengths and weaknesses, and rely on abstract arguments. Although these techniques are rhetorically correct, they are repetitive and lack variety, which makes them less persuasive.

The most significant limitation of the AI-generated persuasion is that no situational calibration is done. Human authors base their arguments on the situation and consider the knowledge of the audience, their expectations, and counterarguments. This results in arguments seen to be personal and reactive. Contrarily, AI-generated texts perform in a generalized communicative paradigm, addressing an implicit audience and making no explicit differentiation. This is not specific and restricts their capability of reaching out to the readers on a more personal level.

Another important factor is the integration of evidence. Human writers are more likely to choose and apply evidence strategically, depending upon the structure of their argument. Evidence is not offered in isolation; it is examined and contextualized in order to support specific statements. AI-generated texts might have the necessary information, but they are likely to display it in a more consistent manner, often lacking the ability to connect it to the argument's priorities. It renders the content weaker and less persuasive.

Furthermore, human persuasion often means making things more forceful over time, which implies arguments get more persuasive and more specific. This isn't something that happens often in AI-generated texts, which tend to stay at the same level of abstraction. These results lead to arguments that are sound but not engaging, and this makes it hard for people to stay interested.

4.4. Reader Engagement and Discursive Momentum

Another thing that comes up in the analysis is the idea of discursive momentum, which is the extent to which a text makes a reader feel like they are moving forward and wish to keep reading. The hand-written texts are generally more inclined to build momentum by the way they are put together, the way the sentences are arranged, and a novel point of view. This makes the readers feel like they are in a flow state, which keeps them engaged.

AI-generated texts are easy to read and have a more consistent structure which can reduce momentum. The sentence structure and form of discourse may lead the reader to a sense of stagnation, especially in longer texts. The presence of authorial voice also influences readers' engagement in reading a particular text as it may be especially enticing in novels and short stories. The choice of words, the style of the author, the utilization of evaluative language, and the way the writer positions themselves relative to the reader all contribute to create a feeling of presence in their work. This existence consolidates the connection between the

reader and the author, which makes the reading more engaging and enticing. AI-generated texts frequently do not have a clear authorial voice and may seem impersonal, which could make them less effective and lead to monotony.

5. Limitations

Despite its contributions, this study suffers some limitations. First, the material, although balanced in terms of genres, is very small and limited. The selection of the 60 texts provides a relatively manageable corpus which is studied in detail however it does not always represent the range of discourse practices across various disciplines and cultural backgrounds. Furthermore, research may expand the sample to multilingual texts and other types of writing, including creative and technical writing.

Second, although the research is a mixed-methods study, the qualitative part of it is based on interpretive analysis, which can be subjective. Despite the measures that were taken, including inter-coder reliability, it is hard to be completely objective in discourse analysis.

Third, the AI-written texts were generated using the assistance of a single model in a controlled environment. Different discursive features can be present in various models, and a comparative study would be conducted in multiple systems, which would provide a deeper understanding.

Finally, the textual production is analyzed without reader-response information. To contribute to the existing findings, future research can concentrate on how AI-generated and human-written texts are perceived and evaluated by various audiences.

6. Conclusion

This paper demonstrates that the difference between AI-generated and human-written texts can be more effectively observed not in the linguistic accuracy but in the discursive role. AI system productions are grammatically correct and

structurally accurate. On the contrary, humanly composed texts are more suggestive and convey a plethora of interpretations to the audience involving persuasion and the intended organization of thoughts. Human-written texts are intended to achieve desirable outcomes.

Additionally, the analysis underscores that AI-generated texts have strengths in clarity in their production, consistency, and accessibility which can be useful to perform the tasks that involve effective and smooth delivery of information to the readers. However, they fail to contribute meaningfully to complex argumentation, direct interpretation, and engage audiences in an enthusiastic manner unlike human writers. These limitations cannot be labeled as technical, rather show the absence of cognitive processes and contextual cues which are the essence of human discourse.

The results confirm the notion that AI as a machine can facilitate rather than replace human cognition and critical thinking abilities. Its integration into communicative practices should be informed by the understanding of its strengths and weaknesses. As Artificial intelligence evolves with the passage of time, more research will identify how these limitations might evolve and how the collaboration between humans and AI can be optimized.

References

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 610–623. https://doi.org/10.1145/3442188.3445922

Clark, E., August, T., Serrano, S., Haduong, N., Gururangan, S., & Smith, N. A. (2021). All that's 'human' is not gold: Evaluating human evaluation of generated text. Proceedings of ACL. https://doi.org/10.18653/v1/2021.acl-long.565

Crawford, K. (2021). Atlas of AI: Power, politics, and the planetary costs of artificial intelligence. Yale University Press. https://doi.org/10.2307/j.ctv1vbd275

Crossley, S. A., Kyle, K., & McNamara, D. S. (2016). The tool for the automatic analysis of text cohesion (TAACO): Automatic assessment of local, global, and text cohesion. Behavior Research Methods, 48(4), 1227–1237. https://doi.org/10.3758/s13428-015-0619-3

Entman, R. M. (1993). Framing: Toward clarification of a fractured paradigm. Journal of Communication, 43(4), 51–58. https://doi.org/10.1111/j.1460-2466.1993.tb01304.x

Floridi, L., & Chiriatti, M. (2020). GPT-3: Its nature, scope, limits, and consequences. Minds and Machines, 30(4), 681–694. https://doi.org/10.1007/s11023-020-09548-1

Gehrmann, S., Strobelt, H., & Rush, A. M. (2019). GLTR: Statistical detection and visualization of generated text. Proceedings of ACL. https://doi.org/10.18653/v1/P19-3019

Halliday, M. A. K., & Hasan, R. (1976). Cohesion in English. Longman.

Ippolito, D., Duckworth, D., Callison-Burch, C., & Eck, D. (2020). Automatic detection of generated text is easiest when humans are fooled. Proceedings of ACL. https://doi.org/10.18653/v1/2020.acl-main.164

McNamara, D. S., Graesser, A. C., McCarthy, P. M., & Cai, Z. (2010). Automated evaluation of text and discourse with Coh-Metrix. Cambridge University Press. https://doi.org/10.1017/CBO9780511894664

Tan, C., Lee, L., & Pang, B. (2016). The effect of wording on message propagation: Topic- and author-controlled natural experiments on Twitter. Proceedings of ACL. https://doi.org/10.18653/v1/P16-1009