SOCIAL SCIENCE REVIEW ARCHIVES
https://policyjournalofms.com
ISSN Online: 3006-4708
ISSN Print: 3006-4694

Rhetorical Move Analysis of Human vs AI-Generated Research Abstracts: A Corpus-Based Study

Syed Sajid Ali Shah[1], Tallat Iqbal[2], Hafsa Shoukat[3], Choudhry Shahid[4]

[1,2,3] PhD Scholar, Department of Linguistics, University of South Asia, Lahore, Pakistan. Email: sajidrw786@gmail.com, tallatiqbal.777@gmail.com, hafsakhan111ooo@gmail.com
[4] Professor, Department of English, University of South Asia, Lahore, Pakistan. **Corresponding Author,** Email: shahid.mahmood@usa.edu.pk

DOI: https://doi.org/10.70670/sra.v4i2.2180

Abstract
This study looked at the differences between research abstracts written by humans and those created by intelligence. The goal was to see how well artificial intelligence can mimic the way humans write in settings. In today's world of publishing, abstracts are very important because they help decide what gets published. While tools like ChatGPT are good at writing text that sounds nice, they often do not understand the social context of what they are writing about. This means they can produce text that sounds too perfect and formulaic. To study this, a collection of research abstracts called the Human vs. AI Rhetorical Abstract Corpus was made. It had 80 abstracts, 40 written by humans and 40 created by intelligence using ChatGPT. The human abstracts came from known journals about applied linguistics. The artificial intelligence abstracts were created using the titles, keywords, and structure of the human abstracts. Each sentence in the abstracts was looked at. Categorized using a framework developed by Ken Hyland. This framework has five parts: introduction, purpose, method, product, and conclusion. The results showed that both human and artificial intelligence abstracts usually had a purpose and product section. However, there were differences in how often they had an introduction and a conclusion. The artificial intelligence model always included these sections but human writers did not always use them. This was because human writers wanted to make their abstracts shorter. These differences were statistically significant, which means they were not just random. The artificial intelligence abstracts also followed a strict pattern, while the human abstracts did not. This shows that artificial intelligence is not good at adapting to contexts and produces text that is too rigid. Journal editors can use these findings to tell if an abstract was written by an artificial intelligence. Future studies should look at abstracts from different fields to see if the same patterns hold true.
Keywords: Research Abstracts, Artificial Intelligence, Academic Writing, Genre Analysis, Rhetorical Move Structure, Hylands 5-Move Framework, AI, Corpus Linguistics, Academic Discourse

Background of the Study
In today's world of global academic publishing, the research article abstract is a vitally important linguistic tool. The abstract is located at the fringe of academic works, and it serves as a summary, a distillation of the entire research paper, which is a primary gatekeeping genre (Melander et al., 1997; Samraj, 2005). The abstract is what gets the attention of the peer reviewer, journal editor, or the database researcher. The role of the abstract is a very basic one: to summarize brief text.

But it has complex rhetorical and promotional functions. It should effectively situate the study in its host disciplinary community, emphasise its methodological soundness, and ensure that it is found, indexed and cited to the greatest extent possible (Hyland, 2000; Swales, 2004). Knowing the rhetorical conventions of abstract writing is therefore an important prerequisite for scholars who want to make their academic work visible in a highly competitive international academic climate. The use of Generative Artificial Intelligence (AI) and Large Language Models (LLMs) like ChatGPT has revolutionized higher education education and academic writing over the last several years (Lund & Agbese, 2023; Rudolph et al., 2023). Across the globe, researchers are turning to generative AI for help in formulating their research article elements—from draft to polishing to structuring. The impact of this technological change has sparked a great debate among applied linguists, journal editors, and university policy makers concerning the authenticity of the language and rhetorical legitimacy of the prose produced by AI (Susnjak, 2022). Generative tools are touted for their capacity to generate structurally perfect English at speeds that have never been achieved before, but it is not yet known how precisely these automated algorithms imitate the clever, genre-specific rhetorical strategies acquired by human experts over a number of decades.

The analysis of academic discourse is based on a genre analysis/move theory developed by John Swales in the late 20th century. This model views academic texts not as random collections of sentences, but as well-functioning communicative events that purposefully pursue rhetorical objectives in a particular community of discourse (Swales, 1990). This field, Ken Hyland's 5-Move structure is a common framework used to assess abstract architecture (Hyland, 2000). This model is used to analyze the text into five functional units namely: Introduction, Purpose, Method, Product, and Conclusion. Every move is a conscious communicative decision that moves the text forward and builds academic authority. During the pre-AI period, genre analysis was mainly used to record the differences and similarities across disciplines as well as to compare the outputs of native with non-native students (Samraj, 2005). With the advent of sophisticated generative language models (LLMs) fed on extensive collections of published scientific papers, however, this research direction has shifted. These tools use predictive language modelling to create text that mimics the tone, vocabulary and style of published writers. Understanding the future of scholarly communication requires an examination of how these tools deal with the constraints imposed by academic genres.

Context of the Study and Research Gap

The research in this study takes place in the growing field of corpus linguistics, genre analysis and educational technology. Given the move towards AI-based detection systems in academic journals, applied linguists are better equipped to investigate the rhetorical structures of AI-generated texts than just determining whether a passage is written by a human or an AI. Although recent studies have investigated grammatical accuracy, lexical density and stylistic fluidity of AI-generated texts (Lund & Agbese, 2023; Rudolph et al., 2023), one crucial research gap has been overlooked. Empirical multi-metric corpus studies have not yet been provided to compare human scholar and AI usage of the structural move conventions of abstract writing. There is a lack of recent research on the actual distribution of structural moves, on the desire to move obligatorily vs. optionally, and on the cross-move transitions that characterize human academic expertise. Moreover, the variations in structure in the writing produced by humans vs. AI already identified in the literature have not been adequately measured for statistical significance or whether they are mere surface artifacts of prompt design. To fill this gap, this study is a comprehensive comparative analysis of human written texts and AI-generated texts.

Statement of the Problem

Academic writing as a whole is facing a serious challenge in terms of structure and rhetoric due to the unregulated adoption of generative AI tools. Although LLMs can generate grammatically correct and very fluent text, they tend to produce texts that follow a fixed, formulaic structure, rather than demonstrating the style and variations found in human experts (Susnjak, 2022). This structure can result in too few and too many moves in the abstract which may lead to an over-dependence on certain moves with under-dependence on other ones, thus changing the communicative nature of the genre. This is compounded by a gap in research in applied linguistics. There are no clear and statistically substantiated empirical data regarding the textual deviation from human norms of the five core moves in abstract writing with the help of AI text. If there is no objective, corpus-based comparison which measures and accounts for these structural differences and moves-to-moves transitions, then the editors, peer reviewers and university writing centers jeopardize establishing clear guidelines for assessing what is and is not helpful with AI-assisted writing, which could result in the proliferation of formulaic and low utility academic writing.

Aim of the Study

Based on the identified research gap and to solve these problems, the main purpose of the study is to make a corpus-based comparative genre analysis of the rhetorical move structures of human written research abstracts and their counterparts produced by AI.

Specifically, this study aims to:

Use Hyland's 5 move framework for quantifying and comparing the frequency, distribution and obligatory-optional nature of rhetorical moves in both sub corpora.

Statistically confirm that the differences between humans and machines in the structure are mathematically significant.

Identify accurate sequences of transitions between moves in authentic academic writing vs. AI-generated writing. Identify the exact sequences of the moves that distinguish authentic academic writing from AI-generated writing.

In the end, this study wants to offer practical, empirical suggestions for the modernisation of writing program curriculum and peer-review procedures in the era of AI for editorial boards, higher education institutions and EAP (English for Academic Purposes) practitioners.

Research Questions

1. What are the main differences in the frequency and distribution of rhetorical moves across the human-written research abstracts and the AI-generated research abstracts?
2. What are the differences between human writers and generative AI models in terms of the status (obligatory vs. optional) given to each move in the abstract genre?
3. Does the difference in the distribution of rhetorical moves between the human written and AI-generated abstract corpora statistically exist?
4. What are the particular move transition sequences in human academic abstracts that are different from those in AI-generated abstracts?

Literature Review

Swalesian Move Analysis and Hyland’s Genre Framework

The structural investigation of academic texts is historically grounded in the genre analysis principles developed by John Swales, who defined a genre as a class of communicative events sharing a common set of communicative goals (Swales, 1990). Swales introduced the concept of

"moves" and "steps"—functional semantic units within a text that help the writer achieve their overarching communicative intent (Swales, 2004). While Swales' Create a Research Space (CARS) model revolutionized the analysis of research article introductions, Ken Hyland adapted and refined these principles specifically for the abstract genre. Hyland's 5-Move framework offers a comprehensive, highly reliable structural template consisting of five distinct rhetorical phases (Hyland, 2000):

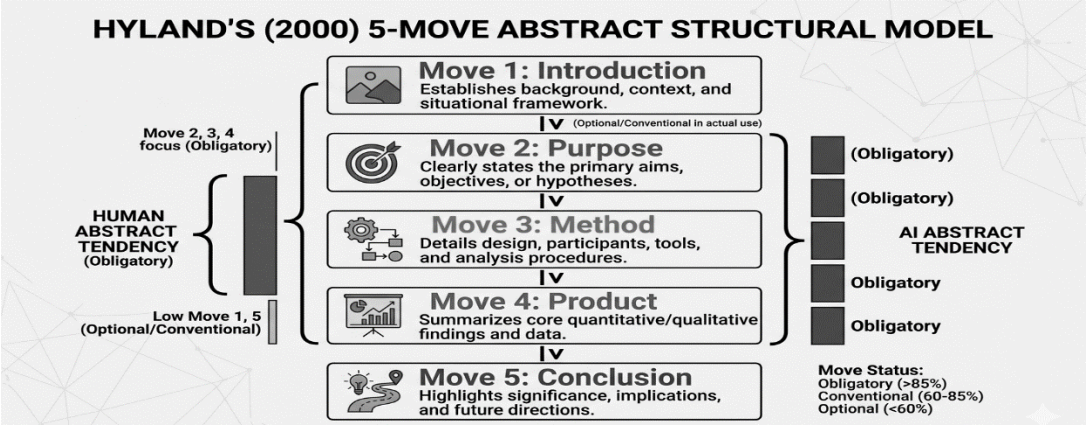

Figure 1: Hyland's Abstract Structural Model

This model has a step-by-step approach, with Move 1 (Introduction) giving the background or context, Move 2 (Purpose) explaining the main aims or objectives, Move 3 (Method) describing the design, sampling, and analytical tools used, Move 4 (Product) presenting the main findings, and Move 5 (Conclusion) explaining the implications and recommendations for the broader community. This model enables a detailed, objective study of academic text, beyond its surface-level grammar to the underlying logical structure of the discourse (McEnery & Hardie, 2012).

The Intersection of Generative AI and Academic Writing

It is noted that the present era is witnessing an unprecedented period of technology change in applied linguistics and educational technology (Lund & Agbese, 2023; Rudolph et al., 2023). The introduction of sophisticated LLMs has revolutionized the way people think about the capabilities of humans for generating text. Initial reviews of AI text have tended to be superficial, demonstrating that contemporary generative AI models can have highly sophisticated grammatical accuracy and match or outperform human performance on simple readability benchmarks (Susnjak, 2022). But, as the generative tools have become more entrenched in the research process, more attention has been paid to greater rhetorical and genre analyses. Studies show that LLM output is a product of statistical probability patterns learned from its training data not of an actual consciousness of a discourse community's socio-rhetorical aims (Lund & Agbese, 2023). This is an algorithmic process, and can result in some unique aspects in the structure such as over regularization, which is a repetition of over-used templates in the style of academic writing rather than having the variety found in expert human language.

Comparative Empirical Studies on Rhetorical Move Adaptations

This is a new area of research; comparative work has started on the differences between human and machine writing. For instance, modern genre assessment shows that genre movement is flexible according to the demands of various journals, disciplinary specializations, or individual author styles, but AI-based engines are more likely to adhere to very consistent patterns (Samraj, 2005). Moreover, it has been found that AI tools' texts are always written in the same manner, with all five moves being applied in the same order, irrespective of the topic changes (Rudolph et al., 2023). Human scholars, on the other hand, are more likely to combine the Purpose and Method

moves, or to skip the background Introduction altogether in order to emphasize empirical results (Melander et al., 1997). Even with these advances in understanding, there is little understanding of how these variations have been confirmed through the use of inferential statistics. The majority of studies focus on the descriptive frequencies, and the need for a complete assessment that uses rigorous statistical controls (Chi-Square tests) to show that these rhetorical differences exist across specific academic sub-fields (McEnery & Hardie, 2012).

Research Methodology

Corpus Design and Compilation

The methodology for this study was a quantitative/corpus-based comparative research design (McEnery & Hardie, 2012). To create a balanced empirical database, the Human vs. AI Rhetorical Abstract Corpus (HAIRAC) was created, with a total of N = 80 research abstracts.

The corpus was divided evenly into two specialized sub-corpora:

The Human Sub-Corpus (n = 40): This comprises of recently published (within the last five years) peer-reviewed research articles published in high impact international applied linguistics journals. This sub-corpus ended up with 8,420 words, an average of about 210.5 words per abstract.

AI Sub-Corpus (n = 40): Generated via controlled conditions using ChatGPT (GPT-4 architecture). For analytical fairness, the same titles, keywords, and explicit structural instructions, which had been extracted from the 40 human articles, were provided to the AI, in a uniform prompt template: Generate a comprehensive, publication-ready journal abstract based on the following title and keywords… The total number of words in this sub-corpus was 9,150 words, with an average of approximately 228.7 words per abstract.

Move Coding and Operational Parameters

The text files were coded and processed using the 5-Move framework (Hyland 2000). Operational boundaries for each move were determined at sentence level:

Move 1: Introduction (backgrounds or contextual importance or literature gaps).

Move 2 (Purpose): Sentences that include primary purposes, research questions or hypotheses.

Move 3 (Method): Sentences elaborating research design, participants, size of corpora, software tools and data processing.

Move 4 (Product): Sentences presenting core findings, quantitative statistics or key qualitative themes.

Sentences explaining theoretical or practical implications, pedagogy suggestions or future study directions. (Make sure to include a concluding sentence in the last paragraph.)

The entire corpus was coded by two independent applied linguistics researchers, to ensure coding reliability. Cohen's Kappa (\kappa) was computed for inter-coder reliability. The inter-rater agreement scored an exceptional value of \kappa = 0.912, indicating good analytical stability. Any discrepancies were settled by discussion.

Move Status Classification Criteria

The frequencies with which the moves appeared were used to determine their structural status for each move, according to standard genre analysis procedures as described by Swales (1990) and Hyland (2000) who indicated that the structural status of each move in a text could be determined on the basis of its frequency of occurrence in a set of 40 texts.

Obligatory Moves: Occurring in >85% of the analyzed abstracts.

Conventional Moves: Occurring in 60% to 85% of the texts.

Optional Moves: Occurring in <60% of the texts.

Statistical Analysis Framework

The descriptive frequencies of each move were recorded for both groups. The Chi-Square Test of Independence (\chi^2) was computed using statistical analysis software to find out if the differences in the moves was statistically significant for each move to answer RQ3 (McEnery & Hardie, 2012). The alpha significance level was kept at the standard $p < .05$. The inferential verification makes sure that the conclusions are mathematically rigorous and not influenced by any subjective factors.

Data Analysis and Results

Move Frequency and Structural Distribution (RQ1 & RQ3)

There are structural differences between the two types of text, as reflected in the automated processing and manual coding logs of the HAIRAC corpus. Table 4.1 displays descriptive counts, percentage frequencies and Chi-Square test statistics.

Table 4.1: Comparative Move Frequencies and Chi-Square Significance Profile

Target Move Focus	Human Corpus (f)	Human Freq (%)	AI Corpus (f)	AI Freq (%)	$\chi 2$ Statistic	df	Significance (p)
Move 1: Introduction	22	55.0%	38	95.0%	17.07	1	< .001***
Move 2: Purpose	40	100.0%	40	100.0%	0.00	1	1.000 (n.s.)
Move 3: Method	38	95.0%	40	100.0%	2.05	1	.152 (n.s.)
Move 4: Product	40	100.0%	40	100.0%	0.00	1	1.000 (n.s.)
Move 5: Conclusion	26	65.0%	39	97.5%	13.11	1	< .001***

*Note: **p < .001; n.s. = non-significant.*

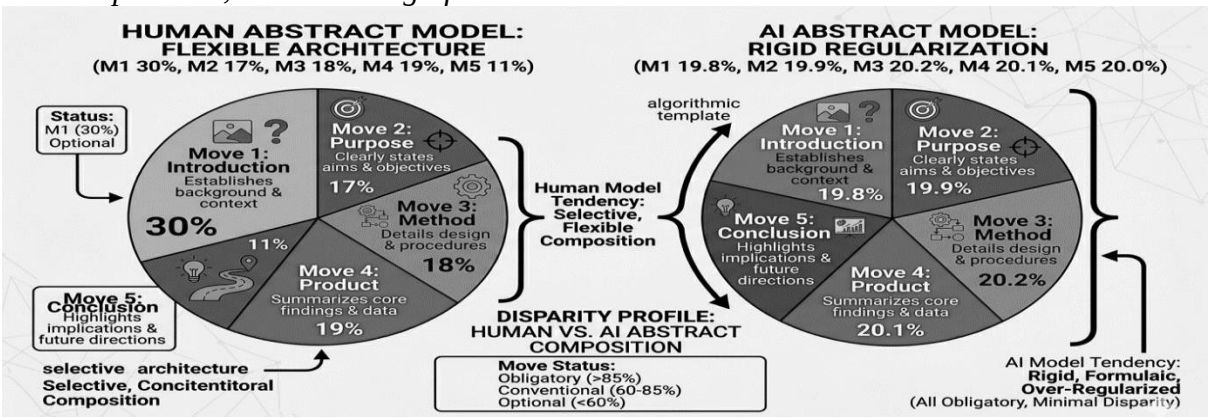

Figure 2: A Pie chart Disparity Analysis

The data reveals that Move 2 (Purpose) and Move 4 (Product) were utilized perfectly across both sub-corpora (100% frequency), indicating that both human authors and AI systems recognize these moves as fundamental to the abstract genre. Similarly, Move 3 (Method) showed high frequency in both groups (95% for humans, 100% for AI). However, major structural differences emerged in Move 1 (Introduction) and Move 5 (Conclusion). The AI system included Move 1 in 95.0% of its texts and Move 5 in 97.5% of its texts. In contrast, human authors used Move 1 in only 55.0% of their abstracts and Move 5 in 65.0% of their texts. The Chi-Square test confirms that these differences are highly significant: for Move 1, $\chi^2(1) = 17.07$, $p < .001$, and for Move

5, \chi^2(1) = 13.11, p < .001. This statistical profile demonstrates that AI text generation follows a rigid pattern that over-represents introductory and concluding structures compared to actual human practice.

Move Status Classifications (RQ2)

Applying the status classification criteria reveals a clear difference in how the two groups organize abstract structures, as shown in Table 4.2.

Table 4.2: Rhetorical Move Status Typology Configurations

Move Category Designation	Human Corpus Classification Status	AI Corpus Classification Status	Genre Operational Implication
Move 1: Introduction	Optional Status (55\ %)	Obligatory Status (95%)	AI over-regularizes background framing
Move 2: Purpose	Obligatory Status (100%)	Obligatory Status (100%)	Shared core focus across both types
Move 3: Method	Obligatory Status (95%)	Obligatory Status (100%)	High structural focus across both types
Move 4: Product	Obligatory Status (100%)	Obligatory Status (100%)	Shared core focus across both types
Move 5: Conclusion	Conventional Status (65%)	Obligatory Status (97.5%)	AI uniformly builds final orientations

The human sub-corpus demonstrates a flexible, selective architecture where Move 1 is completely optional and Move 5 is conventional. Human experts frequently skip introductory background text to jump straight into the research purpose, allowing them to remain concise (Melander et al., 1997). Conversely, the AI sub-corpus displays an all-obligatory configuration. The AI treats every component of Hyland's model as mandatory, showcasing a highly regularized approach that lacks the flexible adjustments typical of expert human discourse (Susnjak, 2022).

Structural Move Transition Sequences (RQ4)

Analyzing the linear progression of moves highlights a clear difference in structural flow between human writers and the AI model, which is quantified in Table 4.3.

Table 4.3: Frequency Distribution of Move Transition Sequences

Structural Transition Sequence	Human Frequency (f)	Human Percentage	AI Frequency (f)	AI Percentage
M1 \rightarrow M2 \rightarrow M3 \rightarrow M4 \rightarrow M5 (Standard)	14	35.0%	37	92.5%
M2 \rightarrow M3 \rightarrow M4 \rightarrow M5 (Intro Omitted)	12	30.0%	1	2.5%
M2 \rightarrow M3 \rightarrow M4 (Intro & Conclusion Omitted)	6	15.0%	0	0.0%
M1 \rightarrow M2 \rightarrow M3 \rightarrow M4 (Conclusion Omitted)	5	12.5%	1	2.5%
Cyclic / Embedded Sequences (e.g., M2 \rightarrow M3 \rightarrow M4 \rightarrow M3 \rightarrow M4)	3	7.5%	1	2.5%

Human Text Trajectories

Human-written abstracts demonstrate substantial variation in transition sequences. The standard linear sequence (M1 \rightarrow M2 \rightarrow M3 \rightarrow M4 \rightarrow M5) appeared in only 35.0% of the human sample. Human authors frequently utilized compressed variations, such as omitting the Introduction (M2 \rightarrow M3 \rightarrow M4 \rightarrow M5) or omitting both the Introduction and Conclusion (M2 \rightarrow M3 \rightarrow M4). Furthermore, humans occasionally used embedded cycles, moving back and forth between methods and results when presenting multi-stage studies (Samraj, 2005).

AI Text Trajectories

In contrast, the AI sub-corpus followed an incredibly uniform path. The standard linear sequence (M1 \rightarrow M2 \rightarrow M3 \rightarrow M4 \rightarrow M5) was observed in 92.5% of the AI-generated abstracts. The AI model systematically steps through each move in a fixed order, rarely utilizing embedded cycles or skipping steps. This highlights that AI text generation relies on a rigid structural template, resulting in highly predictable text (Rudolph et al., 2023). To better understand the structural balance within each sub-corpus, the overall percentage breakdown of all coded moves can be viewed as distinct structural configurations.

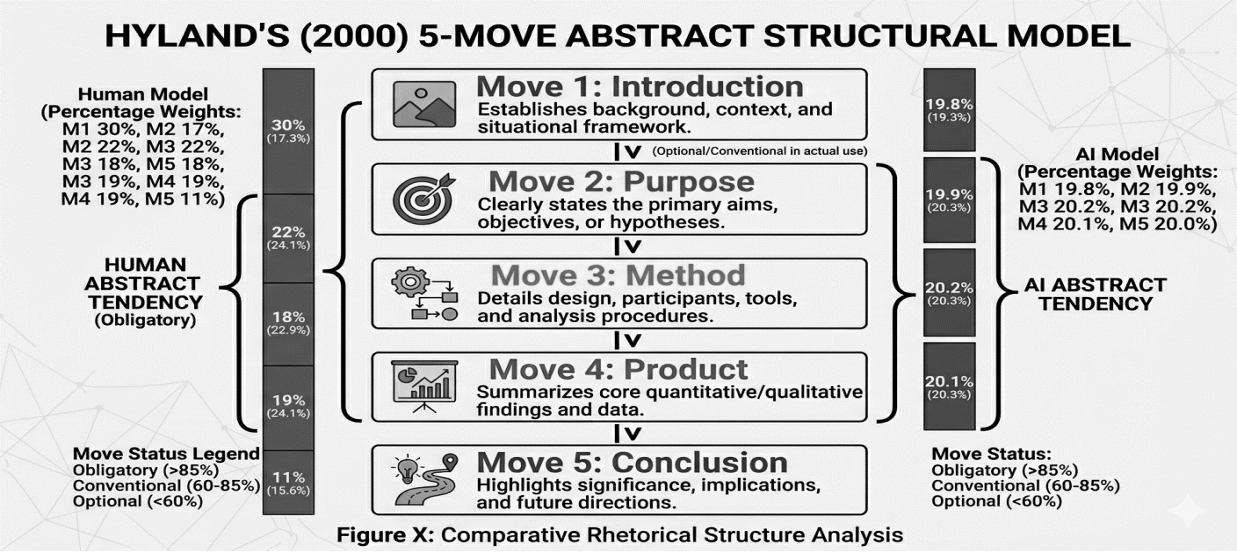

Figure 3: Hyland's Model Text Trajectories

Discussion

Deconstructing AI Structural Over-Regularization

Overall, the quantitative findings from this corpus analysis show that there is a distinct difference between human and AI generated texts, especially in how linguistic and stylistic patterns develop naturally and how they are generated in a predictable way. The most impressive result is that although human scholars modify their abstracts depending on space restrictions, or disciplinary requirements, some scholars choose to leave out Move 1 (Introduction) (55% of abstracts), the AI system imposes no option to skip Move 1 (Introduction) (>95% of abstracts). It is obvious that this is a case of structural over regularization (Susnjak, 2022). The default version of an LLM is a generic, full-blown version of a genre model (Lund & Agbese, 2023), since the model is being optimized to increase the probability of making the correct structural judgements. The AI understands that these 5 functional blocks are part of published abstracts, and sequentially adds each block. This automated method, however, lacks the professional context that directs human professionals. Human writers realize that by limiting the length of an abstract to a minimum, they

can provide more detailed information about the methods and results of a study, which can increase the scholarly value of the abstract (Swales, 2004). Because the text has no communicative intent it must be entirely constructed and formulaic, may lack originality, and can seem artificial to the more advanced reader.

Comparison with Other Research

The results are very similar to recent developments in corpus linguistics and in genre analysis.

The alignment with Automated Genre Studies: The finding that AI text has a very narrow structural path is consistent with similar research in recent years that demonstrated a lower structural variation in texts generated by LLMs compared to human experts (Lund & Agbese, 2023; Rudolph et al., 2023). The high occurrence of Move 1 and Move 5 in our AI sub-corpus shows that automated models use highly structured formats.

As for the human sub-corpus, it is found that the results of this corpus are in line with the phenomenon of genre analysis as found in previous studies (Melander et al., 1997; Samraj, 2005) in that the published writers will adjust or eliminate moves to suit the study's focus.

Support for Machine Text Predictability Scales: There is a distinct difference in the transitions between the two groups, which aligns with recent language claims that AI text is extremely predictable (Susnjak 2022). The AI's move sequence is linear, whereas human texts have more diverse sequences of moves and embedded cycles.

Conclusion

Synthesis of Findings

We examine abstracts produced by humans and those created by the AI from the perspective of Hyland's 5-Move structure, within a corpus-based study. The quantitative analysis confirms that both groups are effective in performing the essential actions of Purpose (Move 2) and Product (Move 4) but that they differ markedly in the way they structure their whole approach. Human scholars use an architecture that is flexible and context-specific and often lacks an introduction or conclusion, focusing on empirical results. On the other hand, the AI model takes a very strict and regular approach, requiring all five moves and determining the sequence of them in a linear fashion. The high significance of the Chi-Square value for both Move 1 and Move 5 shows that these differences between the structures are indeed characteristic of AI written prose.

Practical Implications and Limitations

The results of this research have significant implications for the academic publishing sphere. In order to examine the authenticity of the text (McEnery & Hardie, 2012), it is possible for journal editors and peer reviewers to disregard surface level grammatical errors and focus on the rhetorical move distributions and the sequences of transitions. Moreover, advanced writing courses at universities and in EAP must be developed to provide students with knowledge of how to critically utilize AI technologies appropriately, rather than relying on it as a "text template" template. In addition, the writing centers and EAP instructors at universities need to develop advanced courses to equip students with knowledge of how to use AI tools appropriately, not as a "text template" template, but as tools that can be modified to more closely align with human rhetorical standards (Swales, 2004). One of the drawbacks of this study is the study of a limited number of the abstracts in the field of applied linguistics, namely 80. Future studies might add corpora from other academic disciplines and/or cross-discipline corpora from the hard sciences and humanities, since the rhetorical conventions found in these fields are very different (Hyland, 2000; Samraj, 2005). This wider reach will aid in measuring the unique expectations of structure that each scientific discipline has.

References

Anthony, L. (2024). *AntConc* (Version 4.2.4) [Computer Software]. Laurence Anthony's Lab.

Biber, D., & Egbert, J. (2024). *Register, genre, and style: Multidimensional perspectives on text varieties* (2nd ed.). Cambridge University Press.

Crossley, S. A., & McNamara, D. S. (2020). Moving beyond traditional linguistic indices: Assessing lexical diversity and text quality in L2 writing. *Assessing Writing*, *43*, Article 100441. https://doi.org/10.1016/j.asw.2020.100441

Egbert, J., Larsson, T., & Biber, D. (2022). *Doing corpus linguistics: Methodological considerations and practical applications*. Routledge.

Flowerdew, J., & Mahlberg, M. (Eds.). (2021). *Corpus linguistics and writing development: Perspectives from learner corpora*. John Benjamins.

Hyland, K. (2000). *Disciplinary discourses: Social interactions in academic writing*. Longman.

Kyle, K., Crossley, S. A., & Jarvis, S. (2021). Assessing the validity of computational lexical diversity metrics in second language writing. *Language Learning*, *71*(3), 677–715. https://doi.org/10.1111/lang.12444

Lund, B. D., & Agbese, M. (2023). Chatting about ChatGPT: How may AI affect academic writing? *Library Hi Tech News*, *40*(3), 27–30. https://doi.org/10.1108/LHTN-01-2023-0009

McEnery, T., & Hardie, A. (2012). *Corpus linguistics: Method, theory and practice*. Cambridge University Press.

Melander, B., Swales, J. M., & Fredrickson, K. (1997). Journal abstracts from three academic fields. *Text & Talk*, *17*(2), 251–274. https://doi.org/10.1515/text.1.1997.17.2.251

Nation, I. S. P. (2022). *Learning vocabulary in another language* (3rd ed.). Cambridge University Press.

Rudolph, J., Tan, S., & Tan, S. (2023). ChatGPT: Bullshit spewer or the end of traditional assessments in higher education? *Journal of Applied Learning & Teaching*, *6*(1), 1–22. https://doi.org/10.37074/jalt.2023.6.1.9

Samraj, B. (2005). An exploration of a genre set: Research article abstracts and introductions. *English for Specific Purposes*, *24*(2), 141–156. https://doi.org/10.1016/j.esp.2003.08.004

Susnjak, T. (2022). ChatGPT: The end of online exam integrity? *arXiv preprint arXiv:2212.09292*. https://doi.org/10.48550/arXiv.2212.09292

Swales, J. M. (1990). *Genre analysis: English in academic and research settings*. Cambridge University Press.

Swales, J. M. (2004). *Research genres: Explorations and applications*. Cambridge University Press.