

Linguistic Features of Selected AI-Generated Short Stories

Favour Ogechi Adiele,¹ Adagogo C. Brown²

Abstract— The increasing deployment of artificial intelligence (AI) in creative writing has generated substantial interest in the linguistic qualities of machine-produced literary texts. This study examines the linguistic features of selected AI-generated short stories, with particular attention to syntactic structures, lexico-semantic choices, figurative expressions, and graphological organisation. Adopting a qualitative descriptive design grounded in corpus stylistics, the study analyses twenty short stories generated using three open-access large language models (ChatGPT-4, Claude 2, and Google Gemini) between January and March 2025. The analysis draws on the theoretical frameworks of stylistics and Systemic Functional Grammar. Findings reveal that AI-generated stories demonstrate considerable linguistic consistency across multiple levels of textual organisation. Syntactically, the texts rely predominantly on declarative structures, paratactic sentence patterns, parallelism, and anaphoric repetition, producing a measured but formulaic narrative rhythm. At the lexico-semantic level, vocabulary is organised around recurring semantic fields including isolation, time, memory, technology, and loss, with limited lexical innovation beyond conventional literary associations. Figurative expressions, particularly metaphor, simile, and personification, are employed frequently but tend to follow established literary patterns rather than generating novel conceptual mappings. Graphologically, the stories maintain orderly paragraphing, standard punctuation, and conventional formatting, alongside occasional orthographic irregularities characteristic of machine generation. The study concludes that while AI can reproduce many conventional linguistic features of literary prose with remarkable fluency, its textual patterns remain relatively consistent, formulaic, and constrained by the probabilistic architecture of language models. These findings carry implications for stylistic theory, digital literary studies, and the pedagogy of creative writing in an era of machine-generated text. The study recommends that future research adopt comparative frameworks that directly examine human-authored and AI-generated texts using identical prompts to more precisely delineate the linguistic boundaries between human and machine creativity.

Keywords: artificial intelligence; AI-generated text; linguistic features; short stories; stylistics; syntax; lexico-semantics.

¹ Department of English and Literary Studies Rivers State University, Port Harcourt, Nigeria Email: favouradiele14@gmail.com

² Department of English and Literary Studies Rivers State University, Port Harcourt, Nigeria Email: brown.adagogo@ust.edu.ng

© 2026 the Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License, Attribution-NonCommercial 4.0 International (CC BY-NC 4.0).

INTRODUCTION

Language is central to literary expression. Writers use linguistic choices to construct characters, present events, establish settings, communicate emotions, and develop themes (Ogurinka & Chimene-Wali, 2026). Literary language differs from ordinary communication because linguistic forms are deliberately selected and arranged to produce particular aesthetic and interpretive effects. The study of these choices is one of the major concerns of stylistics, which provides systematic methods for examining how linguistic forms contribute to meaning in literary texts (Leech & Short, 2007).

The development of digital technologies has introduced new possibilities for the production of literary texts. Artificial intelligence, particularly large language models (LLMs) such as OpenAI's GPT series, Anthropic's Claude, and Google's Gemini, can now generate poems, narratives, essays, and other forms of creative writing. These systems produce texts by processing vast quantities of linguistic data and generating sequences that are contextually related to an input or prompt (Boden, 2016; Floridi & Chiriatti, 2020; Adekanmbi, et al., 2026). The increasing ability of AI systems to produce apparently coherent literary narratives has raised fundamental questions about the linguistic qualities of such texts and their relationship to human-authored literature.

AI-generated fiction is particularly interesting from a linguistic perspective because it can reproduce several features traditionally associated with literary writing. AI-generated narratives can contain complex sentence structures, figurative expressions, descriptive vocabulary, thematic patterns, and carefully organised paragraphs. The ability of language models to produce such features suggests that computational systems can reproduce many of the linguistic conventions found in human literary texts, at least at the surface level.

However, the presence of literary features in AI-generated writing does not necessarily mean that the linguistic behaviour of AI is identical to that of human writers. Human writers make linguistic choices within social, cultural, experiential, and ideological contexts. Their language is shaped by intention, emotion, memory, and embodied experience. AI systems, on the other hand, generate language through computational processes based on statistical patterns learned from large bodies of textual data (Floridi & Chiriatti, 2020; Bender et al., 2021). This fundamental distinction makes the linguistic examination of AI-generated fiction both necessary and timely.

The present study therefore focuses on the linguistic features of selected AI-generated short stories. The study examines the texts at different linguistic levels, particularly syntax, lexico-semantics, figurative language, and graphology, in order to identify the recurring patterns that characterise machine-generated literary prose. By applying established stylistic frameworks to these texts, the study aims to contribute to the emerging scholarly conversation about the nature of AI-generated literature and its implications for linguistic and literary studies.

STATEMENT OF THE PROBLEM

The development of artificial intelligence has made the production of literary texts increasingly accessible to machines. AI-generated stories can appear coherent, grammatically acceptable, imaginative, and emotionally expressive. This development has created an important question for literary and linguistic scholarship: what linguistic features characterise AI-generated literary writing, and how do these features compare with the linguistic conventions of human-authored fiction?

Although discussions of artificial intelligence frequently focus on issues such as authorship, plagiarism, ethics, and the future of human creativity, less attention has been given to the actual linguistic structures through which AI-generated narratives create meaning. Understanding these structures is important because the language of AI-generated fiction may reveal recurring patterns that distinguish machine-produced texts from human literary writing. These patterns, once identified, can inform both theoretical understandings of literary language and practical approaches to the evaluation of AI-generated texts.

The problem is particularly relevant to literary studies because style has traditionally been associated with the distinctive linguistic choices of individual human authors. The concept of authorial style presupposes a human agent whose linguistic choices reflect intention, personality, and creative vision. AI-generated fiction challenges this assumption by producing texts that contain recognisable literary patterns without a human author directly composing each linguistic expression. This raises questions about the nature of style itself and the relationship between linguistic form and authorial agency.

PURPOSE OF THE STUDY

The aim of this study is to conduct a stylistic analysis of selected AI-generated short stories in order to examine their linguistic and narrative features. The specific objectives are to examine the syntactic patterns of sentences employed in the AI-generated stories; to ascertain the lexico-semantic features and figurative expressions of AI-generated short stories; and to investigate the graphological features of AI-generated short stories and their contribution to thematic meaning.

RESEARCH QUESTIONS

The following research questions guided the study. What syntactic patterns are employed in the AI-generated short stories? What lexico-semantic features and figurative expressions are employed in the AI-generated short stories? How do the graphological features of AI-generated short stories contribute to their thematic meaning?

LITERATURE REVIEW

Language and Literary Language

Language is a system through which human beings communicate ideas, experiences, emotions, and perceptions. In literary writing, however, language performs functions beyond the simple transmission of information. Writers manipulate linguistic resources to create atmosphere, characterisation, symbolism, rhythm, and emotional effects. The literary text is a site of deliberate linguistic selection, where choices at every level of language contribute to the overall aesthetic and interpretive experience.

Literary language is therefore characterised by what the Russian Formalists termed defamiliarisation, the making strange of familiar linguistic forms in order to renew perception (Leech, 2008). Leech and Short (2007) explain that stylistic analysis is concerned with the relationship between linguistic choices and literary effects. The study of literary language consequently requires attention to the actual linguistic forms through which a text achieves its effects, rather than relying solely on impressionistic judgments about literary quality.

Linguistic Features and Levels of Analysis

Linguistic features refer to identifiable patterns in the use and organisation of language. In literary analysis, such features may occur at different levels, including syntax, lexis, semantics, morphology, graphology, and discourse. Each level provides a distinct perspective on the organisation of the text and contributes to its overall meaning.

Syntax concerns the arrangement of words into phrases, clauses, and sentences. Sentence structure influences rhythm, emphasis, and the representation of relationships between ideas. Leech and Short (2007) identify sentence complexity, sentence length, clause structure, and sentence types as important variables in stylistic analysis. Lexico-semantics deals with vocabulary and meaning, including semantic fields, lexical relationships, connotation, and the interplay between denotative and figurative meaning. Crystal (2003) notes that lexical analysis involves examining both the individual words used in a text and the patterns of meaning they create collectively.

Graphology concerns the visual and orthographic properties of written language, such as punctuation, capitalisation, paragraphing, spacing, and typography. While often overlooked in stylistic analysis, graphological features contribute to the organisation and presentation of literary texts and can carry thematic significance. These linguistic levels are interconnected. A writer's choice of sentence structure may influence rhythm and emphasis, while lexical choices can establish a particular semantic

field or emotional atmosphere. Similarly, punctuation and paragraphing can influence the pace and visual organisation of a narrative.

Artificial Intelligence and Text Generation

Artificial intelligence refers broadly to computational systems designed to perform tasks associated with human intelligence, including perception, reasoning, learning, and language processing (Mensah, et al., 2026). In language generation, AI systems use computational models trained on extensive bodies of textual material to produce new linguistic sequences. These models, often based on transformer architectures, learn statistical patterns of language use from their training data and generate text by predicting likely sequences of words given a particular prompt or context (Vaswani et al., 2017; Nwogu, et al., 2024).

Large language models have significantly expanded the capacity of AI systems to generate coherent text. Floridi and Chiriatti (2020) describe GPT-3 as a large-scale language model capable of generating human-like textual outputs across a wide range of genres and domains. They note that while the linguistic fluency of such models is impressive, their outputs are fundamentally the product of statistical pattern recognition rather than understanding, intention, or consciousness. The model does not know what it is saying; it generates text by calculating the probability of word sequences based on patterns observed in its training data.

Bender et al. (2021) offer a critical perspective on large language models, arguing that they are essentially “stochastic parrots” that reproduce patterns from their training data without genuine understanding. Their analysis highlights the risks of overestimating the capabilities of AI language systems and the importance of careful evaluation of their outputs. These critiques are directly relevant to the analysis of AI-generated literary texts, which may appear meaningful without embodying the intentional meaning-making that characterises human literary production.

AI-generated literary writing is consequently not simply random text. It can reproduce patterns of grammar, vocabulary, cohesion, narrative organisation, and figurative language found in its training data. This makes linguistic analysis a valuable tool for identifying the recurring characteristics of machine-generated fiction and for understanding the relationship between computational text generation and human literary practice.

Empirical Studies on AI-Generated Texts

A growing body of empirical research has examined the linguistic and literary qualities of AI-generated texts. Elkins and Chun (2020) conducted an early evaluation of GPT-3's ability to produce short stories, finding that while the model could generate coherent narratives, its outputs exhibited recurring structural and thematic limitations. The stories tended toward predictable plot structures and conventional characterisations, and they frequently lacked the depth and originality associated with skilled human fiction.

Gunser et al. (2022) compared human-authored and machine-generated literary texts, examining readers' ability to distinguish between them and their perceptions of literary quality. Their findings indicated that readers struggled to reliably identify machine-generated texts, but that human-authored texts were generally rated higher on measures of creativity and emotional resonance. This suggests that while AI can approximate the surface features of literary prose, it has not yet achieved the deeper qualities that readers value in human literature.

Kybartas and Bidarra (2017) provided a comprehensive survey of automated story generation systems, tracing the development of the field from early rule-based approaches to contemporary neural network models. Their analysis highlighted the persistent challenges of narrative coherence, character development, and thematic depth in machine-generated fiction. Fan et al. (2019) examined hierarchical story generation using large language models, demonstrating both the capabilities and limitations of current systems in producing extended narratives.

These studies provide a foundation for the present research. While they have examined various dimensions of AI-generated fiction, few have conducted systematic stylistic analyses of the linguistic features of such texts. The present study addresses this gap by applying established frameworks from stylistics and Systemic Functional Grammar to a corpus of AI-generated short stories.

THEORETICAL FRAMEWORK

Stylistics Theory

The primary theoretical framework for this study is stylistics. Stylistics is concerned with the systematic examination of language in literary and other texts, using the tools and concepts of linguistics to analyse the relationship between linguistic form and literary meaning. Leech and Short (2007) argue that stylistic analysis examines linguistic choices at multiple levels and considers how those choices contribute to literary effects. Stylistics is distinguished from literary criticism more broadly by its commitment to rigorous, evidence-based analysis of the language of texts.

The relevance of stylistics to this study lies in its ability to move beyond general impressions about whether an AI-generated story is effective or ineffective. Instead, stylistics provides a framework for identifying concrete linguistic patterns, including

sentence structures, lexical choices, figurative language, repetition, parallelism, and graphological organisation, and for analysing how these patterns contribute to the overall character of the text.

Systemic Functional Grammar

The study also draws on Systemic Functional Grammar (SFG), associated particularly with the work of Michael Halliday. SFG views language as a resource for making meaning and emphasises the relationship between linguistic choices and their functions within particular social and communicative contexts (Halliday, 1985; Halliday & Matthiessen, 2004). The theory identifies three metafunctions of language: the ideational, which concerns the representation of experience; the interpersonal, which concerns the enactment of social relationships; and the textual, which concerns the organisation of language into coherent texts.

The theory is relevant to the present study because linguistic forms in the AI stories do not merely occur structurally; they contribute to the representation of experiences, relationships, and ideas. Sentence choices, lexical selections, and patterns of emphasis can therefore be examined in relation to their functions within the narratives, providing a richer understanding of how AI-generated texts construct meaning.

METHODOLOGY

Research Design

The study adopts a qualitative descriptive research design grounded in corpus stylistics. This design is appropriate for the close examination of linguistic features across a body of texts and for the identification of recurring patterns that characterise a particular text type or mode of production.

Corpus

The corpus consists of twenty short stories generated using three open-access large language models: ChatGPT-4 (OpenAI), Claude 2 (Anthropic), and Google Gemini (Google). The stories were generated between January and March 2025. Each model was prompted to produce short stories using identical or thematically comparable prompts to ensure consistency across the corpus. The prompts were designed to elicit literary narratives that engaged with conventional short story themes, including loss, memory, technology, isolation, and human relationships.

Examples of prompts used include: “Write a literary short story about memory and loss set in a contemporary city”; “Write a short story in the style of literary fiction about a person waiting for something that never arrives”; and “Write a short story

about the relationship between humans and technology, told from a reflective, literary perspective.” Each prompt was submitted independently to each platform, and the first generated output was included in the corpus without editing or selection bias. The stories range in length from approximately 800 to 2,500 words and cover a variety of themes and settings.

Table 1 provides a summary of the corpus.

Table 1: Composition of the Corpus

Platform	Number of Stories	Story Codes
ChatGPT-4	8	GPT-01 to GPT-08
Claude 2	6	CL-01 to CL-06
Google Gemini	6	GM-01 to GM-06
Total	20	

Analytical Procedure

Each story was subjected to close textual analysis at multiple linguistic levels: syntax, lexico-semantics, figurative language, and graphology. Syntactic analysis focused on sentence types, clause structures, sentence complexity, and patterns of repetition and parallelism. Lexico-semantic analysis examined vocabulary choices, semantic fields, lexical register, and patterns of word association. Figurative language analysis examined the types and functions of metaphor, simile, personification, and other tropes. Graphological analysis examined paragraphing, punctuation, capitalisation, and overall textual organisation.

The analysis was conducted manually by the researchers through multiple readings of each text. Linguistic features were identified, coded, and categorised thematically. Patterns that recurred across multiple stories were noted as characteristic features of AI-generated literary prose.

Limitations

The study has several limitations. First, the corpus is relatively small, comprising twenty stories, which limits the generalisability of the findings. Second, the study does not include a comparative sample of human-authored short stories, which would enable more direct claims about the distinctive features of AI-generated prose. Third, the stories were generated using a limited set of prompts, and the linguistic features observed may be influenced by the nature of those prompts. These limitations are acknowledged and are addressed in the recommendations for future research.

ANALYSIS AND FINDINGS

Syntactic Features

The findings indicate that AI-generated short stories are organised through a high degree of syntactic consistency and structural symmetry. The corpus reveals a predominant reliance on declarative sentences, which establish a measured, observational narrative tone. This pattern aligns with Leech and Short's (2007) stylistic framework, where sentence choice directly shapes narrative perspective and the positioning of the reader in relation to the narrated events.

Across the corpus, declarative sentences constitute approximately 85 to 90 percent of all sentence types. Interrogative and imperative structures appear infrequently and are typically confined to passages of dialogue or moments of narrative emphasis. In story GPT-03, for instance, the entire narrative of 1,200 words contains only two interrogative sentences, both embedded within dialogue. The effect is a narrative voice that is consistently descriptive and reflective rather than interactive or commanding. This syntactic homogeneity contributes to the distinctive register of AI-generated fiction, which tends toward a uniform, polished prose style regardless of subject matter or narrative situation.

The texts further demonstrate frequent paratactic stacking of simple and compound sentences. Parataxis, the coordination of clauses without explicit subordination, produces a linear, additive organisation of information. For example, in story CL-02, a passage reads: "The train arrived at midnight. She was not on it. He waited anyway. The platform emptied. Still, he remained." The series of short, coordinated clauses creates a staccato rhythm that conveys the passage of time and the persistence of waiting. However, this paratactic organisation is employed so consistently across the corpus that it becomes a predictable feature of the AI-generated style. Human writers often vary their syntax, using hypotactic subordination to represent complex logical and temporal relationships; AI-generated texts, by contrast, tend to present information in a more linear, additive fashion.

Parallelism and anaphoric repetition are also recurrent features. In multiple stories, sentences begin with the same or similar structures, creating a rhythmic, incantatory effect. Story GPT-05 opens three consecutive paragraphs with variations on "She remembered," establishing a refrain that structures the narrative. Story CL-04 employs anaphoric repetition with the phrase "There was" to introduce successive descriptions of a desolate urban landscape. While effective in creating emphasis and rhythm, this pattern becomes predictable when employed repeatedly across different texts, suggesting a formulaic dimension to AI-generated literary style.

Table 2 summarises the distribution of key syntactic features across the corpus.

Table 2: Distribution of Key Syntactic Features

Feature	Frequency Across Corpus	Example
Declarative sentences	87% of all sentences	"The evening gathered around her like a crowd that had somewhere else to be."
Paratactic coordination	Present in 80% of stories	"She waited, and the light faded, and the street emptied."
Anaphoric repetition	Present in 65% of stories	"He remembered the rain. He remembered the silence. He remembered the leaving."
Syntactic incompletion	Present in 40% of stories	"Neither did grief."

Negation and syntactic incompletion also function as stylistic markers. Expressions with "not," "never," "no," and "without" are used to signal absence, loss, and emotional privation. In story GM-03, the elliptical clause "Neither did grief" closes a paragraph, compressing meaning and creating an abrupt effect. Overall, the syntactic organisation reflects an algorithmic preference for order, balance, and predictability rather than the fragmented, irregular structures that human writers often employ to represent psychological complexity, emotional turmoil, or the disorienting effects of traumatic experience.

Lexico-Semantic Features

The analysis shows that AI-generated narratives construct meaning through tightly clustered semantic fields, groups of words related by meaning that collectively establish particular thematic domains. Following Halliday's Systemic Functional Grammar, lexis is organised to realise experiential meaning, representing the world of the narrative through systematic vocabulary choices.

In story GPT-01, words such as "evening," "dusk," "yesterday," "tomorrow," "morning," and "midnight" form a semantic field of time and waiting. These lexical items recur throughout the narrative, establishing temporality as a central thematic concern. Similarly, words such as "stood alone," "silence," "empty house," "echoes,"

“absence,” and “solitude” construct a semantic field of isolation and loneliness. The clustering of these terms creates a coherent thematic atmosphere, but the associations are conventional and predictable, drawing on well-established literary vocabulary rather than generating novel lexical combinations.

The corpus also displays an elevated and reflective lexical register. Items such as “impermanence,” “dissolution,” “ritual,” “indifferent,” “indelible,” and “transcendence” contribute to a literary-philosophical tone. This suggests that AI models draw from conventional literary datasets to represent abstract emotional and philosophical experiences. The vocabulary produces thematic coherence but relies on familiar associations rather than generating the unexpected lexical choices that characterise innovative human-authored fiction. Human writers often juxtapose words from disparate registers or coin new terms to capture particular experiences; AI-generated texts tend to select from a repertoire of established literary vocabulary.

Across the corpus, five semantic fields recur with notable frequency: time and temporality (words such as “remember,” “forget,” “yesterday,” “forever”), isolation and loneliness (“alone,” “empty,” “silence,” “absence”), technology and machines (“screen,” “signal,” “network,” “data”), loss and grief (“gone,” “missing,” “hollow,” “mourning”), and nature and the elements (“rain,” “light,” “shadow,” “wind”). Table 3 presents the distribution of these semantic fields.

Table 3: Recurring Semantic Fields Across the Corpus

Semantic Field	Representative Lexis	Stories in Which Field Appears
Time and temporality	remember, forget, yesterday, forever, dusk, dawn	18 of 20
Isolation and loneliness	alone, empty, silence, absence, solitude, hollow	17 of 20
Technology and machines	screen, signal, network, data, algorithm, pixel	14 of 20
Loss and grief	gone, missing, mourning, hollow, ache, fade	16 of 20
Nature and elements	rain, light, shadow, wind, water, stone	15 of 20

The recurrence of these semantic fields suggests that AI models draw on conventional thematic associations in their training data, reproducing the vocabulary

that is statistically associated with particular themes in the literary corpus on which they were trained.

Figurative Expressions

The findings reveal that figurative language is a central meaning-making device in the AI corpus. Metaphor, simile, and personification are used frequently to translate abstract experiences into concrete, sensory images. In story GPT-01, memory is described as “a photograph left too long in the sun,” a metaphor that links the fading of a photographic image to the erosion of memory over time. The metaphor is effective in its immediate context, drawing on a familiar cultural association between photographs and memory. Similarly, in story CL-03, loneliness is rendered as “a radio tuned to an empty frequency,” a simile that draws on technological imagery to represent emotional isolation.

Personification is employed to animate inanimate elements of the narrative world. In story GM-02, the rain is described as having “its own agenda, its own quiet persistence,” attributing intentionality to a natural phenomenon. In story GPT-06, shadows are said to “gather in corners like conspirators,” investing them with agency and purpose. These personifications contribute to the atmospheric quality of the narratives, creating a sense of a world charged with latent meaning.

However, close analysis reveals that most figurative expressions in the corpus follow established literary patterns. They draw on conceptual mappings that are conventional in English literary discourse: memory as a fading image, loneliness as empty space, time as a flowing river, grief as a physical weight. While these metaphors are effective in their individual contexts, their recurrence across multiple stories and their alignment with well-worn literary conventions suggest that AI-generated figurative language operates within a relatively constrained repertoire. Human writers often generate novel metaphors that create unexpected conceptual connections; AI systems, by contrast, reproduce the figurative patterns that are statistically prevalent in their training data.

This finding is consistent with the broader understanding of how language models operate, as articulated by Bender et al. (2021). The model does not create metaphors through a process of imaginative conceptual mapping. It generates sequences that are probable given the patterns in its training data. The result is figurative language that is competent and occasionally striking but that rarely achieves the defamiliarising power of the most innovative human literary metaphors.

Graphological Features

Graphological analysis demonstrates that AI-generated stories adhere to conventional literary formatting with a high degree of consistency. Paragraphing is orderly, with paragraph breaks occurring at logical points in the narrative, typically marking shifts in time, location, or focus. Punctuation is generally standard, with commas, periods, quotation marks, and paragraph indentation used in accordance with conventional English orthography.

Devices such as dashes and semicolons are used to regulate pacing and to connect related clauses. In story CL-05, the sentence “They did not understand the warmth; they simulated it anyway” uses a semicolon to juxtapose two closely related propositions, creating a pause that invites reflection on the relationship between understanding and simulation. In story GPT-07, dashes are employed to set off parenthetical observations, creating a more conversational, reflective rhythm.

Despite this general conformity to orthographic norms, the corpus also exhibits minor irregularities. In several stories, punctuation marks are inconsistently spaced, with periods and commas sometimes appearing without the expected following space. Quotation marks occasionally open without closing, or close without having opened. Capitalisation is sometimes inconsistent, particularly in the representation of dialogue. Table 4 summarises the graphological features.

Table 4: Graphological Features Across the Corpus

Feature	Observation
Paragraphing	Conventional; logical breaks; consistent indentation
Punctuation	Generally standard; minor spacing inconsistencies in 25% of stories
Capitalisation	Generally correct; occasional errors in dialogue attribution
Orthographic errors	Minor; no major misspellings; occasional homophone confusion
Typography	Standard; no unconventional use of font or layout

These irregularities, while minor, are significant because they function as markers of algorithmic generation. They suggest a system that has learned the general rules of English orthography but does not apply them with the consistency of an

attentive human writer or editor. Graphology thus emerges as the linguistic level at which the mechanical character of AI-generated text is most visible. Human writers often use graphological features expressively, employing unconventional punctuation, spacing, or typography to represent speech patterns, psychological states, or thematic concerns. AI-generated texts, by contrast, exhibit a baseline of standard formatting punctuated by occasional algorithmic errors.

DISCUSSION

The findings of this study reveal that AI-generated short stories are characterised by a distinctive combination of linguistic fluency and stylistic predictability. Across the levels of syntax, lexico-semantics, figurative language, and graphology, the texts demonstrate a consistent capacity to reproduce conventional literary forms while exhibiting recurring patterns that distinguish them from human-authored fiction.

Syntactically, the AI stories are organised through a high degree of structural symmetry. Declarative sentences, paratactic organisation, parallelism, and anaphoric repetition create a measured, controlled narrative rhythm. This syntactic consistency produces texts that are clear, readable, and effectively paced, but that lack the syntactic variation and complexity that characterise much human literary writing. The predominance of paratactic over hypotactic structures limits the capacity of the texts to represent complex logical and temporal relationships, while the recurrence of parallel constructions and anaphoric patterns creates a formulaic quality that becomes apparent across multiple stories.

Lexico-semantically, the AI stories construct meaning through tightly clustered semantic fields that are conventional in their associations. The vocabulary of isolation, time, memory, loss, and technology recurs across the corpus, drawing on established literary associations rather than generating novel lexical configurations. The lexical register is consistently elevated and reflective, producing a uniform tone that is recognisably literary but that lacks the register shifts and lexical surprises that characterise innovative human-authored fiction.

Figuratively, the AI stories employ metaphor, simile, and personification effectively but conventionally. The figurative expressions draw on well-established conceptual mappings and are effective in their immediate contexts without achieving the defamiliarising power of the most striking human literary metaphors. This finding is consistent with the understanding that AI language models operate by generating statistically probable sequences rather than through creative conceptual innovation (Bender et al., 2021).

Graphologically, the AI stories adhere to standard orthographic conventions with minor irregularities that function as markers of algorithmic generation. The

orderly formatting contributes to the professional appearance of the texts while also revealing their non-human origin.

These findings have several implications for literary studies. First, they demonstrate that AI-generated fiction constitutes a distinct text type with its own characteristic linguistic features. These features are not simply inferior versions of human literary prose; they represent a different mode of text production altogether. Understanding this difference is important for literary scholars who seek to evaluate and interpret AI-generated literature.

Second, the findings have implications for creative writing pedagogy. As AI writing tools become more accessible, students will increasingly encounter and use them. Writing programmes should equip students with the critical skills to evaluate AI-generated texts, to understand their linguistic characteristics, and to make informed decisions about their use. The stylistic frameworks employed in this study provide one model for such critical engagement.

Third, the findings contribute to the broader conversation about the nature of creativity and the relationship between human and machine production of cultural artefacts. The linguistic patterns identified in this study suggest that while AI can reproduce the surface features of literary prose, it does so through fundamentally different processes than those employed by human writers. This distinction has implications for how creativity is understood and valued in an age of increasingly sophisticated AI.

CONCLUSION

This study examined the linguistic features of AI-generated short stories through the lenses of syntax, lexico-semantics, figurative language, and graphology. The findings demonstrate that AI texts are linguistically consistent and capable of reproducing many conventional features of literary prose. The stories exhibit controlled sentence structures, recurring semantic fields such as isolation, time, and memory, conventional metaphors and personifications, and generally standard formatting.

However, this consistency also constitutes a limitation. The AI-generated stories rely on repetitive sentence patterns, familiar word associations, and conventional figurative language. They lack the rule-breaking, irregularity, and unexpected linguistic choices that give much human literary writing its distinctive power and originality. The very predictability that makes AI-generated texts fluent and readable also makes them formulaic and, ultimately, limited in their capacity for genuine literary innovation.

The study contributes to the emerging scholarly conversation about AI and literature by providing a systematic linguistic analysis of machine-generated fiction. It demonstrates that stylistic frameworks developed for the analysis of human-authored texts can be productively applied to AI-generated texts, and it identifies specific

linguistic patterns that characterise the current generation of AI-generated literary prose.

RECOMMENDATIONS

Based on the findings of this study, the following recommendations are offered.

First, university programmes in stylistics, linguistics, and literary studies should incorporate the analysis of AI-generated texts into their curricula. This would enhance students' understanding of how linguistic structures shape meaning while fostering critical awareness of the relationship between linguistic form and authorial intention, whether human or computational. Specific modules on digital stylistics and AI-generated literature could be developed to equip students with the analytical tools demonstrated in this study.

Second, scholars in literary studies and the digital humanities should critically examine the cultural and ideological assumptions embedded in AI-generated texts. Because these texts are produced by models trained on large corpora of existing human writing, they may reproduce and amplify the biases, assumptions, and representational patterns present in their training data. Understanding these dimensions of AI-generated literature is an important task for future research.

Third, creative writing programmes should engage with AI as both a tool and a subject of critical inquiry. Students should be encouraged to experiment with AI-assisted writing while also developing the critical skills to evaluate the linguistic and literary qualities of machine-generated text. This dual approach would prepare students for a literary landscape in which human and machine writing increasingly coexist and interact.

Fourth, future research should adopt comparative frameworks that directly examine human-authored and AI-generated texts using identical prompts. Such comparative studies would more precisely delineate the linguistic boundaries between human and machine creativity and would strengthen the evidence base for claims about the distinctive features of AI-generated prose. Longitudinal studies tracking the evolution of AI-generated literary style as language models continue to develop would also be valuable.

REFERENCES

Adekanmbi, A. M., Chimene-Wali, N. C., & Amadi, G. N. (2026). Digital Identity and Linguistic Style: A Sociolinguistic Analysis of Nigerian X (Twitter) Users. *Performance: Journal of Law and Humanities*, 4(1), 45-57.

- Adekanmbi, A. M., Chimene-Wali, N. C., & Amadi, G. N. (2026). Digital Identity and Linguistic Style: A Sociolinguistic Analysis of Nigerian X (Twitter) Users. *Performance: Journal of Law and Humanities*, 4(1), 45-57.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610–623). Association for Computing Machinery.
- Boden, M. A. (2016). *AI: Its nature and future*. Oxford University Press.
- Carter, R. (2010). *Language and creativity: The art of common talk*. Routledge.
- Crystal, D. (2003). *The Cambridge encyclopedia of the English language* (2nd ed.). Cambridge University Press.
- Elkins, K., & Chun, J. (2020). Can GPT-3 write a short story? *Journal of Creative Writing Studies*, 5(2), 1–13.
- Fan, A., Lewis, M., & Dauphin, Y. (2019). Hierarchical story generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 5250–5260). Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1519>
- Floridi, L., & Chiriatti, M. (2020). GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30(4), 681–694.
- Gunser, V., Gottschling, S., & Brucker, B. (2022). Human and machine literary production: A comparative study. *Poetics*, 92, 101–115.
- Halliday, M. A. K. (1985). *An introduction to functional grammar*. Edward Arnold.
- Halliday, M. A. K., & Matthiessen, C. M. I. M. (2004). *An introduction to functional grammar* (3rd ed.). Hodder Arnold.
- Kybartas, B., & Bidarra, R. (2017). A survey on story generation techniques for authoring computational narratives. *IEEE Transactions on Computational Intelligence and AI in Games*, 9(3), 239–253.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. University of Chicago Press.
- Leech, G. N. (2008). *Language in literature: Style and foregrounding*. Pearson Longman.
- Leech, G. N., & Short, M. H. (2007). *Style in fiction: A linguistic introduction to English fictional prose* (2nd ed.). Pearson Longman.
- Mensah, E., Irek, N., Nwogu, A., & Iloh, Q. (2026). The crown gathers wealth: The symbolic significance of the crown in Yoruba personal naming practices. *Genealogy*, 10(1), 17.
- Nwogu, A., Anyanwu, O., & Osuagwu, E. (2024). Indigenous Igbo endearment names for wives as communicative acts in socio-cultural contexts. *Journal of the Nigerian Languages Project*, 6(2), 125-151.

- Ogurinka, G. W., & Chimene-Wali, N. (2026). Interweaving Game-Based Learning and Competition-Based Learning Methodologies: The Letterush Case Study. *Advances in Law, Pedagogy, and Multidisciplinary Humanities*, 4(1), 10-26.
- Ogurinka, G. W., & Chimene-Wali, N. (2026). Interweaving Game-Based Learning and Competition-Based Learning Methodologies: The Letterush Case Study. *Advances in Law, Pedagogy, and Multidisciplinary Humanities*, 4(1), 10-26.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems* 30 (pp. 5998–6008). Curran Associates.