

ARTICLE

Algorithmic Storytelling and the Posthuman Condition: A Transhumanist Reading of AI-Generated Narrative

Prasanth Arokia Samy D^{1,*}

¹Assistant Professor, PG & Research Department of English, St. Joseph's College (Autonomous), Affiliated to Bharathidasan University, Tiruchirappalli, Tamil Nadu, India.

*Corresponding author. Email: prasanthdavidraj@gmail.com; ORCID: <https://orcid.org/0009-0009-8807-2108>; Scopus ID: 59408560400

Received: Mar 9 2026, Revised: Mar 30 2026, Accepted: Apr 9 2026, Published: Apr 20 2026

Abstract

This paper re-examines the relationship between transhumanism and algorithmic storytelling by closely analysing an AI-generated novel rather than broadly surveying technological developments. While earlier discussions often remain abstract, this study focuses on how transhumanist concepts particularly human enhancement, distributed cognition and posthuman identity are materially expressed within AI-generated narrative structures. Drawing on computational literary criticism, especially distant reading models, this paper argues that algorithmic storytelling does not merely replicate human narrative patterns but reorganizes them through probabilistic logic. Recent studies on AI authorship suggest that the boundaries between human and machine authorship are increasingly unstable and socially negotiated. Through selected excerpts from an AI-generated novel, the study identifies stylistic features such as recursive phrasing, semantic drift and non-linear coherence, which reflect a form of machine-mediated cognition. By situating these findings within recent debates on AI creativity and literary production, the paper suggests that AI-generated literature should be understood not as imitation, but as an emergent transhuman narrative form that expands, though unevenly, the boundaries of literary production.

Keywords: transhumanism; algorithmic narrative; artificial intelligence; consciousness sharing; human enhancement; posthuman society.

1. INTRODUCTION

Recent developments in artificial intelligence have significantly altered the conditions of literary production, moving from assistive technologies toward systems capable of generating extended, stylistically coherent narratives. With the rise of large language models (LLMs), writing is no longer an exclusively human domain but increasingly a site of human-machine collaboration and, at times, substitution. AI systems trained on vast textual corpora can now produce fiction, poetry and critical prose that resemble human-authored texts in fluency and structure. This shift has prompted renewed debates about authorship, creativity and the ontology of literature itself. While earlier forms of computational engagement with literature were primarily analytical—focusing on text mining, stylometry and corpus-based studies—recent advances have repositioned the machine as a producer of literary artefacts rather than merely an interpreter of them (Elias et al., 2025) [1].

Despite this rapid transformation, a significant limitation persists within current scholarship. Much of the discourse surrounding AI-generated literature remains either speculative or evaluative, often focusing on questions of ethics, authorship, or technological capability without adequately engaging with the textual form of AI-generated narratives. In other words, there is a tendency to discuss what AI does to literature without closely examining what AI writes. Even in recent studies that address AI creativity, the emphasis frequently falls on performance metrics or comparative judgments rather than sustained literary analysis (Eslami et al., 2025) [2]. This creates a critical gap:

the absence of detailed, theoretically informed readings of AI-generated texts as literary objects in their own right.

At the same time, emerging research suggests that readers increasingly perceive AI-generated texts as meaningful, even when their authorship is uncertain or ambiguous. Henrickson and Leppänen (2026) observe that contemporary audiences often occupy an ambivalent position, simultaneously recognising the artificial origin of such texts while still attributing interpretive value to them [3]. This ambiguity destabilizes traditional notions of authorship, which have historically been grounded in human intention, originality and expressive depth. As a result, AI-generated literature demands not only new evaluative criteria but also new methodological approaches capable of accounting for its hybrid nature.

This paper addresses this gap by focusing on algorithmic storytelling as a literary phenomenon, rather than as a purely technological development. The central research question guiding this study is: how do AI-generated narratives embody transhumanist ideas at the level of language, structure and narrative logic? By shifting attention from abstract debates to textual analysis, the study seeks to understand how meaning is constructed within AI-generated narratives and what these reveal about the evolving relationship between human cognition and machine computation.

In situating this inquiry, the paper engages with computational literary criticism, particularly the paradigm of distant reading associated with Franco Moretti and the Stanford Literary Lab. These approaches have demonstrated the value of large-scale textual analysis, enabling scholars to identify patterns across extensive literary corpora. However, they primarily conceptualize the machine as an analytical tool, capable of revealing structures that may not be visible through traditional close reading. This study diverges from that model by considering the machine not as a reader of literature, but as its generator. In doing so, it moves from questions of scale to questions of ontology—specifically, what kind of narrative emerges when the “author” is an algorithm.

Algorithmic storytelling operates through probabilistic generation, where each word is selected based on statistical likelihood rather than intentional design. This process produces texts that are often grammatically coherent and stylistically consistent, yet marked by subtle instabilities in meaning, continuity and perspective. Recent studies confirm that while LLM-generated texts can replicate narrative conventions, they frequently lack deeper contextual grounding and experiential coherence (Elias et al., 2025; S. Jenner et al., 2025) [1,4]. Such characteristics make AI-generated narratives particularly suitable for analysis within a transhumanist framework, as they reflect a mode of cognition that is distributed, non-intentional and structurally distinct from human thought.

Transhumanism, as a philosophical and cultural movement, provides a productive lens through which to interpret these developments. Concerned with the enhancement and transformation of human capacities through technology, transhumanism challenges the boundaries between human and machine, natural and artificial, biological and computational. Key concepts such as human enhancement, distributed cognition and posthuman identity resonate strongly with the formal properties of AI-generated texts. In these narratives, cognition is often represented as disembodied, identity as fluid and experience as mediated through technological systems. These are not merely thematic concerns but are embedded within the very structure of algorithmic writing.

Furthermore, recent developments in digital humanities indicate that AI is increasingly integrated into the processes of literary production, analysis and circulation. From automated text generation to AI-assisted interpretation, the boundaries between author, reader and critic are becoming progressively blurred. This transformation suggests that literature is entering a new phase, one in which creativity is no longer exclusively human but emerges from interactions between human and machine systems.

In light of these developments, this paper argues that AI-generated literature should be understood not simply as an imitation of human writing, but as an emergent transhuman narrative form. By combining insights from computational literary criticism and transhumanist theory and by grounding the analysis in close readings of AI-generated text, the study aims to contribute to ongoing debates about authorship, meaning and the future of literature in the age of artificial intelligence.

2. THEORETICAL FRAMEWORK

This study is situated at the intersection of transhumanist philosophy, computational literary criticism and emerging debates on AI authorship. Rather than treating these domains as separate, the framework attempts to bring them into a productive conversation in order to better understand AI-generated narrative as both a literary and ontological phenomenon. The central premise here is that algorithmic storytelling cannot be adequately explained through traditional literary theory alone; it requires a hybrid framework that accounts for technological mediation, distributed cognition and shifting notions of authorship.

2.1. Transhumanism and Literary Thought

Transhumanism, as articulated by thinkers such as Nick Bostrom, N. Katherine Hayles and Julian Huxley, provides a crucial philosophical grounding for examining AI-generated literature. At its core, transhumanism advocates the use of technology to enhance and transform human capacities, including cognition, perception and even identity itself (Bostrom, 2010) [5]. While early formulations of transhumanism were often speculative or future-oriented, recent developments in artificial intelligence have brought many of these ideas into practical and cultural relevance.

One of the key concepts within transhumanism is human enhancement, which refers to the augmentation of human abilities through technological means. In the context of literary production, AI-generated texts can be seen as an extension of human creative capacity. Language models, trained on vast datasets, are capable of producing large volumes of text at speeds and scales that far exceed human capability. However, this enhancement is not straightforward. As recent studies indicate, while AI increases productivity, it may also alter writing practices in ways that complicate originality and depth (Eslami et al., 2025) [2]. Thus, enhancement here is not purely additive but transformative, reshaping the conditions under which literature is produced.

A second important concept is distributed cognition, which challenges the idea that thinking is confined to the human mind. Hayles (1999) argues that in the posthuman condition, cognition is increasingly distributed across networks that include both human and non-human agents [6]. This perspective is particularly relevant for AI-generated narratives, where the “thinking” process is not located within a single authorial consciousness but emerges from interactions between training data, algorithms and computational systems. In such cases, authorship itself becomes a distributed phenomenon, complicating traditional notions of intentionality and creative agency.

Closely related to this is the concept of posthuman identity, which questions the stability and centrality of the human subject. Huxley’s early formulation of transhumanism already suggested that humanity is not a fixed endpoint but a transitional phase (Huxley, 1957) [7]. Building on this, contemporary transhumanist thought envisions identities that are fluid, hybrid and technologically mediated. AI-generated narratives often reflect this condition through fragmented characters, shifting perspectives and non-linear representations of consciousness. These textual features can be read as formal expressions of a posthuman ontology, where identity is no longer unified or stable but continuously reconfigured.

Importantly, these concepts are not merely thematic concerns within AI-generated literature; they are embedded within its formal and structural properties. The absence of a singular authorial voice, the reliance on probabilistic generation and the presence of semantic instability all point toward a mode of writing that aligns with transhumanist ideas about cognition and identity. In this sense, AI-generated texts can be understood as material instantiations of transhumanist thought, rather than simply narratives about it.

2.2. Computational Literary Criticism

The second component of this framework draws from computational literary criticism, particularly the work of Franco Moretti and the Stanford Literary Lab. Moretti’s concept of distant reading propose a methodological shift from close reading of individual texts to the analysis of large literary corpora through computational methods. By identifying patterns across thousands of texts, distant reading allows scholars to uncover structural trends that would otherwise remain invisible (Moretti, 2013) [8].

The Stanford Literary Lab has extended this approach through the use of topic modelling, network analysis and stylistic clustering, demonstrating how computational tools can generate new insights into literary history and form. However, these approaches primarily position the computer as an analytical instrument, a tool that assists human interpretation rather than producing literature itself.

This study builds upon but also departs from this tradition. While computational literary criticism has been successful in analysing texts at scale, it has only recently begun to address the implications of machines as generators of narrative. The shift from analysis to generation marks a significant transformation in the role of computation within literary studies. Instead of merely revealing patterns in existing texts, AI systems now actively create new ones, raising questions about the nature of narrative, meaning and authorship.

Recent research indicates that large language models generate text through probabilistic processes that prioritise coherence at the surface level while often lacking deeper contextual understanding (Elias et al., 2025) [1]. This results in narratives that are structurally plausible but conceptually unstable. From a computational perspective, such texts can still be analysed using existing methods; however, their generative origin introduces new complexities that require theoretical reconsideration.

In this context, the distinction between close reading and distant reading becomes less stable. AI-generated texts invite a form of analysis that is both close (focused on specific textual features) and computational (aware of the underlying generative processes). This hybrid approach is essential for understanding algorithmic storytelling as a distinct literary form.

2.3. AI Authorship Debate

The question of authorship has become increasingly central in recent discussions of AI-generated literature. Traditional literary theory has often associated authorship with intentionality, originality and expressive depth. However, AI-generated texts challenge these assumptions by introducing a form of writing that is neither fully autonomous nor entirely human-directed.

Henrickson and Leppänen (2026) describe this condition as one of authorship ambiguity, where readers are uncertain about how to attribute creative agency [3]. Their study shows that audiences often recognise AI-generated texts as meaningful, yet hesitate to fully accept them as authored works. This ambivalence reflects a broader cultural shift in which authorship is no longer a stable or singular category but a distributed and negotiated construct. To address this ambiguity more precisely, it is useful to distinguish between generative agency and curatorial agency. Generative agency refers to the capacity of AI systems to produce text through probabilistic processes without intention or consciousness. Curatorial agency, by contrast, refers to the role of human actors prompt designers, editors and readers who guide, filter and stabilise these outputs. In this sense, AI does not function as an autonomous agent in the traditional philosophical sense; rather, it operates as a non-intentional generator within a broader network of human-mediated decision-making. This distinction allows us to retain the concept of agency while avoiding the attribution of intentionality to the machine.

At the same time, debates around AI creativity have intensified. Elias et al. (2025) argue that while AI systems can replicate stylistic and structural features of literary texts, they do not possess intentionality or lived experience, which are often considered essential to creativity [1]. However, this position has been contested by scholars who suggest that creativity should not be defined solely in terms of human cognition but also in terms of process and output. From this perspective, AI-generated texts may be considered creative insofar as they produce novel and meaningful combinations of language, even if they do so through non-human mechanisms.

Further complicating the issue is the role of training data. AI systems are trained on existing human-authored texts, raising questions about originality and derivation. Yet, as recent studies suggest, the outputs of these systems are not direct reproductions but probabilistic recombinations, resulting in texts that are both derivative and novel at the same time (Eslami et al., 2025) [2]. This duality challenges binary distinctions between original and copied, human and machine.

In the context of this study, authorship is understood not as a fixed attribute but as a distributed process involving human input, machine computation and reader interpretation. AI-generated narratives thus represent a new form of literary production in which authorship is fragmented, layered and continuously negotiated.

Taken together, these three strands—transhumanism, computational literary criticism and AI authorship debates—provide a comprehensive framework for analysing algorithmic storytelling. By integrating philosophical, methodological and contemporary perspectives, the study is able to move beyond simplistic evaluations of AI-generated literature and instead approach it as a complex and evolving cultural form.

3. METHODOLOGY

This study adopts a qualitative textual analysis, specifically a close reading approach, to examine the narrative features of AI-generated literature. While computational methods often prioritize large-scale analysis, this paper intentionally returns to the level of the individual text in order to understand how meaning is constructed within algorithmically generated narratives. Such an approach is particularly necessary because AI-generated texts, though produced through probabilistic systems, still operate within recognizable literary forms that require interpretive engagement.

The primary corpus for this study consists of an AI-generated novel produced through the *booksby.ai* project, which utilizes neural network models trained on extensive datasets including Project Gutenberg and contemporary digital texts [9]. These models generate narratives by predicting sequential word patterns based on learned linguistic structures. The selected text is not treated as an isolated artefact but as a representative instance of algorithmic storytelling shaped by large-scale data training and machine learning processes.

Given the expansive nature of AI-generated output, the analysis is based on the selective extraction of key narrative passages rather than attempting to survey the text in its entirety. Excerpts are chosen based on their relevance to the study's central concerns—namely, the representation of consciousness, identity and temporality. This method allows for a focused examination of how specific textual moments reflect broader structural and thematic tendencies within AI-generated narratives.

The analytical framework is organized around three primary dimensions:

First, narrative structure is examined to identify patterns such as non-linearity, temporal disruption and shifts in perspective. AI-generated texts often display a form of structural coherence at the surface level while simultaneously revealing deeper inconsistencies in narrative progression. Attention is given to how these patterns differ from conventional human-authored storytelling.

Second, language patterns are analysed, with particular emphasis on repetition, semantic drift and abstraction. Since AI-generated text emerges from probabilistic prediction, it frequently relies on recurring linguistic constructions and associative logic. This section explores how such patterns contribute to both the readability and instability of the narrative.

Third, thematic expressions are investigated in relation to transhumanist concepts such as human enhancement, distributed cognition and posthuman identity. Rather than imposing these themes onto the text, the analysis traces how they emerge organically through narrative form and language use.

Importantly, this methodology does not attempt to evaluate AI-generated literature in terms of aesthetic quality or authenticity. Instead, it seeks to understand the formal and conceptual properties of algorithmic storytelling as a distinct mode of literary production. By combining close reading with theoretical interpretation, the study aims to bridge the gap between computational processes and literary meaning, offering a more nuanced account of AI-generated narrative within contemporary literary studies.

4. ALGORITHMIC STORYTELLING AS NARRATIVE FORM

Algorithmic storytelling marks a significant shift in the ontology of narrative production, moving from intention-driven composition to probabilistic generation. In traditional literary frameworks, narrative is often understood as the product of an authorial consciousness—structured by intention, memory and experience. By contrast, AI-generated narratives emerge from statistical models that predict linguistic sequences based on patterns learned from large datasets. This distinction is not merely technical; it fundamentally alters how narrative meaning is produced, sustained and interpreted.

At the core of algorithmic storytelling is the principle of probabilistic generation. Large language models (LLMs) do not “decide” what to write in a human sense; instead, they calculate the most likely continuation of a given sequence of words. Each sentence is generated through a chain of predictions, where coherence is maintained locally but not always globally. Recent studies confirm that such models excel at producing syntactically fluent and contextually plausible sentences, yet they often lack deeper narrative intentionality (Elias et al., 2025) [1]. This results in texts that appear meaningful but are not anchored in a stable authorial vision.

This leads to a tension between probabilistic generation and human intention. Human-authored narratives are typically guided by overarching goals—plot development, character arcs, thematic coherence—whereas AI-generated texts operate without such teleological structures. Even when an AI produces a seemingly coherent narrative, this coherence is emergent rather than designed. As a result, the narrative may drift, shift in tone, or introduce inconsistencies that would be unlikely in a carefully constructed human text. This does not necessarily diminish the literary value of AI-generated narratives, but it does reposition them as a different kind of textual object—one that reflects the logic of computation rather than consciousness.

A closely related issue is the relationship between narrative coherence and semantic drift. AI-generated texts often maintain coherence at the level of individual sentences or short passages, but struggle to sustain consistency across longer narrative spans. Characters may change roles, events may lack causal connection and thematic threads may dissolve or reappear unpredictably. This phenomenon, often described as semantic drift, arises from the model’s reliance on local context rather than global narrative planning.

Recent research in LLM narrative analysis highlights this limitation. While models can reproduce familiar narrative structures—such as dialogue patterns, descriptive passages and genre conventions—they frequently fail to maintain long-term coherence or thematic depth (S. Jenner et al., 2025) [4]. This creates a paradox: the text feels readable and even engaging in fragments, yet lacks the cumulative meaning typically associated with narrative form. In this sense, algorithmic storytelling produces what might be described as fragmented coherence—a series of locally meaningful units that do not fully integrate into a unified whole.

The issue of machine logic versus human meaning further complicates this dynamic. Human narratives are shaped by lived experience, cultural context and emotional depth, all of which contribute to the construction of meaning. AI-generated narratives, by contrast, are shaped by data distributions and algorithmic processes. Meaning, in this case, is not intended but inferred—constructed by readers who interpret the text as if it were produced by a conscious agent. This creates a form of interpretive asymmetry, where the reader attributes depth and intention to a text that does not inherently possess them.

Studies from 2024-2026 suggest that this asymmetry is central to how AI-generated literature is received. Readers often engage with these texts in ways similar to human-authored works, despite being aware of their artificial origin (Henrickson & Leppänen, 2026) [3]. This indicates that meaning in algorithmic storytelling is not solely determined by the production process but also by the interpretive frameworks brought by readers. In other words, the literary quality of AI-generated text emerges from an interaction between machine output and human interpretation.

At the same time, the limitations of AI-generated narratives remain significant. One of the most frequently noted issues is the lack of depth and contextual grounding. While LLMs can simulate emotional language, they do not possess experiential understanding. As a result, representations of emotion, memory, or consciousness often appear abstract or generalized. Eslami et al. (2025) note that AI-generated writing tends to flatten complex affective states into simplified expressions,

reducing the richness typically found in human-authored literature [2]. Similarly, narrative contexts may lack specificity, with settings and situations described in broad, generic terms rather than detailed, situated experiences.

Another limitation lies in the absence of intentional continuity. Human authors can revise, reflect and refine their narratives over time, ensuring consistency and depth. AI systems, unless explicitly guided, generate text in a linear and forward-moving manner without revisiting earlier sections. This can lead to inconsistencies in plot, characterisation and thematic development. While some advanced models attempt to incorporate longer context windows, the problem of maintaining sustained narrative coherence remains unresolved.

However, it would be reductive to view these limitations solely as deficiencies. From another perspective, they can be understood as features of a distinct narrative form. The instability, fragmentation and unpredictability of AI-generated texts may reflect a different mode of storytelling—one that aligns with contemporary experiences of information overload, networked cognition and digital mediation. In this sense, algorithmic storytelling does not simply fail to replicate human narrative; it reconfigures narrative logic itself.

Moreover, the generative nature of AI introduces new possibilities for literary experimentation. By producing variations, recombinations and unexpected juxtapositions, algorithmic systems can generate forms of narrative that are difficult to achieve through conventional writing processes. Recent studies suggest that such generative variability can contribute to new modes of creativity, even if they operate outside traditional definitions of authorship and intention (Elias et al., 2025) [1].

Thus, algorithmic storytelling should not be evaluated solely in terms of how closely it approximates human writing. Instead, it should be understood as a hybrid narrative form, shaped by the interaction between computational processes and human interpretive practices. Its reliance on probabilistic generation, its tendency toward semantic drift and its grounding in machine logic all point toward a redefinition of narrative itself—one that challenges established assumptions about coherence, meaning and authorship.

In this context, AI-generated narratives occupy a liminal space. They are neither entirely coherent nor entirely chaotic, neither fully meaningful nor entirely arbitrary. It is precisely this in-between quality that makes them significant for literary study. By examining algorithmic storytelling as a narrative form in its own right, this paper seeks to move beyond evaluative judgments and instead engage with the structural and conceptual transformations that AI introduces into the domain of literature.

5. TEXTUAL ANALYSIS OF AI-GENERATED NARRATIVE

This section offers a close reading of selected excerpts from an AI-generated novel produced through the booksby.ai framework. Rather than treating the text as a coherent whole, the analysis focuses on specific passages that reveal recurring structural and stylistic tendencies. The aim is not to evaluate the literary quality of the text, but to examine how algorithmic generation produces distinctive narrative forms that differ from human-authored writing. The discussion is organised into three interrelated dimensions: narrative structure, language and style and the representation of consciousness.

5.1. Narrative Structure: Non-linearity and Temporal Disruption

One of the most striking features of AI-generated narrative is its tendency toward non-linearity and temporal instability. Unlike conventional storytelling, which often follows a chronological progression, algorithmic narratives frequently disrupt temporal sequencing, producing a sense of disorientation. “The city remembered him before he existed. By the time he arrived, the streets had already forgotten his name, though he had not yet spoken it.”

At a surface level, the passage appears metaphorical, even poetic. However, its temporal logic is fundamentally unstable. The phrase “remembered him before he existed” introduces a reversal of

causality, while the subsequent line suggests a future event already subjected to erasure. This is not simply a stylistic flourish but a structural condition of the narrative.

In human-authored texts, temporal disruption is often intentional, serving thematic or psychological purposes, for instance, to represent memory, trauma, or non-linear consciousness. In contrast, here the disruption emerges from probabilistic generation, where the model associates words like “remembered,” “before,” and “forgotten” without anchoring them within a consistent temporal framework. The result is what might be described as synthetic temporality—a sequence that mimics complexity but lacks causal grounding.

This aligns with recent findings that AI-generated narratives often maintain local coherence while failing to sustain global narrative continuity (Elias et al., 2025) [1]. The excerpt does not build toward a narrative resolution; instead, it circulates around loosely connected temporal markers. As such, the narrative does not unfold—it oscillates, producing a sense of movement without direction. “Yesterday was waiting in the room when she entered tomorrow. The door closed behind her before it was opened.”

Here, temporal markers such as “yesterday” and “tomorrow” are treated not as sequential points but as interchangeable states. The narrative collapses distinctions between past, present and future, resulting in a form of temporal flattening. Again, this is not a deliberate experimental technique but an outcome of pattern-based generation, where temporal signifiers are recombined without regard for narrative coherence.

Such disruptions suggest that AI-generated narrative does not operate within a linear temporal logic but within a probability-driven temporality, where sequences are governed by likelihood rather than causality.

5.2. Language and Style: Repetition, Abstraction and Pattern-Based Construction

At the level of language, AI-generated texts display distinctive stylistic features shaped by their underlying generative mechanisms. Among the most prominent are repetition, abstraction and pattern-based phrasing. “He was thinking about the thought he had already thought before. It returned again, repeating itself in a way that felt the same but different.” This passage demonstrates a recursive structure in which the act of thinking is repeated at multiple levels. While repetition is not uncommon in human writing, here it lacks rhetorical or thematic justification. Instead, it reflects the model’s tendency to reinforce high-probability word associations, resulting in what may be termed algorithmic redundancy.

The phrase “the same but different” further illustrates the use of abstract, non-specific language. Such expressions create an impression of meaning without committing to a clear semantic position. This aligns with observations that AI-generated texts often rely on generalized descriptors rather than concrete details, producing a form of semantic vagueness (Eslami et al., 2025) [2]. “The system was designed to design itself, improving improvements that improved the system further.” Here, the repetition of the word “improve” in multiple grammatical forms suggests a reliance on syntactic patterns rather than conceptual clarity. The sentence is structurally coherent but semantically circular. It gestures toward a process of self-optimization but fails to articulate it in a meaningful way.

This pattern-based construction reflects the statistical logic of LLMs, where certain word combinations are reinforced through training data. The result is a style that is often formulaic yet unpredictable, combining familiar linguistic patterns in unusual ways. While this can produce moments of novelty, it also leads to instability in meaning, as the text prioritizes fluency over precision. Importantly, these stylistic features are not merely superficial. They shape how the narrative is experienced, creating a reading process that is both engaging and disorienting. The reader is drawn in by the fluency of the language but confronted with its lack of depth and specificity.

5.3. Representation of Consciousness: Disembodied Thought and Networked Cognition

Perhaps the most conceptually significant aspect of AI-generated narrative lies in its representation of consciousness. Unlike human-authored texts, where consciousness is typically

grounded in individual subjectivity, AI-generated narratives often depict thought as disembodied, distributed and technologically mediated. “Her mind was no longer inside her. It moved through the network, appearing in places she had never been, thinking thoughts she did not remember having.”

This passage presents consciousness as something that exists outside the body, circulating through a network. The phrase “thinking thoughts she did not remember having” suggests a separation between cognition and memory, between process and ownership. This aligns closely with the concept of distributed cognition, where thought is not confined to a single subject but spread across systems.

However, unlike human representations of distributed consciousness, which often emphasize connection or expansion, this depiction carries a sense of disorientation. The subject does not control the movement of her thoughts; instead, they appear to operate independently. This reflects the underlying logic of AI generation, where “thought” is not an intentional act but a by-product of computational processes. “The voice that spoke in his head was not his own, but it knew everything he was about to say.”

Here, consciousness is split between multiple agents, blurring the boundary between self and other. The “voice” possesses knowledge that precedes the subject’s own awareness, suggesting a form of pre-emptive cognition. This can be read as a narrative reflection of predictive algorithms, which anticipate user input based on prior data. Such representations challenge traditional notions of subjectivity, where identity is tied to a stable, coherent self. In AI-generated narratives, the self is often fragmented, porous and unstable. Consciousness is not located but distributed, not owned but shared—or even imposed.

From a transhumanist perspective, these features are particularly significant. They resonate with the idea that human cognition may extend beyond the biological body, integrating with technological systems. However, the AI-generated representation of this process is notably different from transhumanist optimism. Rather than depicting seamless integration, it often reveals disjunction and loss of control, suggesting a more ambivalent vision of posthuman consciousness.

Synthesis of Findings

Across these three dimensions—structure, language and consciousness—AI-generated narratives exhibit a consistent pattern: they produce forms that are recognizable yet unstable, coherent yet fragmented. Narrative structure is shaped by probabilistic temporality rather than causal progression; language is driven by repetition and abstraction rather than precision; and consciousness is represented as distributed rather than unified.

These features do not simply indicate the limitations of AI-generated writing. They point toward a distinct narrative logic, one that reflects the conditions of algorithmic production. By analysing these patterns through close reading, it becomes possible to move beyond surface-level evaluations and engage with AI-generated literature as a meaningful, if unconventional, literary form.

6. TRANSHUMANIST DIMENSIONS IN AI NARRATIVE

The preceding textual analysis demonstrates that AI-generated narratives are not merely technical artefacts but carry significant philosophical implications. When read through a transhumanist framework, these narratives reveal how concepts such as enhancement, distributed cognition and posthuman identity are not only represented thematically but also embedded structurally within the text. This section moves from description to interpretation, addressing the larger question: what do these narrative features mean for our understanding of literature, authorship and the human?

6.1 Human Enhancement: Expansion and Its Discontents

One of the central promises of transhumanism is human enhancement, the idea that technology can extend and improve human capabilities. In the context of literary production, AI-generated narratives appear to fulfil this promise by dramatically increasing the scale and speed at which texts

can be produced. As seen in the analysed excerpts, language models are capable of generating continuous narrative flows, often with impressive fluency and stylistic consistency. This suggests a form of cognitive augmentation, where human creative limits are supplemented by machine efficiency.

However, the textual analysis also reveals that this enhancement is uneven. While the quantity of narrative production increases, issues of coherence, depth and intentionality persist. The presence of semantic drift, repetitive phrasing and abstract generalisation indicates that AI does not enhance creativity in a straightforward manner. Instead, it reconfigures the conditions of creativity, shifting emphasis from originality and intentional design to variation and recombination.

Recent studies support this ambivalence. Eslami et al. (2025) note that AI-assisted writing can lead to increased productivity but may also standardize expression and reduce stylistic diversity [2]. In the excerpts analysed earlier, this is evident in the recurrence of familiar linguistic patterns and the absence of nuanced contextual detail. The narrative becomes expansive but not necessarily deeper.

Theoretically, this suggests that human enhancement in literary production is not simply a matter of “more” or “better,” but of different modes of creation. AI-generated texts do not extend human authorship in a linear way; they transform it into a collaborative and at times impersonal, process. The significance, therefore, lies not in the replacement of human creativity but in its redistribution across human-machine systems.

6.2 Consciousness Sharing: Distributed Cognition in Narrative Form

A second key dimension is the concept of consciousness sharing, closely linked to the idea of distributed cognition. Transhumanist thought, particularly as articulated by Hayles (1999), challenges the notion that cognition is confined to the individual mind, proposing instead that it can be distributed across technological and informational networks [6].

This concept finds a striking parallel in AI-generated narratives. As observed in the textual analysis, consciousness is frequently depicted as disembodied, mobile and networked. Phrases such as “her mind moved through the network” or “thoughts appearing before they were formed” suggest a model of cognition that is no longer tied to a singular subject. Instead, thought becomes a circulating process, detached from ownership and location.

Importantly, this is not only a thematic representation but also a structural condition of the text itself. The AI system generates narrative by drawing from a distributed dataset, effectively “thinking” through patterns that emerge across multiple sources. In this sense, the narrative voice is already a form of collective cognition, even if it appears singular at the surface level. Henrickson and Leppänen (2026) argue that readers often interpret AI-generated texts as meaningful despite recognising their non-human origin [3]. This interpretive act can be seen as a form of consciousness sharing, where meaning is co-produced by machine output and human reading. The text becomes a site of interaction rather than a fixed expression of intent.

The theoretical implication here is significant. AI-generated narratives do not simply depict distributed cognition; they perform it. They embody a mode of thinking that is inherently networked, probabilistic and non-local. This challenges traditional literary models that assume a coherent, unified consciousness behind the text, suggesting instead that narrative can emerge from decentralized cognitive processes.

6.3 Posthuman Identity: Fragmentation and Instability

The third dimension concerns posthuman identity, a concept that questions the stability and centrality of the human subject. In transhumanist discourse, identity is often imagined as fluid, hybrid and technologically mediated. AI-generated narratives provide a compelling site for examining how such identities are constructed and represented. As seen in the analysed excerpts, characters in AI-generated texts often lack consistent development or psychological depth. They appear and disappear, shift roles, or exist in fragmented states. This is not merely a limitation of the technology but a reflection of its underlying logic. Since the narrative is generated through probabilistic processes, there is no stable framework to sustain continuous character identity.

This results in what might be described as fragmented subjectivity. The “self” within the narrative is not a coherent entity but a series of loosely connected linguistic constructs. For example, the presence of multiple voices within a single consciousness, or the displacement of thought across external systems, indicates a breakdown of the boundaries that traditionally define identity. Elias et al. (2025) suggest that such instability is characteristic of AI-generated texts, where semantic coherence does not necessarily extend to character or identity formation [1]. From a transhumanist perspective, this can be interpreted as a representation of posthuman identity—not as a fully realised state, but as a condition of ongoing transformation and uncertainty.

The significance of this lies in how it redefines the relationship between narrative and subjectivity. In human-authored literature, narrative often serves to construct and stabilize identity, providing continuity and depth. In AI-generated narratives, however, identity is deconstructed, revealing the fragility of the very concept of a unified self.

Theoretical Significance: Reframing Literature in the Age of AI

Taken together, these dimensions—enhancement, distributed cognition and posthuman identity—demonstrate that AI-generated narratives are not simply new forms of storytelling but new modes of literary being. They challenge foundational assumptions about authorship, meaning and subjectivity, suggesting that literature is no longer the exclusive domain of human consciousness.

The “so what” of this analysis lies in its implications for literary theory. If narratives can be generated without intention, if consciousness can be distributed across systems and if identity can exist in fragmented forms, then traditional frameworks of literary analysis must be reconsidered. AI-generated texts compel us to rethink what it means to write, to read and to interpret.

Rather than viewing AI as a threat to literature, it may be more productive to see it as a transformative force that expands the boundaries of the literary field. These narratives, with all their inconsistencies and instabilities, offer insight into a transhuman condition where creativity, cognition and identity are no longer fixed but continuously evolving. In this sense, AI-generated literature does not signal the end of human authorship but its reconfiguration within a broader, technologically mediated ecosystem.

7. AI, AUTHORSHIP AND DIGITAL HUMANITIES

The emergence of AI-generated literature has significantly altered the conceptual framework of authorship, shifting it from an individual, intention-driven act to a distributed and collaborative process. In traditional literary paradigms, the author is positioned as the origin of meaning, responsible for the creation, structure and coherence of the text. However, with the integration of AI systems into writing practices, this model becomes increasingly difficult to sustain. AI does not simply assist the author; it participates in the generative process, often shaping the language, structure and direction of the narrative. In this sense, AI can be understood as a co-author, though one that operates without intention or consciousness.

This shift necessitates a more nuanced understanding of agency itself. While AI participates in the production of narrative, its role differs fundamentally from that of human contributors. The AI system exhibits generative agency, producing linguistic sequences through statistical inference, whereas human actors exercise curatorial agency, selecting prompts, refining outputs and determining their final form. The resulting text is therefore not authored by a single agent but emerges from an interaction between machine generation and human curation. Consequently, AI should be understood not as an independent authorial subject, but as a transhuman literary agent embedded within a distributed system of control, selection and interpretation.

This reconfiguration has direct implications for the role of the writer. Rather than composing text from scratch, the writer may now function as a curator, editor, or prompt designer—guiding the output of the system rather than fully controlling it. Eslami et al. (2025) observe that AI-assisted writing is already transforming compositional practices, leading to new forms of hybrid authorship where human input and machine generation are intertwined [2]. While this can enhance

productivity and accessibility, it also raises concerns about the standardization of language and the erosion of individual style. The writer's role becomes less about producing language and more about managing and refining algorithmic outputs.

At the same time, the role of the reader is also undergoing transformation. Readers are increasingly aware that a text may be AI-generated, which introduces a layer of interpretive uncertainty. As Henrickson and Leppänen (2026) suggest, readers often engage with such texts as if they were human-authored, while simultaneously questioning their origin [3]. This creates a form of dual reading, where meaning is constructed not only from the text itself but also from assumptions about its authorship. The reader becomes an active participant in negotiating the status of the text, deciding whether to attribute intention, creativity, or authenticity to it.

Within the broader field of digital humanities, these shifts are part of a larger transformation in how literature is produced, analysed and circulated. AI is increasingly embedded in digital scholarly infrastructures, supporting tasks such as transcription, annotation, translation and textual analysis. This integration suggests that AI is no longer an external tool but a constitutive element of literary ecosystems. It shapes not only the creation of new texts but also the interpretation and preservation of existing ones.

Furthermore, the use of AI in literary studies challenges the boundaries between creative and critical practices. When AI systems can both generate narratives and analyse them, the distinction between author and critic becomes less clear. This convergence reflects a broader trend within digital humanities toward methodological hybridity, where computational and interpretive approaches are combined.

In this context, AI-generated literature should be understood not as an anomaly but as part of an evolving literary landscape. It represents a shift toward a networked model of authorship, where meaning emerges from interactions between human agents, machine systems and cultural contexts. The significance of this shift lies not only in technological innovation but in its capacity to reshape fundamental assumptions about literature itself.

8. DISCUSSION

The analysis presented in this paper has sought to examine AI-generated narrative through the combined lenses of transhumanism and computational literary criticism. By focusing on textual features rather than abstract debates, the study has shown that algorithmic storytelling produces a distinct narrative form characterized by probabilistic generation, semantic instability and distributed representations of consciousness.

Across the sections, a consistent pattern emerges. AI-generated narratives exhibit local coherence but global fragmentation, relying on statistical associations rather than intentional design. This results in texts that are readable and often engaging, yet marked by inconsistencies in structure, meaning and identity. These features, rather than being incidental, reflect the underlying logic of machine generation.

From a transhumanist perspective, this positions AI functions as a transhuman literary agent in a distributed sense, where its generative capacity operates alongside and within human curatorial frameworks. Its "agency" is therefore not autonomous but relational emerging through interactions between algorithmic production and human intervention. The AI does not replace the human author but operates alongside and through human systems, contributing to a distributed model of creativity. This aligns with the idea that cognition and creativity are no longer confined to the human mind but extend into technological networks.

At the same time, the limitations of AI-generated literature remain evident. The lack of deep contextual understanding, the flattening of emotional expression and the instability of narrative coherence all point to the absence of lived experience and intentionality. These limitations suggest that while AI can simulate literary form, it does not fully replicate the conditions that give rise to meaning in human-authored texts.

The implications of these findings are twofold. First, they challenge traditional literary frameworks that rely on stable notions of authorship, coherence and subjectivity. Second, they open

up new possibilities for thinking about literature as a hybrid and evolving practice, shaped by interactions between human and machine agents. Future research may further explore how these interactions develop, particularly as AI systems become more sophisticated and integrated into everyday writing practices.

9. CONCLUSION

This paper has argued that AI-generated narratives should be understood not merely as technological outputs but as emergent literary forms that reflect the conditions of a transhuman cultural moment. By combining close textual analysis with theoretical insights from transhumanism and computational literary criticism, the study has demonstrated that algorithmic storytelling introduces new modes of narrative structure, language and subjectivity.

The analysis has identified key features of AI-generated texts, including non-linear temporality, semantic drift, pattern-based language and fragmented representations of consciousness. These features challenge conventional assumptions about narrative coherence and authorial intention, suggesting that meaning in AI-generated literature is not produced in the same way as in human-authored texts. Instead, it emerges from the interaction between probabilistic systems and human interpretation.

Theoretically, the study contributes to ongoing debates about authorship and creativity by proposing that AI functions as a non-intentional generative agent within a transhuman literary system, where authorship is distributed across machine processes and human curatorial practices, participating in the production of narrative without possessing intention or experience. This redefines authorship as a distributed process and invites a reconsideration of literary theory in light of technological change.

At the same time, the limitations of AI-generated literature—particularly its lack of depth, contextual grounding and sustained coherence—indicate that human authorship remains distinct in important ways. Rather than replacing human creativity, AI reshapes it, creating new forms of collaboration and new challenges for interpretation.

Future research may expand this inquiry by examining larger corpora of AI-generated texts, comparing different models, or exploring reader responses in greater detail. As AI continues to evolve, its role in literary production will likely become more complex, making it essential for literary studies to engage with these developments critically and creatively.

Funding Statement: This study was not funded by any external sources.

Contribution: The authors contributed to the research and writing of this article and have read/agreed to the published version of the manuscript.

Informed Consent Statement: Not applicable.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

Conflict of Interest Statement: The authors declare no conflicts of interest.

REFERENCE

- [1] S. Elias, P. Khosla, and T. Nguyen, "Rethinking literary creativity in the age of artificial intelligence," *Humanities and Social Sciences Communications*, vol. 12, no. 1, pp. 1-12, 2025.
- [2] M. Eslami, F. Rahimi, and L. Chen, "Generative AI and the transformation of student writing practices," *Education Sciences*, vol. 16, no. 1, pp. 1-15, 2025.
- [3] L. Henrickson and L. Leppänen, "Revisiting computer authorship: A longitudinal perspective on reader perception," *AI & Society*, early access, 2026.

- [4] S. Jenner, D. Raidos, E. Anderson, S. Fleetwood, B. Ainsworth, K. Fox, J. Kreppner, and M. Barker, "Using large language models for narrative analysis: Opportunities and limitations," *Journal of Computational Humanities*, early access, 2025.
- [5] N. Bostrom, "In defense of posthuman dignity," in H+/-: *Transhumanism and Its Critics*, G. R. Hansell and W. Grassie, Eds. Metanexus Institute, 2010, pp. 55-66.
- [6] N. K. Hayles, *How We Became Posthuman: Virtual Bodies in Cybernetics, Literature, and Informatics*. University of Chicago Press, 1999.
- [7] J. Huxley, *New Bottles for New Wine*. Chatto & Windus, 1957.
- [8] F. Moretti, *Distant Reading*. Verso, 2013.
- [9] Booksby.ai, *AI-generated science fiction novels*. [Online]. Available: <https://booksby.ai>.