



The Author as Machine, the Ghost as Reader: Calvino's *If on a Winter's Night a Traveller* and the Prehistory of Large Language Models

Dr. Naila Anjum

Associate Professor,

English,

Bharati College (University of Delhi), Delhi, India.

naila.anjum@bharati.du.ac.in

 <https://orcid.org/0009-0004-4326-2040>

<https://doi.org/10.66376/criterion.v17.n4.8>

Abstract:

This paper examines Italo Calvino's vision of the "literature machine" as recorded in his 1967 "Cybernetics and Ghosts" and its fictional realisation in his 1979 novel *If on a Winter's Night a Traveller*. Nearly half a century before the advent of large language models, Calvino conceptualised writing as a combinatorial process governed by finite rules yet capable of generating infinite narrative variations. Through an analysis of the novel's interrupted beginnings, its counterfeiter figure Ermes Marana, and its framing of reading as a fundamental human act, the paper demonstrates how Calvino anticipated key features of contemporary AI text generation: the dissipation of the author into a procedural system, the treatment of existing literature as a training corpus, and the productive tension between automated permutation and the "ghost" of unconscious meaning. At the same time, Calvino insisted that no machine can replace the desiring reader, an insight that remains pertinent to debates about artificial intelligence and creativity. By reading Calvino's theoretical writings alongside his narrative practice, the paper argues that he was not merely ahead of his time but had already mapped the conceptual terrain on which current discussions of AI and literature still unfold.

Keywords: Italo Calvino, Barthes, Literature Machine, AI, LLM, *If on a Winter's Night a Traveller*.

Article History: Submitted-13/06/2026, Revised-27/07/2026, Accepted-03/08/2026, Published-31/08/2026.

Copyright vests with the author. Licensing: Distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>)

Introduction

Soon after Granata published the regional winners of the Commonwealth Short Story Prize in May 2026, controversy erupted over the Trinidadian writer Jamir Nazir's short story "The Serpent in the Grove" as being entirely AI-generated. Pangram flagged stories by two of the five winners-- Sharon Aruparayil and Malta's John Edward DeMicoli (Chakraborty). Similarly, Polish Nobel laureate Olga Tokarczuk ignited controversy after her remarks in Poznan were widely interpreted as an admission that she uses artificial intelligence in her writing process. She later issued a clarification, stating that her forthcoming novel was not written with AI and that she mainly uses the technology for preliminary research and fact-checking, not as a co-author (Temple). Such debates may lead us to wonder if the literature machine that can take AI-generated tales to prestigious award lists is finally here.

In November 1967, Italo Calvino stood before an audience in Turin and made a claim that sounded preposterous. He declared that literature could be understood as a machine. The writer, far from being a vessel of inspired genius, was "already a writing machine" who followed rules that could be extracted, formalised, and eventually programmed into a mechanical device. "Nothing prevents us from foreseeing a literature machine," he wrote, adding with characteristic calm that his own place "could perfectly well be occupied by a mechanical device" (Calvino "Cybernetics" 16, 15). However, Calvino was never a pure mechanist. The same essay that dismantles the cult of originality insists on something that no machine can supply alone. Myth, the unconscious, the unsayable: these are "the hidden part of every story", the region "still unexplored because there are as yet no words to enable us to get there" (Calvino "Cybernetics" 18). The combinatorial game, he argues, becomes poetry only when it inadvertently stumbles upon what the conscious mind could not have deliberately produced. For Calvino, literature lives precisely in the space between

the finite and the infinite, the rule and the rupture, the machine and the ghost. Twelve years later, the intricate and whimsical representation of this dual perspective, both computational and profoundly human, was most fully realised in *If on a Winter's Night a Traveller*.

The novel, inter alia, is a working model of a literature machine in action. The novel opens by addressing “you”, the Reader, who has just bought Calvino’s latest book. You find a comfortable position, adjust the light, and begin. A few dozen pages in, the story simply breaks off: a printing error has bound your copy to another novel. You return to the shop, exchange the defective volume, and restart the process. Only, to be interrupted again. Ten novels follow in this way, each with its own atmosphere and genre, each surfacing and then vanishing at the very height of its suspense. Threaded through these fragments is the Reader’s double pursuit: of the missing continuations, and of Ludmilla, the Other Reader, a woman who reads “the way some animals dig holes or build anthills” (Calvino *Night* 189). Along the way we encounter Ernes Marana, a translator who manufactures counterfeit novels, passes plagiarisms off as originals, and floods the world with apocrypha.

In retrospect, he is an astonishingly precise figure for the large language model: a system fed on the corpus of existing literature, capable of generating plausible continuations, and fundamentally indifferent to the distinction between the authentic and counterfeit.

The Italian critic Anna Botta has argued that Calvino’s engagement with the Oulipo group of mathematicians and writers who experimented with formal constraints was not a betrayal of his “Italianness” but rather a deepening of his lifelong interest in combinatorial poetics (Botta 84-85). What Botta calls the “Italian ghost in the combinatory machine” is precisely the element that Calvino refused to automate: the historical, embodied, desiring presence of the reader. This study takes Botta’s insight as a starting point and extends it into the contemporary moment. I argue that

Calvino's theoretical writings on the literature machine and his practical demonstration in *If on a Winter's Night a Traveller* constitute a prescient prototype for understanding large language models. Calvino did not predict specific technologies, but he grasped the underlying principles that make them conceivable: the discrete, combinatorial nature of language; the generation of infinite variations from finite rules; the displacement of the author by a procedural system; and the paradoxical persistence of the reader as the site where meaning is finally secured.

In "Cybernetics and Ghosts," Calvino showed how he synthesised information theory, structuralist narratology, and combinatorial poetics into a vision of writing as a rule-governed generative system. A close reading of *If on a Winter's Night a Traveller* shows the literary manifestation of that vision, focusing on the ten embedded novels as generated outputs, Marana as the proto-LLM counterfeiter, and the Oulipian constraints that organise the novel's labyrinthine architecture. Systematic parallels can be drawn between Calvino's model and the operation of modern large language models to examine the role of training data, token prediction, and hallucination. The conclusion reflects what Calvino's premature vision offers to contemporary debates about AI and creativity. It is not a prophecy but a framework, a way to think about the relationship between the combinatorial and the ghostly without falling into either technoutopianism or Luddite despair.

Literature Survey

Scholarship on Italo Calvino's engagement with computational and combinatorial approaches to literature spans literary theory, digital humanities, and AI studies, and the conversation among these fields is unusually rich. It returns again and again to his 1967 essay "Cybernetics and Ghosts", in which he reframes literature as a combinatorial system governed by formal rules, and

in doing so, anticipates the mechanisation of writing itself. Cronin (2017) and Berkman (2017, 2020) follow Calvino's movement from the mythopoetic narratives of *Cosmicomics* to a rigorous algorithmic poetics shaped by structuralism and semiotics. His years in the intellectual ferment of the 1970s Paris, and especially his ties to the Oulipo group, hardened this commitment to formal constraint as a generative device, as Botta (1997) shows.

His early experiments with machine-assisted narratives make this point concretely. Working with IBM researcher Bruce Skyvington on "The Burning of the Abominable House" (1968), Calvino set out to map the plot permutations available to four characters across twelve situations entirely by computational means (Cronin 2017). The attempt was genuinely pioneering, but it also exposed the limits of the method: the combinatorial logic kept producing outcomes that were illogical or simply lifeless, and Calvino came to see that literary vitality depended on the *clinamen*, the creative swerve or deviation that no rulebook could supply (Duncan 2012; Cronin 2017).

More recent scholarship treats this body of work as a forerunner of the large language model (LLM). Meier-Vieracker (2026) and Franceschelli et al. (2023) draw the resemblance between Calvino's Literature Machine and today's transformer-based architectures, which generate text through probabilistic token prediction across enormous corpora. However, the same comparison exposes a deep divergence. Where Oulipo bends its combinatorial constraints precisely to manufacture surprise, LLMs optimise for statistical predictability, minimising entropy and, with it, the room for genuine creativity (Berkman 2020; Franceschelli et al. 2023).

Ukpaka's philosophical enquiry pushes this analysis further, arguing that LLMs are not creative in the full sense because they lack the first-person experience of the world and the subjective intentionality necessary to be genuine authors (Ukpaka 2024). This is a philosophical position that resonates with Calvino's own view, expressed in "Cybernetics and Ghosts", that the "ghost" is not

a mystical addition to the text but the structural trace of a consciousness that has a relationship to the world—a history of encounters, desires and losses that exceeds its textual record. In this sense, the LLM's training data is all it has; a human author's writing is a selection from and transformation of a life that the writing can only partially render.

Stylometric analysis sharpens this divide. Jemama and Kumar (2025) show that even when LLMs convincingly imitate surface features of style, their output carries measurably lower textual entropy than human writing—the signature of a constrained, homogenised voice. Mikros (2025) and Alvero et al. (2024) reach the same conclusion from different angles, finding that AI-generated texts cluster in an intermediate stylistic space and tend to erase individual and demographic variation, reinforcing dominant linguistic norms rather than widening the range of expressions.

The pattern holds at the level of story. Rettberg and Wigers (2025) document a marked drift toward formulaic narrative structures in AI-generated fiction. They call it the “synthetic imaginary,” a tendency to preserve cultural stability rather than imagine change. The same flattening appears in collaborative settings: Anderson et al. (2024) and Wenger and Kenett (2026) find that LLMs narrow semantic diversity across the users who work with them, potentially dampening divergent thinking.

For all this critical attention, Calvino's *If on a Winter's Night a Traveller* remains surprisingly underexplored as a working prototype of the literature machine, particularly in relation to how AI actually generates text. No existing reading fully draws together the novel's fragmented structure, Ermes Marana as a figure for LLM-style counterfeiting, and the contrasting roles of its readers Ludmilla and Lotaria as embodiments of human engagement set against mechanised evaluation (Botta 1997; Calvino 1979).

This omission invites a fresh reading of Calvino through the lens of contemporary AI, foregrounding the enduring tension between combinatorial automation and the “ghostly” human dimension of reading and desire. Such a reading corroborates with Calvino’s own later conviction that literature draws its vitality not from mechanical generation alone but from the embodied, interpretive work of readers who keep surprise and meaning alive beyond anything an algorithm can predict (Duncan 2012; Botta 1997).

Research Gap

Existing scholarship has firmly established Calvino’s theoretical engagement with cybernetics, his Oulipian commitment to combinatorial constraint, and the formal machinery of his late novels.

Scholars such as Cronin, Berkman, and Botta have convincingly traced how Calvino’s relocation to Paris and his immersion in structuralism and semiotics led him to conceive writing as a rule-governed generative system. Moreover, recent computational and stylometric studies have illustrated that modern large language models operate through probabilistic optimisation that minimises textual entropy, producing a predictable “synthetic imaginary” that stands in sharp contrast to the Oulipian ideal of the *clinamen*—the intentional, creative deflection from constraint that generates genuine artistic surprise.

However, a significant lacuna remains. While critics have acknowledged *If on a Winter’s Night a Traveller* as a key example of Calvino’s combinatorial poetics, no sustained analysis has been undertaken that reads this novel specifically as a practical prototype of the literature machine whose operations can be directly compared to those of contemporary LLMs. Existing scholarship tends to treat Calvino’s theoretical essays and his fictional practice as separate domains, or to analyse the novel’s formal architecture without connecting it to specific behaviours such as

hallucination, apocryphal generation, reader positioning, and narrative interruption, which now define AI text generation.

Also, while scholars like Rettberg and Wigers have identified the “synthetic imaginary” as a flattening force in AI-generated narratives, no research has asked whether Calvino’s novel offers an implicit counter-strategy a way of preserving the ghost of the unsayable within the combinatorial machine. The novel’s conclusion, in which the Reader and Ludmilla marry and continue reading side by side, suggests that Calvino located the irreplaceable human element not in authorship but in the desiring, embodied act of reading itself. This insight remains absent from current debates about AI and creativity, which tend to focus either on the machine’s output or on the author’s displacement, while neglecting the reader’s persistent agency.

This study addresses these gaps by undertaking a close, comparative reading of *If on a Winter’s Night a Traveller* alongside the operational characteristics of modern large language models. In doing so, this paper bridges Calvino scholarship and digital humanities criticism, offering a framework for understanding what is achieved and what is compromised when the ghost gives way to the algorithm.

Methodology

This study employs a comparative close-reading methodology situated at the intersection of Calvino scholarship and digital humanities criticism. The analysis is organised around four figures from the novel who embody different aspects of text generation and reception: Silas Flannery (the blocked writer), Ermenegildo Zegna (the counterfeiting translator), the readers Ludmilla and Lotaria (competing models of literary attention), and Mr Cavedagna (the publishing employee)

Silas Flannery: The Author Who Loses Authority

Flannery appears by Chapter Eight, but his shadow falls across the entire novel. He is the writer behind *In a Network of Lines that Enlace* and the writer whose spiritual crisis delays Marana's translations. His diary entries reveal a man who has stopped believing in his own originality. He watches a woman reading in a deck chair and imagines that she is reading his true book, the one he will never manage to write. "I would like to be able to write a book that is only an incipit", he confesses (Calvino *Night* 193). This is precisely what the larger novel has already done: it has generated ten beginnings and stopped. Flannery wants to be the machine but lacks its indifference. He suffers from jealousy, longing and inadequacy.

When Lotaria explains how she uses a computer to read novels by listing word frequencies, Flannery is appalled. "Every time I write a word, I see it spun around by the electronic brain, ranked according to its frequency" (Calvino *Night* 188-189). His fear is that the machine's reading will reduce his labour to a statistical artefact. However, Japanese counterfeiters have already replicated his style. Flannery responds by fantasising about collaborating with Marana to "flood the world with apocrypha" (Calvino *Night* 193). He is a writer who has accepted that his name no longer guarantees authenticity.

Flannery represents the modern realisation that an author's "intentionality" is often just a sophisticated recycling of existing linguistic patterns. This shows that the author is no longer the sovereign source of a message but a repository of stylistic habits. In AI terms, Flannery represents the loss of first-person subjective experience. His crisis mirrors the "dearth of the author" in LLMs, where text is generated not to say something, but because it is statistically probable. His desire to write a book that is "only an incipit" reveals an intentionality that has been reduced to pure potentiality without the will to conclude.

To comprehend Calvino's literary and theoretical outlook, it's important to consider his connection with the Oulipo, the Workshop of Potential Literature, founded in Paris in 1960 by Raymond Queneau and Francois Le Lionnais. This group proposed a straightforward yet powerful idea: literature could be created through strict formal constraints, and rather than stifling creativity, these constraints would reveal new possibilities that freeform writing might miss. Although Calvino officially became a foreign member in February 1973, Anna Botta's archival work shows that his intellectual affinity with Oulipo's aims had begun years earlier (Botta 82). Marcel Benabou described the group's recognition of Calvino as a kind of "pre-established harmony" between his creative approach and their combinatorial style. This harmony extended beyond style. Calvino had been deeply influenced by Roland Barthes's seminars at the Ecole des Hautes Etudes and A. J. Greimas's structural semantics lectures, which framed literary texts as systems of structured differences rather than mere expressions of an author's inner voice (Botta 1997). Within this context, Calvino came to see narrative not as a personal expression but as the functioning of a code.

Dennis Duncan has detailed the two main combinatorial models Calvino explored during his time in Paris. The first, inspired by the medieval thinker Ramon Llull, views combinatorial literature as the exhaustive and systematic creation of every possible combination within a finite set—a closed, all-encompassing system. The second, linked to the Atomist philosopher Lucretius, imagines literary creation as the random merging of limited elements to produce endless, open-ended variations (Duncan 2012). Duncan suggests that between 1967 and 1973, Calvino shifted from the Lullian to the Lucretian model, embracing the idea of literary constraints not as restrictive cages but as what K. A. Vinoj calls a "liberating paradox", a structural framework that transforms the author from a "romantic expresser" into a "structural engineer" (Vinoj et al. 2025).

Susie Cronin's research into Calvino's exchanges with programmers adds a valuable biographical layer to this theoretical picture. In 1977, Calvino corresponded with William Skyvington, an ex-IBM programmer, about a possible computational project. Skyvington was sceptical, and pressed the practical limits of computational text generation. Calvino appears to have found these limits illuminating rather than discouraging (Cronin 2017). That exchange, between a writer pushing constraint-based composition as far as it would go and a programmer who knew exactly where the machinery gave out, captures the kind of dialogue that lies behind the Organization for the Electronic Production of Homogenized Literary Works (OEPHLW): a fictional apparatus convincing enough to feel plausible, even threatening, yet whose very operation throws into relief what remains stubbornly human. Ernes Marana in the novel introduces himself as a representative of the OEPHLW to the blocked author Silas Flannery.

Calvino's Oulipian method, along with Natalie Berkman's empirical research in digital humanities, reveals that the group's algorithmic techniques fundamentally resemble computer algorithms: both rely on the systematic, rule-based manipulation of a limited set of elements (Berkman 2017). This idea of intentional deviation is key to distinguishing Calvino's fictional machine from a genuine author, and by extension, to differentiating the probabilistic outputs of large language models from true literary creativity.

OEPHLW is more than just a satirical narrative element in the novel. It serves as a fully developed, albeit implicit, theory of literary style that engages in a nuanced and productive dialogue with Roland Barthes's contemporary ideas on authorship and textuality. In his 1967 essay "The Death of the Author", Barthes argued that the concept of the Author as a biographical, intentional source of a text's meaning is a modern, bourgeois invention that critical analysis should set aside. For Barthes, writing represents "the destruction of every voice, of every point of origin"; the author

“enters into his own death” at the moment of writing, since writing functions as a system of differences without a fixed origin or ultimate purpose (Barthes 49). Meaning is not implanted by an authorial consciousness for the reader to uncover; rather, it emerges through the reader’s engagement, arising from the interplay of codes beyond any single author’s control. Calvino’s OEPLW enacts a literal and importantly distorted version of this theory. The machine works by identifying the stylistic and conceptual patterns of a living author in the novel, *Silas Flannery* and uses these patterns to generate new text. Here, the author is not metaphorically “dead” as Barthes suggests, but is diminished: reduced to a data set, a statistical profile, a collection of habits and preferences that a machine can learn and reproduce. This distinction carries real philosophical weight. Barthes’s “death of the author” is an epistemological claim about where meaning resides; the OEPLW makes an ontological one, insisting that the author is nothing more than her data and that style exhausts personhood. Calvino takes up this proposition only to refine it, and he does so through the novel’s treatment of the reader. The “You” addressed directly in the second person is no passive recipient of meaning but an active, yearning, and irreducibly individual consciousness. This is precisely the one thing that the machine cannot reproduce.

When the Reader encounters a text produced by the OEPLW, there is an undeniable absence, even if it is not immediately recognisable. Although the machine achieves “perfect fidelity” in replicating the surface style, the experience still feels artificial. This discrepancy between stylistic imitation and genuine literary presence is the “ghost” Calvino describes in “Cybernetics and Ghosts”: the reader’s consciousness, which contributes an experiential background, emotional depth, and interpretive unpredictability that no machine can fully replicate. The OEPLW’s approach to style is implicitly distributional, viewing style as an aggregate of an author’s measurable textual tendencies. It records their sentence lengths, grammatical choices, preferred

metaphors, vocabulary, and structural forms. This resembles the way modern Natural Language Processing (NLP) researchers define style for imitation and attribution. However, what the fictional machine fails to capture are the repressed emotions, traumas, taboos and the surprisingly varied ways in which they express their thoughts.

Stylometric Fidelity: Surface-Level Success

The most straightforward evaluation of the OEPHLW's claim to replicate an author's style "with perfect fidelity" is reflected in recent stylometric analyses of large language model (LLM)-generated text. Mikros' study makes the same point in empirical terms. Mikros tested GPT-4o's ability to imitate the prose of Ernest Hemingway and Mary Shelley across three different prompting methods, drawing on distance-based authorship attribution: the 1,000 most frequent words, cosine distance metrics, and random forest classifiers. The model captured some surface features of each style, but consistently failed to reproduce the full complexity of an author's stylometric profile. Imitations generated through in-context learning, where the model is shown examples of the target style before it writes, are more closely aligned with the original, yet even these retain a heavy overlap with what Mikros calls "generic GPT outputs"—the model's own intrinsic voice surfacing beneath the disguise (Mikros 2025). There is a quiet irony here that Calvino would have relished. In the novel, Ludmilla's sister Lotaria reads by feeding books into a machine that prints out the frequency of their words, certain that the tally lays bare a text's meaning faster than reading ever could. Mikros's classifiers, which sort authors by their thousand most common words, are Lotaria's method made operational, and they arrive at the very conclusion Calvino reached in fiction: the frequency list catches the surface and misses the voice.

This outcome closely mirrors Calvino's fictional scenarios. The OEPLW attains surface-level fidelity; the generated prose sounds like Silas Flannery but cannot entirely mask the machine's inherent logic, which emerges as a kind of stylistic genericness, smoothing out the distinctive quirks of the author's style. Similarly, Jemama et al.'s framework-based evaluation of five LLMs across four prompting strategies reaches a comparable conclusion from a different perspective: high stylometric fidelity does not guarantee human-like unpredictability. Their study found that human essays averaged a perplexity of 29.5, whereas LLM-generated texts with matched style averaged only 15.2 (Jemama et al. 2025). In NLP, perplexity measures how surprising a language model finds a sequence of text; lower perplexity indicates higher statistical predictability. Thus, human writing is inherently less predictable than even the most accurate LLM imitations.

This statistical unpredictability in human writing corresponds, in theoretical terms, to the *clinamen* Calvino's Oulipian framework identifies as the hallmark of authentic literary creation: a meaningful deviation from established patterns that a purely distributional model cannot produce because it lacks any basis for choosing to deviate. It represents the structural imprint of a consciousness that cannot be reduced to textual antecedents. While surface-level stylistic imitation is achievable, the deeper concern highlighted by Calvino's novel is the tendency of large language models to drive homogenisation: a systematic narrowing of literary diversity toward a dominant, statistically prevalent mode. Wenger et al.'s study in PNAS Nexus directly examined this by eliciting creative outputs from numerous humans and LLMs through standardised creativity tasks and comparing diversity across populations. Their results were clear: LLM-generated responses resemble each other far more closely than human responses do, even after accounting for key confounding factors (Wenger et al. 2026). This effect goes beyond individual model style traits; it

represents a population-level convergence and homogenisation of the overall creative output landscape.

Similarly, Anderson et al.'s research on LLMs as creativity support tools revealed that users collaborating with ChatGPT produced ideas that were less semantically varied than those using an alternative tool, indicating homogenisation at both the individual and aggregate creative levels. Sourati et al.'s review in *Trends in Cognitive Sciences* integrates findings from linguistics, psychology, and computer science to argue that LLMs both reflect and amplify dominant styles while marginalising alternative voices. They warn that widespread LLM use risks “flattening the cognitive landscapes that drive collective intelligence and adaptability” (Sourati et al. 2025).

In Calvino's novel, the OEPHLW is explicitly dedicated to producing “homogenised” literature. The term appears in its very name. The homogenisation it achieves is twofold: smoothing out individual authorial idiosyncrasies within each text and converging all generated texts across the literary field toward what sells, what patterns, and what statistically dominates the training data. Calvino identifies this as an intentional organisational aim rather than an unintended side effect. This reflects the market logic of literature in reducing writing to a commodity expressed through the machine’s distributional logic. Contemporary research suggests that LLMs, even without such organisational goals, reach a similar convergence because their training on massive corpora is weighted towards privileges dominant over marginal, central over peripheral.

Despite the standardised, homogenised answers that AI produces, its continuing popularity shows the cognitive ease that is missing in more complex, nuanced works produced by litterateurs. Most humans prefer easy-to-decipher texts over literature with complex metaphors and layered interpretations. The language in which the LLMs are trained is that of a university-educated, privileged, and dominant Western male. These factors make the machine output predictable. It is

feared that too much exposure to AI may kill our creativity by internalising a homogenised and sanitised vocabulary.

Ermes Marana: The Counterfeiter as Generative Model

Marana is the most direct figure for the large language model in Calvino's novel *Winter's Night*. He produces false translations, versions that have no original, and continuations that break off before reaching an ending. When Mr Cavedagna confronts him, Marana replies with a letter asking, "What does the name of an author on the jacket matter?" (Calvino *Night* 101). He points out that some great works are anonymous, that some famous authors have left no surviving works, and that Homer may name a collective rather than a person (Calvino *Night* 101). He concludes that authorship is a temporary convenience.

Marana's letters from around the world describe how he inserts himself into existing networks of literary production and deranges them. In Arabia, he translates novels for a Sultana, whose husband has banned books. To keep her from finishing one novel and having nothing to read, he breaks off each translation at the moment of greatest suspense and starts another novel within it (Calvino *Night* 125). The result is a regressus ad infinitum of interrupted narratives, exactly the structure of *If on a Winter's Night a Traveller* itself.

Marana also founds two secret societies, the Wing of Light and the Wing of Shadow, each hunting apocryphal books for opposite reasons. The Wing of Light believes that among false books, a few contain genuine truth. The Wing of Shadow believes that only counterfeiting can represent absolute value (Calvino *Night* 130-131). Marana plays both sides. His real motive is personal: he resents Ludmilla's reading and wants to make her doubt every word she reads. An LLM does not

experience such jealousy. It generates text because it has been trained to minimise prediction errors. This is both less interesting and, in some ways, more insidious.

The Readers: Ludmilla, Lotaria, and the Pursuit of Surprise

The Reader (the “you” of the frame narrative) wants closure. His frustration drives the plot, but he is also a limited character. The narrator notes that keeping him unnamed allows any reader to identify with him, yet this universality is an illusion (Calvino *Night* 141). The Reader is a specific character with specific blind spots, including a tendency to confuse his desire for Ludmilla with his desire to finish reading books.

Ludmilla reads not for closure but for the experience of being carried along the narrative. She tells Flannery that she likes novels which “create an illusion of transparency around a knot of human relationships as obscure, cruel, and perverse as possible” (Calvino *Night* 192). When two professors argue over whether a text is Cimmerian or Cimbrian, Ludmilla ignores them; she only cares whether the story can continue (Calvino *Night* 74). Ludmilla represents the site where intentionality is ransomed. If the machine has no intent, the human reader must project it to make the text art. As the desiring user, Ludmilla’s role is labour to maintain the possibility of surprise. Her intentionality is directed toward experience rather than information.

Calvino argues that the “poetic result” of a machine is the effect it has on a “man endowed with a consciousness and an unconscious” (Calvino “Cybernetics” 22). Meaning is said to be “dumb” until it chimes with the reader’s own “ocean of the unsayable” (Calvino *Cybernetics* 19) Dumb meaning is a term used in modern linguistics to describe the output of Large Language Models (LLMs), which mimic meaningful choices through statistical probability rather than intent. Meier-Vieracker credits this term to Hannes Bajohr (92). In modern AI, the user’s prompt is the only

source of intent. Ludmilla is the Other Reader who leans forward to peer beyond the printed page, looking for a voice that comes from the unsaid. She proves that intentionality resides in the destination of the text, not its origin. Her sister Lotaria accuses her of passive escapism, but Calvino seems to side with Ludmilla. She embodies what the literature machine cannot replace: the embodied, desiring act of reading that transforms words into an inner experience.

Lotaria does not read books; she has a computer read them for her. The machine outputs frequency lists from which she infers a book's themes. She shows Flannery a list of words appearing 19 times in a novel: "blood, cartridge belt, commander, do, have, immediately, it, life, seen, sentry, shots, spider, teeth, together, your" (Calvino *Night* 187). From this, she concludes that it is a war novel. Her method flattens every text into a distribution of tokens, erasing what makes a particular novel worth reading: its timing, voice, and refusal to do what the reader expects. Lotaria's computer is a precursor to the stylometric tools that now detect AI-generated text. Those tools measure frequency and entropy, but they cannot tell us whether a sentence surprises us or changes us.

Mr Cavedagna and the Publishing Machine

Cavedagna works in a publishing house overwhelmed by manuscripts and constantly making errors. The bindery mixes signatures, the printer sets the wrong title page and the translators deliver forgeries. Cavedagna manages this chaos without losing sight of the fact that books are also just objects. However, he dreams of retiring to his village and reading again, as he did as a boy, hiding in the chicken coop (Calvino *Night* 97). He holds two truths without resolving them: that books are made through a mundane, commercial process and that for readers, books are portals to other worlds. When the Reader complains about defective copies, Cavedagna tells the whole story of Marana's fraud. He is almost relieved to talk to someone who still reads for pleasure. "I come

across fewer and fewer readers”, he says (Calvino *Night* 97). Cavedagna represents the infrastructure that makes the literature machine run. He does not write and barely reads. He keeps the gears turning.

Parallels with Large Language Models

These four figures map onto components of a large language model. Flannery is the training data: a corpus of existing texts that the model has absorbed and can reproduce but no longer control. Marana is the generation algorithm itself: the process of sampling from probability distributions to produce plausible continuations. Ludmilla is the user who prompts the model not for information but for experience. Lotaria is the evaluator who measures outputs against benchmarks, caring about frequency and perplexity rather than affect. Cavedagna is the deployment infrastructure: the servers, the APIs, the content filters.

None of these parallels is perfect, and Calvino insists on their interplay. Marana’s counterfeits would be meaningless without Ludmilla’s desire for authentic surprise. Flannery’s blockage matters because someone expects him to produce. Cavedagna’s exhaustion matters because he still remembers what it was like to read in the chicken coop. Large language models today can mimic genres, imitate authors, and continue sentences plausibly. What they cannot do is want. They cannot be frustrated when a story breaks off. They cannot feel relieved when a sentence turns out differently than expected. Calvino’s literature machine was never designed to replace that feeling. It was designed to make us notice it.

Discussion

The progression of Italo Calvino's thought moves from a 1967 theoretical optimism in a "Literature Machine" to a fictional interrogation of that machine in *If on a Winter's Night a Traveller*, ultimately culminating in the realisation that true literature requires a human "Ghost" to disrupt the mechanical algorithm.

Calvino's actual attempts to collaborate with computer experts such as Paul Braffort and William Skyvington led to a decisive shift in his perspective (Cronin 78, 81). He discovered that literature could not be produced by a machine alone for several reasons. His experiments showed that pure combinatorial logic produced scenarios that were "completely stupid" or logically impossible, such as a character being blackmailed after they had already been poisoned (Cronin 78). Calvino concluded that a computer can only liberate the author from the "slavery of a combinatory search" (Cronin 79; Duncan 104). The actual "work of art" only emerges through the *clinamen*, an intentional "swerve" or error that shatters the mechanical constraint (Cronin 79; Duncan 104).

In his later realisation, the "Ghost" is the author's own repressed material (traumas and taboos) that "speaks" through the act of writing (Duncan 106). Since a machine has no childhood, no trauma, and no "hairy claws scratching in the dark", it is mathematically incapable of the "swerve" that makes literature human (Duncan 108).

Therefore, while the algorithm can map the territory of language, Calvino realised it cannot navigate the human experience, which relies on the "spasms of an I" to give the mechanical output meaning (Botta 87).

The analysis above shows that Calvino's *If on a Winter's Night a Traveller* models a distributed system of text generation in which no single agent-- author, translator, publisher, or reader-- controls the flow of meaning. Instead, meaning emerges from the friction between these roles.

Flannery cannot write without imagining Ludmilla reading over his shoulder. Marana cannot counterfeit without the publishing industry's hunger for new titles. The Reader cannot finish a book because the bindery has mixed up the signatures. None of these failures is merely a mistake. These are the conditions under which the literature machine produces its effects.

What makes this model relevant to contemporary large language models is the way it disperses agency. When a human writes a sentence, they decide what comes next based on their intention, memory, and sense of the whole. When an LLM writes a sentence, it computes the probability distribution of the next token based on the preceding context. The result can look similar, but the underlying process is different. However, Calvino's novel suggests that human writing has always been more distributed than we would like to admit. The romantic figure of the solitary genius is a late invention and not a very useful one. Writers borrow, misremember, imitate, and sometimes counterfeit without knowing it. The literature machine simply makes this visible.

One of the novel's most unsettling insights is that readers, too, contribute to the machine's operation. Ludmilla's open-ended curiosity keeps the cycle of interrupted narratives going. If she demanded closure, if she refused to start another book each time the previous one broke off, the whole structure would collapse. But she does not refuse. She is willing to be led, again and again, into new beginnings. Her willingness is not passivity. It is a form of labour. She works to maintain the possibility of surprise. While Flannery and Maran represent the training data and algorithm that produce a predictable output, reminiscent of the Lucretian atoms falling in straight, parallel lines, only Ludmilla (User) or a human author using a *clinamen* can provide the swerve that Calvino calls a "true work of art" (Duncan 104-105). Lotaria, by contrast, tries to short-circuit the machine with her frequency lists. She wants to know what a book is without reading it, to extract

its meaning without undergoing its duration. Her method succeeds in its own terms, but it also kills what she claims to study. The computer cannot be surprised, and neither can she.

The publisher Cavedagna occupies a middle position. He handles books as objects, problems to be solved, and inventory to be managed. However, he also remembers a time when reading was different, when he hid in the chicken coop with a cheap edition and coloured the engravings with crayons. That memory does not help him do his job. It only makes the job harder because he knows what is being lost. Cavedagna's predicament is that of anyone working in the culture industry today.

Large language models intensify this tension without resolving it. They are trained on a corpus that includes the entire history of literature, but they have no memory of reading any of it. They can generate a plausible continuation of a novel, but they cannot feel the frustration of having it break off. They can mimic the style of an author, but they cannot want that mimicry to succeed or fail. Calvino's novel does not lament this. It observes it. The ghost in the machine is not the author, who has already vanished. The ghost is the reader who remains.

Conclusion

Italo Calvino did not predict large language models but he did foresee a structural shift in how literature would be produced, circulated, and received. The romantic author, he argued in "Cybernetics and Ghosts," was a "spoiled child of ignorance." The future belonged to a more thoughtful person, someone who knew that the author was a machine and understood how that machine worked. *If on a Winter's Night a Traveller* is that person's report from the future. It shows us a world in which texts are generated combinatorially, distributed through error-prone networks, and judged by readers who cannot agree on what they are looking for.

In the digital age intentionality has shifted from the origin (author) to the destination (reader).

The parallels between the characters in Calvino's novel and LLMs mean that literature is no longer a top-down message but a distributed system where the reader's desire is the final component that redeems the text from being mere statistical "AI slop" (Rettberg and Wigers 14).

The novel's ten interrupted beginnings are not a gimmick. They demonstrate the core operation of the literature machine: generate, interrupt, replace, repeat. The machine does not produce endings because they would stop the process. What keeps it running is the reader's desire for closure, which is perpetually deferred. The Reader never finishes a story, yet he never stops trying. Ludmilla never asks for an ending, yet she never stops reading. Between them, the machine finds its rhythm. Ermes Marana, the counterfeiter, is the machine's most disturbing component because he has no conscience. He produces false texts not out of malice, exactly, but out of a kind of aesthetic nihilism. If everything can be faked, then nothing is sacred, and he can finally stop worrying about whether Ludmilla prefers another author. Marana wants to win an argument, not to make something true. This makes him different from a large language model, which does not want anything at all. It is not malevolent but empty. The danger is not that the machine will lie, but that we will stop caring about whether it is lying.

The novel ends with the Reader and Ludmilla married, lying in bed, each reading a separate book. He says he is almost finished with *If on a Winter's Night a Traveller* by Italo Calvino. The line is comic because it breaks the frame. We are reading that sentence now. He is reading it too. The book has caught up with the world. It ends not with a resolution but with a pause, a breath before the next beginning. This is the only ending that the literature machine can provide.

Calvino's legacy for thinking about AI and creativity is not a set of predictions. It is a set of questions. The novel does not answer these questions but acts them out. This is what literature can do that machines cannot. It can perform its own uncertainties without resolving them. It can hold the ghost and the machine together in the same sentence and let them look at each other.

The broader implication for literary studies is this: Calvino's novel, read against the current landscape of generative AI, suggests that the most important critical category for evaluating AI-generated literature is not quality which is a surface measure but intentionality. The question is whether a text is the product of a consciousness that had reasons for its choices. This is not a mystical criterion. It is a structural one. It is a question that Calvino's fiction, nearly half a century before the age of LLMs, had already learned to ask.

Works Cited:

Alvero, A. J., et al. "Large Language Models, Social Demography, and Hegemony: Comparing Authorship in Human and Synthetic Text." *Journal of Big Data*, vol. 11, article 138, 2024, pp. 1–28. <https://doi.org/10.1186/s40537-024-00986-7>.

Anderson, Barrett R., et al. "Homogenization Effects of Large Language Models on Human Creative Ideation." *Proceedings of the 16th Conference on Creativity & Cognition*, 2024.

Antony, Lima. "The Ethics of AI-Generated Literature: Originality, Authorship, and Literary Value in the Age of Large Language Models." *International Journal of Arts and Liberal Studies*, 2025.

Barthes, Roland. "The Death of the Author." *The Rustle of Language*, translated by Richard Howard, Blackwell, 1986, pp. 49–55.

Berkman, Natalie. "Digital Oulipo: Programming Potential Literature." *Digital Humanities Quarterly*, 2017.

Berkman, Natalie. "Italo Calvino's Oulipian Clinamen." *MLN*, vol. 135, no. 1, 2020, pp. 161–185.

Botta, Anna. "Calvino and the Oulipo: An Italian Ghost in the Combinatory Machine?" *MLN*, vol. 112, no. 1, 1997, pp. 81–89.

Calvino, Italo. "Cybernetics and Ghosts." *The Uses of Literature*, translated by Patrick Creagh, Harcourt Brace, 1986, pp. 3–27.

Calvino, Italo. *If on a Winter's Night a Traveller*. Translated by William Weaver, Harcourt Brace 1979.

Chakrabarty, Tuhin, et al. "Art or Artifice? Large Language Models and the False Promise of Creativity." *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024.

Chakraborty, Abhinav. "AI Detection Tools, Granta, Pangram, and the Writing Controversy, Explained." *The Indian Express*, 1 July 2025, <https://indianexpress.com/article/explained/explained-ai/ai-detection-tools-granta-pangram-writing-controversy-10720880/>. Accessed 2 July 2026.

Cronin, Susie. "Cybernetic Collaborations, Literature Machines and Italo Calvino's 'L'incendio della casa abominevole'." *The Italianist*, vol. 37, no. 3, 2017, pp.69-85.

Duncan, Dennis. "Calvino, Lull, Lucretius: Two Models of Literary Combinatorics." *Comparative Literature*, vol. 64, no. 1, 2012, pp. 93–109.

Alvero, A. J., et al. "Large Language Models, Social Demography, and Hegemony: Comparing Authorship in Human and Synthetic Text." *Journal of Big Data*, vol. 11, article 138, 2024, pp. 1–28. <https://doi.org/10.1186/s40537-024-00986-7>

Jemama, Rebira, et al. "How Well Do LLMs Imitate Human Writing Style?" *Proceedings of the IEEE 16th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference*, 2025.

Meier-Vieracker, Simon. "No Choice: On the Stylistics of AI-Generated Texts." *Journal für Medienlinguistik*, vol. 8, no. 2, 2026, pp. 73–101. <https://doi.org/10.21248/jfml.2026.85> 2026.

Mikros, Georgios. "Beyond the Surface: Stylometric Analysis of GPT-4o's Capacity for Literary Style Imitation." *Digital Scholarship in the Humanities*, 2025.

Rettberg, Jill Walker, and Hermann Wigers. "AI-Generated Stories Favour Stability over Change: Homogeneity and Cultural Stereotyping in Narratives Generated by GPT-4o-mini." *Open Research Europe*, vol. 5, article 202, 2025, pp. 1–17. <https://doi.org/10.12688/openreseurope.20576.1>

Sourati, Zhivar, et al. "The Homogenizing Effect of Large Language Models on Human Expression and Thought." *Trends in Cognitive Sciences*, 2025.

Temple, Emily. "Olga Tokarczuk Has Responded to the Controversy over Her Reputed Use of AI." *Literary Hub*, 19 May 2026, <https://lithub.com/olga-tokarczuk-has-responded-to-the-controversy-over-her-reputed-use-of-ai/>. Accessed 2 July 2026.

Ukpaka, P. "The Creative Agency of Large Language Models: A Philosophical Inquiry." *AI and Ethics*, 2024

Vinoj, K. A., et al. "Calvino's Labyrinths: Constraint as the Postmodern Degree Zero." *RESEARCH REVIEW International Journal of Multidisciplinary*, 2025.

Wenger, Emily, et al. "Large Language Models Are Homogeneously Creative." *PNAS Nexus*, 2026.