SP Publications
International Journal Of English and Studies (IJOES)
An International Peer-Reviewed and Refereed Journal; **Impact Factor:**8.337(SJIF)
ISSN: 2581-8333|**Volume 8, Issue-11(November);2026**

Can an Interpretation Exist Without an Interpreter? Artificial Intelligence and the Problem of Literary Reading

Anmol Deep[1]
Designation: Ph D Research Scholar, Affiliation: Department of English and Foreign Languages, Guru Jambheshwar University of Science and Technology, Hisar, Haryana, India.
Email: anmolnavu77@gmail.com
Dr. Sakshi Jain[2]
Affiliation: Department of English and Foreign Languages, Guru Jambheshwar University of Science and Technology, Hisar, Haryana, India.
Email: dr.sj.mittal@gmail.com

rticle Received: 09/10/2026
Article Accepted:11/11/2026
Published Online:12/11/2026
DOI:10.47311/IJOES.2026.8.11.158

Abstract

A language model can now produce something that looks remarkably like literary criticism: it can notice contradictions, select evidence, discuss form, identify ambiguity, and construct an argument about meaning. But a more difficult question begins where this apparent success ends: if a machine can produce a convincing reading, what—or who—is doing the interpreting? This paper examines the distinction between interpretive adequacy and interpretive agency through AI-generated readings of Robert Frost's "The Road Not Taken" and Toni Morrison's "Recitatif." Drawing on hermeneutics, reader-response theory, interpretive-community theory, and posthumanist approaches to agency, it considers how literary meaning emerges through relations among textual structures, readers, and established interpretive practices. Through a qualitative close analysis of the AI-generated responses, the paper asks whether a reading can be evaluated as textually grounded and critically persuasive without that evaluation establishing the agency of its producer. It argues that interpretive adequacy and interpretive agency require distinct forms of assessment. Rather than resolving the question of whether AI is an interpreter, this distinction clarifies how AI-generated criticism challenges the assumption that interpretation originates in a single, autonomous subject.

Keywords: Literary Interpretation; Large Language Models; Interpretive Agency; Reader-Response Theory; Posthumanism; Artificial Intelligence

SP Publications
International Journal Of English and Studies (IJOES)
An International Peer-Reviewed and Refereed Journal; **Impact Factor:**8.337(SJIF)
ISSN: 2581-8333|**Volume 8, Issue-11(November);2026**

Introduction

A language model can now do something that once seemed distinctly human: it can read a poem or story and produce something that looks remarkably like literary criticism. It can identify patterns, select textual evidence, notice contradictions, discuss form, recognise ambiguity, and build an argument about meaning. The result can be persuasive enough to make a deceptively simple question difficult to answer: if a machine can produce a convincing literary reading, what exactly are we recognising as "interpretation"? Recent research has begun to examine the ability of artificial systems to perform literary interpretation, raising questions about both the reliability of their readings and the status of the systems that produce them (Ichien et al. 296–97). But the more interesting problem is not simply whether an AI can produce a reading that resembles one written by a human critic. It is what happens to the category of interpretation when a nonhuman system can perform many of the operations through which a reading becomes recognisable as interpretation.

Artificial intelligence complicates these assumptions because an AI system can generate a literary interpretation without occupying the conventional position of a human reader. This creates a distinction that is often obscured in discussions of whether AI "understands" literature. **Interpretive adequacy** concerns what a reading accomplishes as a textual and critical product: whether it attends to textual evidence, identifies formal and semantic tensions, recognizes ambiguity, and develops a coherent account of the text. **Interpretive agency**, by contrast, concerns the status of the interpreter to whom those operations can be attributed. A reading may therefore be interpretively adequate without the production of that reading necessarily establishing who or what has interpreted.

The human reader does not provide a simple answer to this problem either. Literary theory has long questioned the idea of the reader as an isolated, autonomous source of meaning. Stanley Fish, for example, shows how interpretation is shaped by the shared practices of interpretive communities. Rosi Braidotti similarly challenges the humanist image of the self-determining individual, arguing for a more relational account of subjectivity grounded in embodiment, affectivity, and relations with others (26–27). These perspectives make the AI problem more complicated: if human interpretation is already shaped by social, material, and cultural relations, then the contrast between an autonomous human reader and a dependent machine is less straightforward than it first appears.

This paper therefore sets aside the question of whether AI is conscious or whether it "really" understands literature. Such questions make an unobservable mental state the final test of interpretation, while leaving the interpretive performance itself largely unexplained. The more productive question is what sophisticated AI-generated readings reveal about interpretation as a category. To explore this question, the paper turns to two texts that make

different dimensions of reading unusually visible: Robert Frost's "The Road Not Taken" and Toni Morrison's "Recitatif." Frost exposes the tension between contradiction, ambiguity, form, and retrospective narration; Morrison makes the reader's role in producing and revising meaning visible through textual gaps, withheld racial identification, competing memories, and uncertainty surrounding Maggie.

The paper begins by separating two questions that are often treated as one: whether an AI system can produce the kind of reading we recognise as interpretation, and whether producing such a reading makes the system an interpreter. The distinction is then tested through close readings of Frost and Morrison, before the discussion turns to a relational account of interpretation informed by posthumanist theory. The argument is deliberately neither a defence nor a dismissal of AI interpretation. The point is more unsettling: once a machine can produce a reading that is textually sophisticated and critically persuasive, we are forced to reconsider whether interpretation ever belonged to a single, autonomous interpreter in the first place.

1. What Makes an Interpretation?

What, exactly, do we require before we call something an interpretation? Literary theory offers no single answer. Some accounts locate meaning in authorial intention or textual structure; others emphasise history, the reader, or the communities through which meaning becomes intelligible. Rather than choosing among these positions, this section asks what each makes visible about the problem of AI interpretation.

E. D. Hirsch's distinction between meaning and significance provides an important starting point. In *Validity in Interpretation*, Hirsch distinguishes an author's intended meaning from the significance that meaning may acquire for later readers. His account preserves a degree of stability in verbal meaning while allowing interpretation to generate historically changing significance (Hirsch 8–9). For AI, the distinction creates an immediate difficulty: a language model may produce a textually defensible significance without having access to authorial intention. The question is therefore not simply whether its reading is intelligible, but whether producing an intelligible and defensible significance is enough to count as interpretation—or whether interpretation also requires an identifiable interpreter.

In his work *Truth and Method*, Gadamer claims understanding is never detached from history and tradition: the interpreter approaches a text from within a horizon shaped by prior experience, and understanding emerges through the encounter between the interpreter's horizon and that of the text (Gadamer 299, 305). Interpretation therefore cannot be separated from the situation in which it occurs (398). If such situatedness is constitutive of understanding, an AI response cannot establish its own interpretive agency simply by producing a successful reading. Yet neither can the absence of conventional human

situatedness simply be treated as proof that no interpretation has taken place, because the output itself may participate in historically inherited forms of interpretation.

Iser shifts the problem from what a text contains to what happens when someone reads it. In *The Act of Reading*, the literary text is not a finished store of meaning but a potential effect realised through the reader's activity; meaning emerges through the interaction between textual structures and acts of comprehension (ix–x). His account of textual blanks makes this especially clear: indeterminacies invite readers to connect patterns, fill gaps, and form relationships that the text does not explicitly provide (Iser 182–83). Jauss widens the frame by asking us to also consider the history of those encounters. His aesthetics of reception locates a work's historical life in the reciprocal relationship between the work and its recipients, shifting attention from the text alone to the horizons through which it is received and interpreted (15–16). An AI-generated reading enters this history rather than standing outside it: its language, assumptions, and interpretive moves are drawn from a much larger archive of previous readings. The machine therefore does not simply produce a reading of a text; it produces one from within a history of reading.

Fish makes the problem of the autonomous interpreter even harder to sustain. His account of interpretive communities challenges the idea that meaning originates in an isolated reader confronting a self-contained text. What readers recognise as meaningful depends on socially established assumptions and practices that make an utterance intelligible in the first place (Fish 194–95). Those same practices shape the standards by which interpretations are understood and evaluated, even though such standards remain embedded in their social contexts (Fish 202–03). This matters for AI because a language model does not invent the conventions through which its readings become recognisable as literary criticism. It draws on an enormous inheritance of human language, critical vocabulary, and interpretive practice. Yet this does not make the resulting reading equivalent to a human interpretation. It shifts the question: if the conventions that make an interpretation possible are already socially produced, how much of the interpretive act can legitimately be attributed to the system that reproduces and recombines them?

Posthumanist theory makes this difficulty productive. Hayles's work on embodiment and the posthuman challenges the assumption that information can be separated from material and technological processes, while Braidotti's posthumanism questions the universal human subject as the unquestioned centre of knowledge and agency. From this perspective, AI should not merely be asked to prove that it possesses the same qualities traditionally associated with the human. Instead, its emergence can be used to examine the assumptions through which those qualities have been defined in the first place.

The discussion so far leaves us with two questions that should not be confused. First: can an AI system perform the operations we recognise as literary interpretation—selecting evidence, tracing patterns, identifying ambiguity, and constructing an argument? Second: if it can, does that performance give us sufficient grounds to call the system an interpreter? A sophisticated reading can answer the first question without settling the second.

This distinction also reflects how literary criticism actually works. We judge interpretations by what they do with a text: what evidence they select, how they connect it, whether they account for contradictions and formal features, and whether their argument illuminates something that might otherwise remain unnoticed. We do not directly observe an interpreter's agency; we infer it from the relation between the interpreting subject and the interpretation produced. The case studies that follow therefore evaluate the AI responses first as critical texts—before asking what, if anything, their success tells us about the entity that produced them.

2. AI and Interpretive Performance

To test the distinction between interpretive performance and interpretive agency, I examine two LLM-generated close readings: one of Robert Frost's "The Road Not Taken" and one of Toni Morrison's "Recitatif." Both responses were generated using GPT-5.6 Luna on 1 September 2026. The literary texts themselves were not supplied; the model received the titles together with the prompts reproduced in the Appendix. One response was generated for each work, and the original outputs were analysed without editing. The two readings are treated as qualitative case studies rather than as a representative sample of AI-generated literary criticism. The analysis therefore asks a deliberately limited question: how adequately can these particular responses perform literary interpretation, and what—if anything—does that performance allow us to infer about interpretive agency? Because the literary texts were not supplied in full, the study does not independently control for the model's prior textual knowledge or establish the provenance of each interpretive observation.

This distinction is also important in light of Mitchell and Krakauer's warning that fluent performance does not, by itself, establish humanlike understanding, and that methods designed to assess human cognition may not adequately evaluate AI systems (2–4). The present study therefore does not treat fluency as evidence of understanding. Instead, it evaluates the readings through the criteria that literary criticism ordinarily applies to interpretation: textual evidence, formal analysis, attention to ambiguity, and the ability to develop a coherent argument.

3. Frost: Interpretation as Performance

Robert Frost's "The Road Not Taken" is an especially useful text for testing interpretive performance because its apparent simplicity has generated an influential but

unstable cultural interpretation. As Katherine Robinson observes, the poem has often been read as a celebration of individualism, although its language complicates that interpretation. The speaker must choose between two roads without establishing a meaningful difference between them, but later imagines narrating the choice as distinctive. The contradiction is not incidental. It is built into the poem's language.

The AI's close attention to the poem's diction demonstrates a substantial degree of interpretive adequacy. It identifies the significance of "just as fair," "perhaps the better claim" ("Road Not Taken" 7), and "really about the same" ("Road Not Taken" 10), reading these qualifications as evidence of the speaker's continual revision of his own account. The assertion that "both that morning equally lay" ("Road Not Taken" 11) becomes particularly consequential because it sits in tension with the later claim that one road was "less traveled." Rather than resolving this contradiction, the response makes it constitutive of the poem's meaning.

The response similarly recognises that the poem's uncertainty operates at both semantic and formal levels. Its attention to the four five-line stanzas, the ABAAB rhyme scheme, and the relation between formal regularity and conceptual instability culminates in the formulation, "The form is stable; the thought is not." The observation is significant because it converts formal description into an interpretive claim: the poem's formal coherence provides a counterpoint to the instability of the meaning retrospectively assigned to the choice.

The same is true of its attention to syntax and qualification. Words such as "perhaps," "though," "really," "yet," and "I doubted" are treated as signals through which the speaker continually qualifies or revises what he has said. The first-person voice likewise does not automatically guarantee psychological certainty. The AI distinguishes between the fact that the speaker makes a choice and the uncertainty surrounding the meaning later attributed to that choice.

Most importantly, the AI notices the temporal structure of the final stanza. "I shall be telling this with a sigh" ("Road Not Taken" 16) places the interpretation of the decision in an imagined future. The meaning of the choice is therefore not simply present at the moment it is made; it is retrospectively constructed through a future act of narration. The response therefore reads the poem not merely as about choosing between roads, but as about how a person later tells the story of a choice whose alternatives could not initially be distinguished.

The response's treatment of the final "sigh" provides a further test of interpretive adequacy. It recognizes the expression's ambiguity, but the range of possible emotional meanings does not by itself explain its rhetorical function. In a letter to Edward Thomas

dated 26 June 1915, Frost described the poem's final sigh as a "mock sigh, hypo-critical for the fun of the thing" (Frost, *Letters* 321). This comment complicates readings that treat the sigh simply as regret or nostalgia, suggesting instead that the speaker's retrospective account of his choice may involve an element of performance. The AI response identifies the retrospective construction of meaning, but the historical evidence allows us to specify more sharply the possible relation between self-narration, irony, and the cultural appeal of individual choice.

The AI's reading demonstrates substantial but not exhaustive interpretive adequacy. It selects relevant textual evidence, identifies contradictions, connects formal regularity with conceptual instability, and recognises the retrospective construction of meaning. However, its treatment of the final sigh remains incomplete: although it recognises ambiguity, it does not connect the sigh to Frost's description of it as "mock" and "hypo-critical." This context complicates a straightforward reading of regret and reveals a dimension of the poem's irony that the AI's reading leaves underdeveloped. The reading therefore demonstrates what the system can accomplish as a critical text without settling what kind of agency, if any, should be attributed to its producer. Its sophistication makes that distinction more necessary, not less.

4. *Recitatif*: Readerly Position and Metacriticality

Morrison's 'Recitatif' poses a more demanding problem: can an AI-generated reading register not only what a text withholds, but the interpretive activity that those absences provoke? The significance of Morrison's strategy lies not simply in the story's ambiguity but in its capacity to make the reader's attempt to resolve that ambiguity part of the interpretive structure. Phelan's rhetorical and cognitive reading emphasises Morrison's sustained frustration of attempts to stabilise the racial identities of Twyla and Roberta (122–25). Goldstein-Shirley extends this reader-centred account, arguing that the story's ambiguous racial codes require readers to confront their own assumptions and participate in 'co-creat[ing] meanings along with the text' (89–93). Kim similarly argues that the story's indeterminacy redirects attention from solving the racial puzzle towards examining the assumptions that make such categorisation seem necessary (134–35). The relevant criterion for evaluating the AI response, therefore, is not simply whether it recognises uncertainty, but whether it apprehends the reader's implication in producing that uncertainty.

The AI response identifies this mechanism almost immediately. Rather than treating the missing racial labels as an informational deficit awaiting resolution, it asks how their absence structures the reader's activity. It observes that readers search elsewhere for classificatory clues and that the withholding of explicit identification generates the desire to

classify (Morrison 1–3, 11–12, 14–15). Uncertainty consequently appears not as a deficiency in the text but as an effect produced by its design.

Maggie exposes a different dimension of the story's indeterminacy. The AI describes her as 'the story's largest textual gap' and attends to her silence, vulnerability, and inability to provide an authoritative account of what happened (Morrison 5–7). It then places Twyla's and Roberta's contradictory memories in tension rather than treating either as a definitive reconstruction of events (Morrison 22–24). Twyla's movement from denying that she kicked Maggie to admitting that she wanted to participate in the violence—'I didn't kick her; I didn't join in with the gar girls and kick that lady, but I sure did want to' (Morrison 34–35)—further unsettles the distinction between remembered action, imagined participation, and acknowledged desire. Roberta's memory cannot simply function as an objective correction either, because her own account changes. The AI thus refuses the conventional interpretive desire to identify one account as finally authoritative.

But the AI's emphasis on racial indeterminacy leaves another dimension of the story comparatively underdeveloped: the conditions under which Maggie becomes an object of interpretation. Scholarship on Maggie foregrounds her disability, silence, partial representation, and uncertain status as an object of memory, suggesting that the desire to render her fully intelligible is itself implicated in the story's representational difficulty (Sklar 147–52). The AI's attention to contradictory memories is therefore textually productive, but its interpretive weighting remains uneven: racial classification receives greater explanatory force than the problem of how Maggie is represented and made knowable.

The AI's final formulation of this process is particularly revealing: "absence → inference → apparent certainty → contradiction → revision → renewed uncertainty." The sequence does more than summarise the story's ambiguity; it models interpretation as a recursive process in which readers construct hypotheses from textual gaps, temporarily stabilise those constructions, encounter contradictory evidence, and revise them. This is what may be termed metacriticality: the capacity to represent not merely what a text means, but the operations through which meaning is pursued, stabilised, and unsettled. The more defensible point is epistemological: as with human criticism, the internal status of the interpreting entity is not directly observable from the interpretive product alone.

5. From Interpreter to Interpretive Relation

The two case studies make a simple human/machine hierarchy difficult to sustain. Human interpretation cannot be treated as automatically authentic while machine interpretation is dismissed as derivative, because the preceding analysis has already shown that human reading is itself historically, socially, and materially mediated.

Sofie Vlaad provides a useful point of comparison, although her question is different from the one pursued here. She argues that the absence of a conventional author does not necessarily prevent an AI-generated text from sustaining literary meaning, ultimately defending postulated author hypothetical intentionalism as an account of how such meaning may be fixed (7–10). Her concern is the attribution of meaning to AI-generated texts, not the attribution of interpretive agency to an AI system reading an existing work. The distinction matters. A text may sustain meaning without a conventional author; it does not follow that a machine producing a reading of an existing text thereby becomes an interpreter in the same sense as a human reader. The question is instead how interpretive activity and responsibility are constituted when a machine produces the reading.

Interpretive possibility is never generated from nowhere. What counts as a meaningful textual observation is shaped by disciplinary expectations, learned conventions, and communal practices. In the case of AI, this accumulated inheritance becomes unusually visible: its outputs are generated from patterns learned across vast bodies of human-produced language and are further conditioned by prompts, textual inputs, and user expectations. This does not collapse the distinction between human interpretation and machine generation. Rather, it makes the distinction harder to locate in the output alone.

Bender et al. caution against treating linguistic fluency as evidence of natural-language understanding, arguing that language models learn patterns of linguistic form to which readers may attribute meaning and communicative intent (615–17). This caution is directly relevant to literary criticism: a coherent and textually persuasive reading does not, by itself, establish an autonomous interpreting subject. At the same time, the datasets on which language models are trained can reproduce dominant cultural viewpoints and existing biases (613–14). An AI-generated reading must therefore be considered both as an interpretive performance and as an artefact of inherited linguistic and cultural patterns.

Barad's agential realism offers the paper a more radical account of agency. She argues that phenomena, rather than independently existing entities, are ontologically primary, and that intra-actions enact the boundaries through which entities become determinate. Agency, accordingly, is not an attribute possessed by a pre-existing subject or object but an ongoing reconfiguration of the world (139–41). Applied to literary interpretation, Barad's framework shifts attention from whether AI independently possesses agency to how interpretive agency emerges through intra-actions involving the text, reader, machine, prompt, and accumulated critical practices. This application does not establish that an AI system interprets in the same sense as a human reader; it offers a way to examine interpretation as a relational process without presuming that agency must originate in a single, autonomous interpreter.

Jörg Noller's account of AI as an extension of human agency further complicates the opposition between autonomous human interpretation and machine performance. Rather than locating agency exclusively in either the human subject or the technological object, Noller understands AI as an extension of human agency through the interrelation of human purposes and intentions, socially grounded databases, and algorithms (4–5). Applied to literary interpretation, this account directs attention towards the conditions under which an AI-generated reading emerges, including the literary text, the prompt, learned critical conventions, and human evaluation. It does not establish that AI possesses independent interpretive agency; rather, it provides a framework for examining interpretation as a process mediated by human and technological relations. Hayles's reading of the Turing test offers a complementary caution against treating observable performance as a transparent index of the agent behind it. She argues that the test places perception and judgment within a technological circuit that joins represented and embodied subjects, thereby complicating the boundary between human and machine (xiii–xiv). The insight does not settle whether an AI has interpreted a poem or story. It sharpens instead the distinction between what an interpretive performance makes observable and what can legitimately be inferred about the agency behind it.

From this perspective, the interpreter becomes part of an interpretive relation rather than its sole autonomous source. This shifts attention towards how agency is distributed and constituted, without assuming that all interpretations are equivalent or resolving which participants should receive credit for an interpretive act.

Conclusion

The question of whether an interpretation can exist without an interpreter cannot be answered by treating AI either as a failed human reader or as a newly autonomous one. The two case studies show that an LLM can produce sophisticated literary analysis: it can identify contradiction, attend to formal and textual detail, and model the processes through which readers construct and revise meaning. Yet the sophistication of the reading does not, by itself, establish the presence of an autonomous interpreter behind it.

AI's significance therefore lies less in proving that machines can replace human readers than in exposing the assumptions embedded in the category of the interpreter itself. Literary interpretation has long been understood as historically situated, socially mediated, and shaped by relations among texts, readers, and interpretive communities. AI makes these conditions unusually visible because its readings emerge through interactions among textual input, prompts, learned linguistic patterns, critical conventions, technological systems, and human evaluation.

The distinction between interpretive performance and interpretive agency therefore need not culminate in another human/machine opposition. The AI-generated reading is neither simply an autonomous machine interpretation nor merely a meaningless simulation; it is an interpretive event produced through a network of human and nonhuman conditions. The question, then, is no longer simply whether a machine can generate a reading, but how—and through whom or what—interpretation comes into being.

Appendix: Prompts Used to Generate the AI Readings

Prompt 1: Robert Frost's "The Road Not Taken"

Give me a close reading of Robert Frost's "The Road Not Taken." Focus specifically on contradiction, ambiguity, formal features, and textual evidence. Do not give me a predetermined interpretation or tell me what the poem is "really" about. Develop the interpretation from the text itself.

Prompt 2: Toni Morrison's "Recitatif"

Give me a close reading of Toni Morrison's "Recitatif." Focus specifically on contradiction, ambiguity, textuales gaps, readerly uncertainty, and textual evidence. Do not give me a predetermined interpretation or tell me what the story is "really" about. Develop the interpretation from the text itself.

Works Cited

Barad, Karen. *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning*. Duke UP, 2007.

Bender, Emily M., et al. "On the Dangers of Stochastic Parrots: ¿Can Language Models Be Too Big?" *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, Association for Computing Machinery, 2021, pp. 610–623. https://doi.org/10.1145/3442188.3445922.

Braidotti, Rosi. *The Posthuman*. Polity, 2013.

Fish, Stanley. "'Is There a Text in This Class?'" Translated by Rafael Eugênio Hoyos-Andrade, *Alfa*, vol. 36, 1992, pp. 189–206.

Frost, Robert. *The Letters of Robert Frost*. Edited by Donald Sheehy, Mark Richardson, and Robert Faggen, vol. 1, 1886–1920, Belknap Press of Harvard UP, 2014.

Frost, Robert. "The Road Not Taken." *The Poetry of Robert Frost*, edited by Edward Connery Lathem, Henry Holt, 1969, pp. 105–06.

Gadamer, Hans-Georg. *Truth and Method*. 2nd rev. ed., translated by Joel Weinsheimer and Donald G. Marshall, Continuum, 2004.

Goldstein-Shirley, David. "Race/[Gender]: Toni Morrison's 'Recitatif.'" *Journal of the Short Story in English*, no. 27, 1996, pp. 83–96. https://doi.org/10.3406/jsse.1996.1453.

Hayles, N. Katherine. *How We Became Posthuman: Virtual Bodies in Cybernetics, Literature, and Informatics*. U of Chicago P, 1999.

Hirsch, E. D., Jr. *Validity in Interpretation*. Yale UP, 1967.

Ichien, Nicholas, Dušan Stamenković, and Keith J. Holyoak. "Large Language Model Displays Emergent Ability to Interpret Novel Literary Metaphors." *Metaphor and Symbol*, vol. 39, no. 4, 2024, pp. 296–309. https://doi.org/10.1080/10926488.2024.2380348.

Iser, Wolfgang. *The Act of Reading: A Theory of Aesthetic Response*. Johns Hopkins UP, 1978.

Jauss, Hans Robert. *Toward an Aesthetic of Reception*. Translated by Timothy Bahti, U of Minnesota P, 1982.

Kim, Miehyeon. "Sympathy and Indeterminacy in Toni Morrison's 'Recitatif.'" *Feminist Studies in English Literature*, vol. 23, no. 1, 2015, pp. 133–166. https://doi.org/10.15796/fsel.2015.23.1.005.

Mitchell, Melanie, and David C. Krakauer. "The Debate over Understanding in AI's Large Language Models." *Proceedings of the National Academy of Sciences*, vol. 120, no. 13, 2023, article e2215907120. https://doi.org/10.1073/pnas.2215907120.

Morrison, Toni. *Recitatif: A Story*. Introduction by Zadie Smith, Alfred A. Knopf, 2022.

Noller, Jörg. "Extended Human Agency: Towards a Teleological Account of AI." *Humanities and Social Sciences Communications*, vol. 11, 2024, article 1338.

Phelan, James. "Rhetorical Theory, Cognitive Theory, and Morrison's 'Recitatif': From Parallel Play to Productive Collaboration." *The Oxford Handbook of Cognitive Literary Studies*, edited by Lisa Zunshine, Oxford UP, 2015, pp. 120–135. https://doi.org/10.1093/oxfordhb/9780199978069.013.0007.

Robinson, Katherine. "Robert Frost: 'The Road Not Taken.'" *Poetry Foundation*, 27 May 2016, www.poetryfoundation.org/articles/89511/robert-frost-the-road-not-taken.

Sklar, Howard. "'What the Hell Happened to Maggie?': Stereotype, Sympathy, and Disability in Toni Morrison's 'Recitatif.'" *Journal of Literary and Cultural Disability Studies*, vol. 5, no. 2, 2011, pp. 137–154. https://doi.org/10.3828/jlcds.2011.12.

Vlaad, Sofie. "Texts without Authors: Ascribing Literary Meaning in the Case of AI." *The Journal of Aesthetics and Art Criticism*, vol. 83, no. 1, 2025, pp. 4–11. https://doi.org/10.1093/jaac/kpae047.