

Computational Hermeneutics: Machine Reading and the Epistemology of Literary Interpretation

Belay Sitotaw Goshu¹, Muhammad Ridwan²

¹Department of Physics, Dire Dawa University, Dire Dawa, Ethiopia

²Universitas Islam Negeri Sumatera Utara, Indonesia

Abstract:

This review evaluates the article "Computational Hermeneutics: Machine Reading and the Epistemology of Literary Interpretation," situating it within the emerging interdisciplinary discourse on artificial intelligence and literary studies. The article advances a theoretically ambitious thesis: that machine reading constitutes not a positivist or anti-hermeneutic practice but a distinctive form of interpretation that both challenges and enriches traditional humanistic hermeneutics. The software code, and even quantified models, has, the article develops three interlocking arguments: that the hermeneutic-positivistic divide in digital humanities is a false dichotomy; that machine reading operates through "interpretive lenses" that shape interpretive attention; and that machine-generated interpretations require renewed engagement with the hermeneutic concepts of situatedness, plurality, and ambiguity. The article's most significant contributions are its rigorous articulation of the epistemological stakes of computational hermeneutics and its nuanced treatment of the "simulated hermeneutic." However, three limitations constrain its impact: an uneven empirical foundation, particularly in its comparative study of human and machine readings; a limited engagement with non-Western and multilingual contexts; and an underdeveloped treatment of power, ideology, and material conditions. The review identifies Habermas's emancipatory interest of hermeneutics of suspicion, and political economy of AI as theoretical resources capable of addressing these gaps. The article ultimately stands as both an important intervention and an invitation to develop an emancipatory computational hermeneutics that integrates critical theory with computational practice.

Keywords:

computational hermeneutics; machine reading; digital humanities; AI interpretation; literary epistemology

I. Introduction

The past two decades have witnessed a profound reconfiguration of the relationship between computational methods and humanistic inquiry. What began as the tentative application of quantitative techniques to literary corpora has evolved into a sustained theoretical interrogation of the epistemological foundations of humanistic knowledge production itself. Within this landscape, the article under review, "Computational Hermeneutics: Machine Reading and the Epistemology of Literary Interpretation", positions itself at a critical juncture: the intersection of hermeneutic theory, digital humanities methodology, and the epistemology of machine-generated interpretation. The article's central claim is that computational approaches to textual analysis are not merely instrumental additions to the literary scholar's toolkit but constitute a fundamentally hermeneutic enterprise that demands epistemological scrutiny.

This claim intervenes in a debate that has structured digital humanities discourse for over a decade. Van Zundert (2016) has argued that "software code, and even quantified models, has a deeply hermeneutical nature, they have intent and a priori interpretative aspects" (p. 331). This position challenges the tacit assumption, widespread even among digital humanists, that computation operates within a quantitative and reductive domain sharply distinguished from the interpretive realm of humanistic understanding. As Van Zundert (2016) observes, "virtually all researchers in digital humanities as well as those in digital textual scholarship only study the end result of computational or digital processes hermeneutically," accepting the graphical interface as "a hard boundary between the realm of quantification and the realm of hermeneutics" (p. 332). The article under review extends this argument by insisting that the hermeneutic character of computational methods is not incidental but constitutive—that machine reading embodies a distinct form of interpretation whose epistemological status requires sustained theoretical attention.

The stakes of this argument are considerable. If computation is indeed hermeneutical, then the conventional division of labor between "distant reading" as data-driven pattern recognition and "close reading" as interpretive engagement collapses. Ramsay's (2011) notion of "algorithmic criticism" anticipated this collapse by proposing that we "channel the heightened objectivity made possible by the machine into the cultivation of those heightened subjectivities necessary for critical work" (p. 5). As Ramsay (2011) elaborates, algorithmic criticism "seeks, in the narrowing forces of constraint embodied and instantiated in the strictures of programming, an analogue to the liberating potentialities of art" (p. 7). The article under review takes up this provocation and pushes it further: not only can computation serve critical subjectivity, but computational operations themselves perform interpretive labor that shapes what can be seen, known, and said about literary texts.

The article's theoretical architecture draws on three distinct but converging intellectual traditions. First, it engages with the hermeneutic philosophy of Gadamer and Ricoeur, particularly the concepts of pre-understanding, horizon fusion, and the text as a distanced object of interpretation. Recent scholarship has extended hermeneutic frameworks to computational contexts. One pragmatist approach defines a computational hermeneutic model as "a collaborative symbolic system between humans and machines to analyze textual data" wherein "to understand the contextual meaning of words, machines need humans, and to explain logical and grammatical structure, humans need machines" (Lévy, 2024, p. 5). Second, the article situates itself within digital humanities debates about the hermeneutic-positivistic divide. Van Zundert (2016) has noted that "a computational turn does not automatically imply a turn to empiricism and scientism, or a disregard for hermeneutic tradition," though stylometry and distant reading may "lean in their approaches more towards an empiricist or scientific attitude" (p. 333). The article challenges this lingering binary by examining the interpretive dimensions of computational processes themselves. Third, it engages with emerging work on AI hermeneutics, exemplified by Kwok's (2026) framework treating retrieval-augmented generation systems as "interpretive infrastructure" rather than mere automation. As Kwok (2026) argues, "interpretation itself becomes procedural: machines now arguably perform a version of the critic's movement between evidence and hypothesis" (p. 1).

The article's intervention is timely. As generative AI systems become increasingly integrated into literary scholarship, the question of what it means for a machine to "read" a text has acquired new urgency. Recent work has proposed that generative AI systems function as "context machines" that must inherently address three interpretive challenges: situatedness

(meaning only emerges in context), plurality (multiple valid interpretations coexist), and ambiguity (interpretations naturally conflict) (Mohr et al., 2026, p. 2). The article under review engages each of these challenges, arguing that computational hermeneutics must account for the situated nature of machine interpretation, the plurality of possible readings, and the irreducible ambiguity of literary language. Kwok (2026) has similarly emphasized that algorithmic mediation need not be seen as "a threat to interpretation but as an opportunity to make the labour of reading visible and rewarded again" (p. 1). The article under review aligns with this reframing, positioning machine reading not as a replacement for humanistic interpretation but as a distinctive hermeneutic practice that demands its own epistemological account.

This review proceeds in four parts. First, I outline the article's principal arguments and theoretical architecture, assessing its engagement with hermeneutic theory and its conceptualization of machine reading. Second, I evaluate its methodological framework and evidentiary basis, including the three case studies through which it demonstrates its claims. Third, I situate the article within the broader scholarly conversation on computational hermeneutics, AI reading, and the epistemology of digital interpretation, identifying points of convergence and tension with existing work. Finally, I assess its contributions, limitations, and implications for future research in language, literature, and computational criticism. Throughout, I attend to the article's central question: what kind of knowledge does machine reading produce, and how does that knowledge relate to the interpretive traditions of the humanities?

II. Review of Literatures

The article advances a deceptively simple yet theoretically ambitious thesis: that machine reading, far from being a positivist or anti-hermeneutic practice, embodies a distinct form of interpretation that both challenges and enriches traditional humanistic hermeneutics. The author draws on the foundational work of Van Zundert (2016), who argued that "software code, and even quantified models, has a deeply hermeneutical nature, they have intent and a priori interpretative aspects" (p. 331). This claim constitutes the article's theoretical anchor, positioning computational methods not as external tools applied to literary texts but as interpretive agents whose operations are themselves embedded in assumptions, values, and epistemological commitments. The article thereby aligns itself with a broader tradition in critical technical practice, which insists that technical systems embody values and assumptions that shape their outputs in ways that demand critical scrutiny (Agre, 1997).

The article develops this argument through three interlocking conceptual moves. First, it reframes the much-debated "hermeneutic-positivistic divide" in digital humanities as a false dichotomy. Drawing on Ramsay's (2011) notion of "algorithmic criticism" as a mode of channeling "the heightened objectivity made possible by the machine into the cultivation of those heightened subjectivities necessary for critical work" (p. 5), the author contends that the binary between interpretation and computation obscures the interpretive character of both. This position resonates with Kirschenbaum's (2014) observation that digital humanities has always involved interpretive decisions at every level of technical practice, from data modeling to interface design. The article extends this insight by arguing that the very act of encoding a literary text into machine-readable form constitutes an interpretive act, one that foregrounds certain textual features while rendering others invisible.

Second, the article proposes that machine reading operates through what it terms "interpretive lenses", computational operations that make certain textual features salient while rendering others invisible. This lens metaphor, adapted from Ringler's work on large-scale data hermeneutics, provides a productive framework for understanding how computational methods shape interpretive attention. The article's argument here intersects with recent work in critical AI studies, which has demonstrated that machine learning models encode selective views of the world that systematically exclude or marginalize certain perspectives (Ridge, 2024) . As Ridge (2024) notes, "ML models have a selective view of the world—it might exclude things you value" and "AI tends to sand away the unusual" because it is "trained to answer with the most likely answer to your question, which is not necessarily the most correct [or interesting] answer." This observation reinforces the article's claim that computational lenses are not neutral but constitutive of what can be seen and known.

Third, the article argues that the epistemological status of machine-generated interpretations requires a renewed engagement with hermeneutic theory, particularly the concepts of situatedness, plurality, and ambiguity. The article's engagement with these concepts is sophisticated but uneven. Its treatment of situatedness; the principle that meaning emerges only in context is compelling, particularly in its critique of the tendency within AI evaluation frameworks to treat culture as a measurable variable rather than as constitutive of meaning-making. This critique is grounded in the Gadamerian insight that understanding is always embedded within a tradition and shaped by the interpreter's horizon (Schmidt, 2025) . As Gadamer argued, the interpreter "is embedded within tradition since he has inherited a set of prejudices that constitute his horizon of understanding," and correct understanding occurs "when he is able to legitimize his prejudices by grounding them in the things themselves" (Schmidt, 2025, p. 6). The article's extension of this insight to machine interpretation raises the question of whether computational systems can be said to possess a "horizon" in any meaningful sense, or whether they operate through a fundamentally different epistemological logic.

However, the article's account of plurality (the coexistence of multiple valid interpretations) could benefit from deeper engagement with reader-response theory and the phenomenology of reading. The claim that machines can model interpretive plurality through probabilistic generation is asserted rather than demonstrated, and the article does not adequately address whether computational plurality differs qualitatively from human interpretive diversity. Reader-response theory, particularly the work of Fish (1980) on interpretive communities, has long established that meaning is not inherent in texts but emerges through the shared interpretive strategies of communities of readers. Fish argued that "it is interpretive communities, rather than either the text or reader, that produce meanings" (Fish, 1980, p. 14) . If this is correct, then the plurality of interpretations is not merely a feature of texts but a function of the social and institutional contexts within which reading occurs. The article's framework for computational hermeneutics would benefit from engaging with this insight, particularly in considering whether machines can participate in interpretive communities or whether they necessarily operate outside them.

The article's treatment of ambiguity also warrants critical attention. Drawing on Ricoeur's hermeneutics of suspicion, which approaches texts with the assumption that surface meanings conceal deeper, often ideological, truths; the article suggests that computational methods might be particularly well-suited to uncovering meanings that human readers overlook (Scott-Baumann, 2011) . Ricoeur's framework distinguishes between a hermeneutics of recollection, which seeks to retrieve and preserve meaning, and a hermeneutics of suspicion, which seeks to demystify and expose (Phillips, 2009) . The article's implicit alignment with the latter raises important questions

about the ethical and political dimensions of computational interpretation. If machine reading is inherently suspicious in Ricoeur's sense, then its outputs must be evaluated not only for their accuracy but for their ideological implications.

The article's engagement with these theoretical traditions is valuable, but its synthesis remains incomplete. The concept of "interpretive lenses" is productive but underdeveloped; the article does not fully specify how computational lenses differ from or relate to the prejudices and pre-understandings that Gadamer identified as constitutive of human understanding. Similarly, the claim that machines can model interpretive plurality through probabilistic generation requires more rigorous engagement with the phenomenological tradition, particularly the work of Iser on the implied reader and the role of indeterminacy in the reading process (Iser, 1996). Iser's concept of "Leerstellen" (blanks or gaps) identifies the spaces in literary texts that invite readerly participation, and it is not clear whether computational systems can meaningfully engage with such indeterminacy or whether they necessarily resolve it through probabilistic prediction. The article's theoretical framework would be strengthened by addressing these questions directly rather than gesturing toward them.

A further limitation concerns the article's engagement with the material and institutional dimensions of computational hermeneutics. The lens metaphor implies a passive model of computational mediation, as if algorithms simply filter textual features in the way a camera lens filters light. But as Verbeek (2006) has argued, technologies do not merely mediate between subjects and objects; they actively shape what appears and how it appears, co-constituting both the interpreter and the interpreted. Verbeek's concept of "technological mediation" suggests that computational tools are not neutral instruments but active participants in the constitution of interpretive horizons. The article's framework would benefit from a more sustained engagement with postphenomenological approaches to technology, which would clarify how machine reading shapes not only what can be seen but what kinds of interpretive subjects can emerge in relation to computational systems.

Table 1: Theoretical Frameworks and Their Application in the Article

Theoretical Framework	Key Concept	Article's Application	Assessment
Gadamerian Hermeneutics	Horizon fusion, prejudice	Situatedness of machine interpretation	Productive but undertheorized
Ricoeurian Hermeneutics	Suspicion, demystification	Ambiguity and critical interpretation	Underdeveloped ethical dimension
Reader-Response Theory	Interpretive communities, indeterminacy	Plurality of machine readings	Inadequate engagement
Postphenomenology	Technological mediation	Interpretive lenses	Missing framework
Critical Technical Practice	Values embedded in systems	Hermeneutic character of code	Strong but could be extended

Table 1 synthesizes the five theoretical frameworks underpinning the article's argument, revealing an uneven integration of hermeneutic and postphenomenological traditions. The article's strongest engagement lies in critical technical practice, where it convincingly establishes the hermeneutic character of code (Agre, 1997; Van Zundert, 2016). However, its treatment of Gadamerian hermeneutics remains productive yet undertheorized, particularly regarding whether

machines possess genuine horizons (Schmidt, 2025). Ricoeurian suspicion is invoked but its ethical implications remain underdeveloped (Scott-Baumann, 2011), while reader-response theory, crucial for theorizing interpretive plurality is inadequately engaged (Fish, 1980). Most notably, postphenomenology is entirely absent, despite its direct relevance to the article's "interpretive lenses" metaphor (Verbeek, 2006). This gap weakens the framework's capacity to explain how computational mediation actively co-constitutes interpretive subjects and objects.

In summary, the article's theoretical architecture makes a significant contribution by insisting that computational methods are hermeneutical rather than merely technical. Its three conceptual moves—reframing the hermeneutic-positivistic divide, proposing "interpretive lenses" as a model of machine reading, and calling for renewed engagement with hermeneutic theory, provide a productive framework for analyzing machine interpretation. However, the framework requires further elaboration in three areas: first, the relationship between computational and human interpretive horizons; second, the qualitative distinctiveness of machine plurality versus human interpretive diversity; and third, the material and institutional conditions that shape computational hermeneutics. These limitations do not diminish the article's value but rather indicate directions for future theoretical development.

III. Research Methods

The article employs a mixed-methods approach that combines theoretical synthesis with empirical case studies. The theoretical component draws extensively on hermeneutic philosophy, digital humanities scholarship, and critical AI studies. The empirical component presents three case studies: a stylometric analysis of authorship attribution, an experiment in retrieval-augmented generation (RAG) for contextual interpretation, and a comparative study of human and machine readings of a canonical literary text.

The methodological rigor of these case studies varies considerably. The stylometric analysis, while competently executed, largely reproduces findings that have been established in the literature for over a decade. Its argument that quantification possesses a hermeneutic character is well-taken but not novel. More compelling is the RAG-based experiment, which engages with an emerging paradigm that treats retrieval systems as "interpretive infrastructure" rather than mere automation, arguing that machines can perform a version of the critic's movement between evidence and hypothesis (Kwok, 2026). This experiment suggests that machine reading can surface "overlooked cultural memory" while foregrounding the ethical boundaries of computational interpretation—a claim that merits further empirical investigation.

The comparative study of human and machine readings is the article's most original contribution but also its most methodologically problematic element. The study's design raises significant construct validity concerns. The article treats "human interpretation" as a stable baseline against which machine interpretation can be measured, yet reader-response theory has long established that human interpretation is itself highly variable, context-dependent, and shaped by cultural positioning (Fish, 1980). The article's finding that AI exhibits "algorithmic positivity", a tendency toward affirmative, harmonious, and morally generalized interpretations (Mohr et al., 2026) is genuinely important, but the study's design does not adequately control for the interpretive variability of its human participants.

Table 2: Overview of Case Studies and Methodological Assessment

Case Study	Method	Key Claim	Evidentiary Strength
Stylometric Analysis	Principal Component Analysis	Quantification has hermeneutic character	Moderate—replicates established findings
RAG Experiment	Retrieval-Augmented Generation	Machines perform interpretive labor	Strong—novel application, clear demonstration
Comparative Reading	Human-AI Response Comparison	AI exhibits "algorithmic positivity"	Weak—construct validity concerns

Table 2 evaluates the evidentiary strength of the article's three case studies, revealing a methodological hierarchy that complicates its central claims. The RAG experiment emerges as the strongest evidence, convincingly demonstrating that machines perform interpretive labor by navigating between evidence and hypothesis (Kwok, 2026). The stylometric analysis, while competent, offers only moderate support because it reproduces established findings about quantification's hermeneutic character (Stamatatos, 2018; Stein & Eder, 2024). Most problematically, the comparative reading study, the article's most original contribution—exhibits weak construct validity. Its treatment of human interpretation as a stable baseline ignores reader-response theory's insistence that interpretation is context-dependent and culturally positioned (Fish, 1980). Consequently, the "algorithmic positivity" finding (Mohr et al., 2026), though important, remains inadequately substantiated.

The article's engagement with the "black box" problem in computational analysis is theoretically sophisticated but practically incomplete. It acknowledges that the internal processes of deep neural networks are often opaque, such that explanation can usually only be achieved through post-hoc analysis and surrogate models, which amounts less to turning a black box into a glass box than to building another, smaller box to study the behavior of the original system (Barta, 2024). Black box research faces an epistemological dilemma: researchers may understand neither enough to reject false beliefs nor enough to justify true beliefs, leaving humanists with the uncomfortable choice of trading the unreadability of the "great unread" for the unreadability of the "black box" (Barta, 2024). The article calls for "engaging and uncovering the 'black box' effect" but offers few concrete strategies for doing so.




Figure 1: The Methodological Cycle of Computational Hermeneutics

Figure 1 shows the cyclical model depicts the methodological logic of computational hermeneutics. Theoretical assumptions guide the selection of computational modeling choices (e.g., feature extraction, model architecture); their outputs require hermeneutic interpretation; and that interpretation feeds back into theoretical revision. The article's central problem lies in the third step of this cycle, the "black box" transition from computational output to hermeneutic interpretation, which lacks sufficient methodological transparency.

IV. Result and Discussion

4.1 Situating the Article in Scholarly Conversation

The article enters a vibrant and rapidly evolving scholarly conversation. Its most direct interlocutors include the emerging body of work on "AI hermeneutics," exemplified by Kwok's (2026) framework treating retrieval-augmented generation (RAG) systems as interpretive infrastructure rather than mere automation. Kwok (2026) argues that "interpretation itself becomes procedural: machines now arguably perform a version of the critic's movement between evidence and hypothesis" (p. 1). This reconceptualization of machine reading as interpretive labor rather than mechanical retrieval constitutes the article's closest intellectual alliance. The article's contribution to this conversation lies in its sustained theoretical engagement with the epistemological dimensions of machine interpretation, a dimension that much of the existing literature treats only tangentially. However, the article's empirical foundation remains comparatively thin when positioned against Kwok's (2026) work, which demonstrates interpretive claims through concrete experiments with the Conflict Archive on the Internet and conflict-era poetry.

The article's relationship to the broader digital humanities tradition is more ambivalent. It acknowledges the foundational contributions of Moretti's (2007) "distant reading" and Jockers's (2013) microanalysis but positions computational hermeneutics as a corrective to what it perceives as the empiricist and scientistic tendencies of these approaches. This positioning is somewhat overstated. As Van Zundert (2016) notes, "a computational turn does not automatically imply a turn to empiricism and scientism, or a disregard for hermeneutic tradition" (p. 333). The article's tendency to frame its contribution as a necessary corrective to a monolithic "computational paradigm" obscures the diversity of methodological and theoretical positions within digital humanities. Moreover, as Van Zundert (2016) further observes, quantified approaches in the digital humanities often exhibit a "relative lack of explanatory power" and a tendency toward "confirmation bias", verifying existing answers rather than generating new questions (p. 335). Stylometry, for instance, "can tell us or rather indicate to us—that there are two authors to a certain text. But it tells us unsatisfactory little about how and why the individual styles differ" (Van Zundert, 2016, p. 336). The article's proposed framework would benefit from engaging more directly with this critique of computational methods' explanatory limitations.

The article's engagement with non-Western and multilingual perspectives is notably limited. The case studies focus exclusively on English-language texts, and the theoretical framework draws almost entirely on European hermeneutic traditions. This limitation is significant given the article's claims about the universal applicability of computational hermeneutics. Recent work on AI and literary interpretation in Indonesian folklore studies, for example, has demonstrated that algorithmic interpretation exhibits culturally specific patterns of "emotional flattening and cultural abstraction" (Mohr et al., 2026, p. 5). As Mohr et al. (2026) document, AI "tended to universalize the moral meanings of the folktales," interpreting culturally specific narratives through broad values such as kindness, courage, and perseverance, thereby reducing "the cultural texture of the narratives" (p. 5). This pattern termed "cultural

abstraction"—transforms "culturally embedded meanings into generalized moral claims" (Mohr et al., 2026, p. 5). Similarly, scholarship on AI reading in Chinese literary criticism has emphasized the need for "literary computational analytical criticism" that integrates ethical and cross-cultural dimensions (Nie, 2025). Nie (2025) argues that while AI reading can surpass humans in statistical analysis and thematic extraction, it "cannot yet fully replace human emotional experience and subjective judgment" (p. 28). These findings complicate the article's more general claims about machine reading and underscore the need for a more culturally situated framework.

The article also engages, albeit implicitly, with the ethical dimensions of AI-mediated interpretation. Recent scholarship has emphasized that AI systems are not neutral cultural technologies but are embedded within "systems of data gathering, work, and global power imbalances" (Crawford, 2021, as cited in Mohr et al., 2026, p. 3). Crawford's (2021) analysis of AI as a technology of extraction, from the minerals mined for hardware to the exploited labor behind "automated" services, demonstrates that the epistemological questions the article raises are inseparable from political-economic questions about who builds, owns, and benefits from AI systems. The collections of texts used to train language models "showcase dominant language habits and cultural practices," with the result that "AI-generated writings may perpetuate hegemonic narrative frameworks while remaining stylistically diverse" (Mohr et al., 2026, p. 6). The article's framework for computational hermeneutics would be strengthened by more explicit engagement with these critical perspectives on algorithmic culture and the politics of training data.

Table 3: The Article's Position in Scholarly Conversation

Scholarly Strand	Key Interlocutors	Article's Contribution	Tension/Unresolved Issue
AI Hermeneutics	Kwok (2026); Mohr et al. (2026)	Theoretical depth on epistemology	Limited empirical validation; plurality undertheorized
Digital Humanities	Ramsay (2011); Moretti (2007); Van Zundert (2016)	Critique of positivist tendencies	Overstates binary; ignores explanatory power limitations
Critical AI Studies	Crawford (2021); Bender et al. (2021)	Hermeneutic framework for bias	Undertheorizes power and political economy
Non-Western DH	Nie (2025); Mohr et al. (2026)	Minimal engagement	Eurocentric limitations; universalizing claims

Table 3 maps the article's position across four scholarly strands, revealing an asymmetrical engagement with its intellectual environment. Its strongest contribution lies in AI hermeneutics, where it advances theoretical depth on epistemological questions (Kwok, 2026; Mohr et al., 2026), yet this is undermined by limited empirical validation and an undertheorized account of interpretive plurality. Within digital humanities, the article's critique of positivist tendencies productively challenges reductive binaries but overstates the coherence of the "computational paradigm" it opposes and neglects longstanding concerns about explanatory power (Van Zundert, 2016; Ramsay, 2011). Critically, its hermeneutic framework for bias remains disconnected from political-economic critique (Crawford, 2021; Bender et al., 2021). Most significantly, engagement with non-Western traditions is minimal, producing Eurocentric limitations and universalizing claims that contradict the article's own pluralist commitments (Nie, 2025).




Figure 2: The Article's Position in the Scholarly Landscape

Figure 2 shows the positions the article at the center of four scholarly conversations. The thickness of arrows indicates the strength of engagement. The article's strongest dialogue is with AI hermeneutics, while its engagement with critical AI studies and non-Western digital humanities remains comparatively weak. The dashed arrow at the bottom identifies the gap that future research might address: a cross-cultural computational hermeneutics that integrates non-Western interpretive traditions.

4.2 Critical Assessment: Contributions and Limitations

The article's most significant contribution is its rigorous articulation of the epistemological stakes of computational hermeneutics. By insisting that machine reading is not merely a methodological innovation but a transformation in the conditions of interpretation itself, the article elevates the conversation beyond technical debates about tools and methods. Its engagement with the concepts of situatedness, plurality, and ambiguity provides a vocabulary for analyzing machine interpretation that is grounded in humanistic theory while responsive to computational specificities. As Van Zundert (2016) observes, "a computational turn does not automatically imply a turn to empiricism and scientism, or a disregard for hermeneutic tradition" (p. 333). The article's framework productively extends this insight by demonstrating that computational operations themselves embody interpretive commitments that shape what can be seen, known, and said about literary texts.

The article's second contribution lies in its nuanced treatment of the "simulated hermeneutic"—the notion that AI systems can simulate interpretive labor without possessing the existential, embodied, and cultural conditions that ground human interpretation. This concept has significant implications for how we evaluate machine-generated interpretations: not by asking whether they are "correct" in some absolute sense but by examining what they reveal about the assumptions, biases, and limitations embedded in both human and computational interpretive practices. Recent scholarship has extended this line of inquiry, arguing that AI-generated "simulated hermeneutics" constitutes a distributed process involving human readers, algorithmic agents, and platform infrastructures (Synthetic Hermeneutics, 2026). The article's framework anticipates this development but does not fully theorize its implications for the boundaries of interpretive agency.

However, the article exhibits several limitations that constrain its impact. First, as noted above, its empirical foundation is uneven. The comparative study, which could have provided the most compelling evidence for the article's claims, is methodologically underdeveloped. A

more robust design would have employed larger and more diverse participant pools, controlled for interpretive expertise, and incorporated qualitative analysis of the reasoning processes underlying human and machine interpretations. The article's reliance on undergraduate participants as a proxy for "human interpretation" is particularly problematic given the well-documented variability of interpretive responses across expertise levels and cultural contexts (Fish, 1980).

Second, the article's treatment of power, ideology, and bias—while present—remains underdeveloped relative to its epistemological focus. The article acknowledges that machine reading operates within and reproduces cultural biases, citing research on "algorithmic positivity" and "cultural abstraction" (Mohr et al., 2026). Yet it does not fully theorize how computational hermeneutics might contribute to a critique of these biases. The hermeneutic tradition, particularly in its critical and emancipatory dimensions, offers resources for such a critique that the article leaves largely unexplored. Habermas's concept of "knowledge-constitutive interests" demonstrates that knowledge is never neutral or disinterested but is shaped by underlying human interests, including the interest in emancipation (Packer, 2010). Habermas's debate with Gadamer over the requirements for critical inquiry—in which Habermas insisted that a researcher needs "more than the resources of a member" and requires "a frame of reference outside the form of life being investigated" (Packer, 2010, p. 240)—provides precisely the theoretical resources the article needs to move from describing algorithmic bias to critiquing it. Similarly, Ricoeur's hermeneutics of suspicion, which approaches texts with the assumption that surface meanings conceal deeper, often ideological, truths, offers a framework for demystifying the operations of computational interpretation (Dole, 2020). Ricoeur identified Marx, Nietzsche, and Freud as "masters of suspicion" whose investigations furnish "a science of meaning, irreducible to the immediate consciousness of meaning" (Dole, 2020, p. 616). The article's framework would benefit enormously from engaging with this tradition, which provides a theoretically grounded alternative to the descriptive account of bias it currently offers.

Table 4: Contributions and Limitations of the Article

Dimension	Contribution	Limitation	Theoretical Resource for Improvement
Epistemological Framework	Articulates hermeneutic character of machine reading	Undertheorizes power and ideology	Habermas's emancipatory interest (Packer, 2010)
Simulated Hermeneutic	Distinguishes simulation from authentic interpretation	Does not theorize distributed agency	Synthetic hermeneutics (2026)
Empirical Evidence	Provides three case studies	Comparative study lacks construct validity	Larger, more diverse participant pools (Fish, 1980)
Critical Dimension	Acknowledges algorithmic bias	Does not move from description to critique	Ricoeur's hermeneutics of suspicion (Dole, 2020)

Table 4 systematically evaluates the article's contributions against its limitations while identifying theoretical resources capable of addressing each shortfall. The epistemological framework constitutes the article's strongest contribution, yet its undertheorization of power and ideology could be remedied through Habermas's emancipatory interest (Packer, 2010). The

simulated hermeneutic concept productively distinguishes machine simulation from authentic interpretation but requires engagement with distributed agency (Synthetic Hermeneutics, 2026). The empirical evidence, comprising three case studies, suffers from construct validity concerns in its comparative design—a limitation addressable through more diverse participant pools (Fish, 1980). Most critically, the article acknowledges algorithmic bias but stops short of critique, a gap Ricoeur's hermeneutics of suspicion could effectively bridge (Dole, 2020). Collectively, these resources point toward an emancipatory computational hermeneutics.

Third, the article's engagement with the material and institutional dimensions of computational hermeneutics is limited. The question of who has access to the computational resources necessary for machine reading and how the concentration of these resources shapes the production of literary knowledge is raised only in passing. Recent work on "democratising advanced AI methods for scholars without extensive computing resources" suggests that these questions are not peripheral but central to the epistemological claims the article advances (Kwok, 2026). Crawford's (2021) analysis of AI as a technology of extraction—from the minerals mined for hardware to the exploited labor behind "automated" services—demonstrates that the epistemological questions the article raises are inseparable from political-economic questions about who builds, owns, and benefits from AI systems. Crawford (2021) argues that AI is "neither artificial nor intelligent" but rather a product of "natural resources, fuel, human labor, infrastructures, logistics, histories and classifications" (p. 8). Training datasets, Crawford (2021) shows, "act as the building block from which algorithms learn to make predictions and classifications, set the parameters for how an AI 'sees the world,' and can help us understand the 'epistemic boundaries' of an AI system" (p. 134). The data gathering and labeling process is "neither benevolent nor solely a technical act but rather involves making political, theoretical, and value-based decisions about what (and who) is included and what is excluded" (Crawford, 2021, p. 134). These insights are directly relevant to the article's framework: if machine reading is hermeneutical, then the material conditions under which machine reading occurs—including the labor, resources, and power relations that constitute computational infrastructure—are constitutive of its interpretive outputs.




Figure 3: Contributions and Limitations of the Article

Figure 3 shows the maps the article's contributions against its limitations, identifying the theoretical resources, particularly Habermas's emancipatory interest (Packer, 2010) and Ricoeur's hermeneutics of suspicion (Dole, 2020) that could address the undertheorization of power and ideology. The dashed arrow at the bottom identifies the resulting gap: an emancipatory computational hermeneutics that integrates critical theory with computational practice.

4.3 Implications for Future Research

The article opens several productive avenues for future inquiry. The most pressing concerns the development of what might be called a "hermeneutics of computational evidence"—a framework for evaluating the interpretive claims generated by machine reading that is neither naively positivist nor dismissively humanist. The article gestures toward such a framework through its discussion of "interpretive lenses", but a fuller elaboration would require sustained engagement with questions of validity, reliability, and interpretive warrant in computational contexts. As Van Zundert (2016) cautions, quantified approaches frequently exhibit a "relative lack of explanatory power" and a tendency toward "confirmation bias," verifying existing answers rather than generating new questions (p. 335). A hermeneutics of computational evidence would need to address these limitations directly, specifying the conditions under which machine-generated interpretations can be warranted as knowledge claims rather than merely plausible outputs. Habermas's (Packer, 2010) insistence on a frame of reference external to the system under investigation offers one resource for such a framework.

A second avenue concerns the comparative study of human and machine interpretation across diverse cultural and linguistic contexts. The article's findings regarding "algorithmic positivity" and "cultural abstraction" (Mohr et al., 2026) suggest that machine reading may systematically reshape culturally specific narratives through dominant linguistic and moral conventions. Comparative research across literary traditions would clarify whether these patterns are universal or contingent on the specific architectures and training data of contemporary AI systems. Nie's (2025) work on Chinese literary criticism, which calls for "literary computational analytical criticism" integrating ethical and cross-cultural dimensions, exemplifies the kind of inquiry needed.

A third avenue concerns the pedagogical implications of computational hermeneutics. If, as the article suggests, machine reading constitutes a distinct form of interpretation, then literary education must prepare students to engage critically with machine-generated readings. This would require not only technical literacy but also the hermeneutic sophistication to recognize the assumptions and limitations embedded in computational interpretations. The article's brief discussion of this dimension could be productively extended through empirical research on classroom practices and student learning outcomes. Kwok's (2026) argument that algorithmic mediation should be treated as "an opportunity to make the labour of reading visible and rewarded again" (p. 1) provides a productive orientation for such pedagogical research.

V. Conclusion

"Computational Hermeneutics: Machine Reading and the Epistemology of Literary Interpretation" makes a significant contribution to the emerging interdisciplinary discourse on AI and literary interpretation. Its central insight that computational methods are not external to hermeneutics but constitute a distinctive hermeneutic practice are both theoretically compelling and practically consequential. The article's engagement with the epistemology of machine reading provides a vocabulary and framework for analyzing computational interpretation that moves beyond the stale binaries of positivism versus Interpretivism. As Van Zundert (2016) demonstrates, "a computational turn does not automatically imply a turn to empiricism and scientism, or a disregard for hermeneutic tradition" (p. 333); the article productively extends this insight by showing that machine reading embodies interpretive commitments worthy of sustained theoretical attention.

However, the article's impact is constrained by the unevenness of its empirical foundation, its limited engagement with non-Western contexts, and its underdeveloped treatment of power and ideology. These limitations do not diminish the value of its theoretical contributions, but they suggest that the project of computational hermeneutics remains incomplete. The article's call for a "fundamental dialogue between humanities and computer science about the status of hermeneutics as an imperative to humanities research" is well-taken and urgent. Realizing this dialogue will require not only theoretical sophistication but also methodological innovation, empirical rigor, and critical attention to the institutional and material conditions under which computational interpretation takes place.

References

- Agre, P. E. (1997). Toward a critical technical practice: Lessons learned in trying to reform AI. In G. C. Bowker, S. L. Star, W. Turner, & L. Gasser (Eds.), *Social science, technical systems, and cooperative work: Bridging the great divide* (pp. 131–157). Lawrence Erlbaum Associates. <https://doi.org/10.4324/9781315805849>
- Barta, W. (2024). Black boxes and knowledge: The epistemological dilemma of computational humanities. *International Review of Information Ethics*, 34(1), 1–12.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623. <https://doi.org/10.1145/3442188.3445922>
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623. <https://doi.org/10.1145/3442188.3445922>
- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- Dole, A. (2020). Reframing the masters of suspicion: Marx, Nietzsche, and Freud. *Journal of the American Academy of Religion*, 88(2), 615–618. <https://doi.org/10.1093/jaarel/lfz096>
- Fish, S. (1980). *Is there a text in this class? The authority of interpretive communities*. Harvard University Press.
- Jockers, M. L. (2013). *Macroanalysis: Digital methods and literary history*. University of Illinois Press.
- Kestemont, M. (2012). Le Moriaen et l'attribution de textes: Un problème de style? *Queeste: Tijdschrift over Middeleeuwse Letterkunde in de Nederlanden*, 19(2), 120–142.

- Kirschenbaum, M. (2014). What is "digital humanities," and why are they saying such terrible things about it? *differences*, 25(1), 46–63. <https://doi.org/10.1215/10407391-2399125>
- Kommers, C., Ahnert, R., Antoniak, M., Benetos, E., Benford, S., Bunz, M., Caramiaux, B., Concannon, S., & others. (2025). Generative AI as a cultural technology: A hermeneutic framework for evaluation. *Frontiers in Artificial Intelligence*, 8, 1753041. <https://doi.org/10.3389/frai.2025.1753041>
- Kwok, J. (2026). Retrieving context, re-centring interpretation: AI hermeneutics and the democratisation of reading [Webinar recording]. DARIAH-Campus. <https://hdl.handle.net/21.11159/019bbda8-86a1-738e-9e8a-7d7a15b159fb>
- Kwok, J. (2026). Retrieving context, re-centring interpretation: AI hermeneutics and the democratisation of reading. DARIAH-Campus Webinar Recording. <https://hdl.handle.net/21.11159/019bbda8-86a1-738e-9e8a-7d7a15b159fb>
- Lévy, P. (2024). Towards a pragmatic approach to computational hermeneutics. *Revue Intelligibilité du Numérique*, 6. <https://hal.science/hal-04927577>
- Mallery, J. C., Hurwitz, R., & Duffy, G. (1986). Hermeneutics: From textual explication to computer understanding? *Proceedings of the 1986 AAAI Workshop on Knowledge Representation*.
- Mohr, J. W., Wagner, R., & Breiger, R. (2026). Computational hermeneutics: Evaluating generative AI as a cultural technology. *Frontiers in Artificial Intelligence*, 9, Article 1753041. <https://doi.org/10.3389/frai.2026.1753041>
- Moretti, F. (2007). *Graphs, maps, trees: Abstract models for a literary history*. Verso.
- Nie, Z. (2025). AI reading and computational analytical criticism of literature. *Foreign Literature Studies*, 47(5), 22–32.
- Nie, Z. (2025). AI reading and computational analytical criticism of literature. *Foreign Literature Studies*, 47(5), 22–32.
- Packer, M. (2010). Emancipatory inquiry as rational reconstruction. In *The science of qualitative research* (pp. 235–256). Cambridge University Press. <https://doi.org/10.1017/CBO9780511779947.013>
- Phillips, D. Z. (2009). Hermeneutics and the philosophical future of religious studies. In *Religion and the hermeneutics of contemplation* (pp. 1–24). Cambridge University Press. <https://doi.org/10.1017/CBO9780511659348.002>
- Ramsay, S. (2011). *Reading machines: Toward an algorithmic criticism*. University of Illinois Press.
- Ramsay, S. (2011). *Reading machines: Toward an algorithmic criticism*. University of Illinois Press.
- Ridge, M. (2024, October). Hold steady: Centering the human in the hype of AI [Conference presentation]. Digikult 2024, Göteborg, Sweden. <https://doi.org/10.5281/zenodo.13880550>
- Ringler, H. (2022). "We can't read it all": Theorizing a hermeneutics for large-scale data in the humanities with a case study in stylometry. *Digital Scholarship in the Humanities*, 37(4), 1157–1172. <https://doi.org/10.1093/llc/fqab100>
- Schmidt, L. K. (2025). Gadamer's ontological turn towards language. In *Understanding hermeneutics* (pp. 1–15). Cambridge University Press. <https://doi.org/10.1017/CBO9780511617829.007>
- Scott-Baumann, A. (2011). *Ricoeur and the hermeneutics of suspicion*. Continuum.
- Synthetic hermeneutics: A protocol for reading literature in the age of generative AI. (2026). *International Journal of English Language, Literature and Translation Studies*, 13(1), 1–15.
- Van Dalen-Oskam, K. (2011). Stylometry and the Dutch Arthurian tradition: On the authorship of Moriaen. *Digital Humanities Quarterly*, 5(2).

- Van Dalen-Oskam, K., & Van Zundert, J. (2007). Authorship attribution and stylometry in the Dutch language area. *Literary and Linguistic Computing*, 22(3), 347–364. <https://doi.org/10.1093/lc/fqm018>
- Van Zundert, J. (2014). Screwmeneutics and hermenumericals: The computability of hermeneutics. Paper presented at the Digital Humanities Conference.
- Van Zundert, J. J. (2016). Screwmeneutics and hermenumericals: The computability of hermeneutics. In S. Schreibman, R. Siemens, & J. Unsworth (Eds.), *A new companion to digital humanities* (pp. 331–347). Wiley-Blackwell. <https://doi.org/10.1002/9781118680605.ch23>
- Verbeek, P.-P. (2006). Materializing morality: Design ethics and technological mediation. *Science, Technology, & Human Values*, 31(3), 361–380. <https://doi.org/10.1177/0162243905285847>