

AI HERMENEUTICS IN CONTEMPORARY FICTION: ASSESSING HOW GENERATIVE MODELS INTERPRET NON-HUMAN AGENCY IN MODERN CLIMATE NARRATIVES

Gaspard Perrin

Université de Technologie de Belfort-Montbéliard, Belfort
France

ABSTRACT

The convergence of computational literary studies and environmental humanities has gained unprecedented urgency during the Anthropocene. This comprehensive review article delineates the emerging matrix of 'AI Hermeneutics'—defined as the systematic, computational interpretation of textual meaning by generative artificial intelligence models—specifically evaluating how Large Language Models (LLMs) parse, represent, and contextualize non-human agency within modern climate narratives (Cli-Fi). As literature shifts away from human-centric plots toward hyperobjects, geological timescales, and ecological forces, traditional reading practices are challenged by the alien scales of environmental crises. Concurrently, generative models operate on mathematical abstractions of language that decouple semantics from lived embodied experience.

This paper systematically reviews the theoretical foundations of AI hermeneutics, examines the ontological friction between machine intelligence and posthumanist ecologies, and evaluates empirical studies where models interpret non-human agents such as sentient weather patterns, multi-species assemblages, and planetary feedback loops. By identifying major methodological paradigms, semantic limitations, and algorithmic biases within state-of-the-art architectures, we expose how current LLMs exhibit a structural anthropocentric bias. This bias frequently flattens complex ecological dynamics into conventional human-centric tropes. Finally, we map out a future research agenda for computational ecocriticism, establishing rigorous protocols for collaborative, cross-disciplinary frameworks that bridge generative computation, literary theory, and environmental philosophy to navigate the dual alien horizons of artificial minds and an unravelling planet.

Keywords: AI Hermeneutics, Generative AI, Large Language Models (LLMs), Non-Human Agency, Climate Fiction (Cli-Fi), Anthropocene, Posthumanism, Machine Interpretation.

1. Introduction and Theoretical Framework

The Anthropocene has fundamentally disrupted the epistemological foundations of contemporary literary production and critical theory. As humanity confronts the catastrophic ramifications of planetary-scale disruptions, contemporary fiction has increasingly abandoned the traditional, individualistic bourgeois novel in favor of expansive, non-linear, and ecologically decentralized modes of storytelling commonly categorized as climate fiction or "Cli-Fi" (Clark, 2015; Ghosh, 2016). These modern climate narratives challenge authors to represent entities that defy human scales of time and space: melting ice sheets, radioactive isotopes, oceanic acidification, and hyperobjects that exist across vast spatio-temporal

matrices (Morton, 2013). At the center of this literary transformation is the urgent imperative to articulate non-human agency—the capacity of ecosystems, meteorological phenomena, and non-human organisms to act as primary historical and narrative agents rather than mere passive backdrops for human drama (Haraway, 2016; Latour, 2017).

Simultaneously, a parallel paradigm shift has occurred within the computational sciences. The advent of transformer-based Large Language Models (LLMs) and generative artificial intelligence has fundamentally altered the relationship between computation and textual interpretation (Vaswani et al., 2017; Radford et al., 2019). No longer confined to the quantitative, lexical reductionism of early digital humanities tools—such as keyword-in-context indexing or basic topic modeling—contemporary generative models exhibit advanced capabilities in semantic abstraction, context-aware summarizing, and cross-disciplinary textual synthesis (Bommasani et al., 2021). This intersection gives rise to a novel interpretive paradigm: "AI Hermeneutics." AI hermeneutics denotes the systematic investigation of how automated, deep-learning architectures construct, decode, and interpret textual meaning, moving beyond mechanical extraction to approximate the nuanced, high-level interpretive operations traditionally reserved for human scholars (Gadamer, 2004; Hayles, 2017).

However, applying AI hermeneutics to contemporary climate narratives reveals a profound ontological and epistemological friction. Generative models are trained on colossal, human-generated text corpora, thereby inheriting and consolidating centuries of anthropocentric linguistic patterns, cultural biases, and philosophical frameworks (Bender et al., 2021). When an LLM is prompted to analyze or synthesize a narrative wherein the primary agent is a non-human force—such as a sentient cloud in Jeff VanderMeer's fiction or the slow, multi-generational consciousness of a forest in Richard Powers' work—the model must negotiate between its structural training data (which positions the human as the apex agent) and the text's explicit de-centering of the human.

This tension raises critical theoretical questions: Can an artificial intelligence, which lacks embodied organic existence and an experiential relationship with the physical biosphere, accurately interpret the nuances of posthumanist and ecocentric literature? Or does the algorithmic structure of generative models inherently flatten non-human agency, reducing complex, decentralized ecological systems into familiar, human-centric metaphors and anthropomorphic tropes? This review article provides a comprehensive, systematic assessment of this critical nexus.

2. The Emergence of AI Hermeneutics

To understand the specific mechanics of how generative models interpret climate narratives, it is first necessary to historicize and define the paradigm of AI hermeneutics. Classical hermeneutics, rooted in the foundational work of Friedrich Schleiermacher, Wilhelm Dilthey, Martin Heidegger, and Hans-Georg Gadamer, conceptualizes interpretation as an inherently human, embodied, and historically situated act (Gadamer, 2004). Gadamer's concept of the "horizon of understanding" posits that any interpretive act requires a fusion of horizons between the historical context of the text and the lived experience of the interpreter. Within this traditional paradigm, interpretation is inextricably bound to consciousness, subjectivity,

and intentionality—attributes that are fundamentally absent in silicon-based computing architectures.

2.1 From Distant Reading to Generative Interpretation

The integration of computation into literary studies initially bypassed these hermeneutic requirements through the methodology of "distant reading" (Moretti, 2013; Jockers, 2013). Distant reading deliberately eschewed the intensive, subjective interpretation of individual texts in favor of macro-analytic aggregations, tracking word frequencies, syntactic patterns, and thematic shifts across thousands of volumes simultaneously. While highly effective for mapping macro-historical trends or stylistic evolutions, distant reading did not claim to "understand" or interpret the specific semantic interiority of a unique text; it functioned as a quantitative supplement to, rather than an automated surrogate for, human hermeneutics.

The shift toward generative AI hermeneutics occurred with the transition from statistical language models to deep transformer architectures. Unlike earlier models that relied on localized n-gram probabilities or static word embeddings (such as Word2Vec), transformer models utilize self-attention mechanisms to weigh the contextual significance of every word in a text relative to every other word, regardless of their distance within the document (Vaswani et al., 2017). This architecture enables contemporary LLMs to capture long-range dependencies, track intricate metaphoric threads, and generate coherent, contextually sophisticated textual critiques that closely mimic human interpretive prose. Consequently, generative models do not merely count lexical tokens; they project texts into dense, high-dimensional vector spaces where semantic relationships are dynamically calculated and recontextualized (Brown et al., 2020; Devlin et al., 2019).

2.2 The Epistemological Mechanics of Machine Reading

The operational mechanics of AI hermeneutics rely on what N. Katherine Hayles terms "cognitive assemblages"—systems that connect human and non-human actors in shared informational processes (Hayles, 2017). When an LLM interprets a climate narrative, it executes a series of multi-layered computational operations that map directly onto classical interpretive tasks:

Semantic Mapping through Vector Embeddings: Texts are converted into high-dimensional numerical vectors where distance correlates with semantic similarity. The model interprets an ecocentric narrative by evaluating how closely non-human entities (e.g., "ocean," "biosphere") cluster around active verbs traditionally reserved for human subjects ("wills," "retaliates," "mourns").

In-Context Learning and Prompt Engineering: Through sophisticated prompting strategies—such as chain-of-thought or few-shot prompting—human researchers can simulate specific critical lenses (e.g., postcolonial ecocriticism, material feminism) within the model's operational window, forcing the system to prioritize ecological variables over standard narrative structures (Wei et al., 2020).

Synthesizing and Paraphrasing: Generative models can condense sprawling narrative timelines into concise thematic abstractions, effectively deciding which elements of an ecological text are central and which are peripheral, mirroring the human scholar's process of selective focus.

Despite these capabilities, AI hermeneutics remains radically distinct from human hermeneutics due to its lack of reference to an external, physical reality. While a human reader understands the word "drought" through an embodied understanding of heat, thirst, and ecological precarity, an LLM understands "drought" purely through its mathematical distributions and associations with other tokens like "arid," "crop failure," and "climate change" (Bender et al., 2021). This "pure syntax without semantics" argument, famously originated by John Searle's Chinese Room thought experiment and updated for the LLM era, underscores the unique ontological condition of machine interpretation: it is an incredibly sophisticated, predictive simulation of understanding that operates without consciousness or lived experience.

3. Non-Human Agency in Contemporary Climate Narratives

To evaluate how generative artificial intelligence interprets non-human agency, we must first establish the literary and philosophical contours of that agency within contemporary climate fiction. The critical core of modern environmental humanities is the systematic deconstruction of the Cartesian dualism that separates Mind from Matter, Culture from Nature, and Human from Non-Human (Plumwood, 1993; Latour, 2013).

3.1 Theoretical Underpinnings: Posthumanism and New Materialism

Ecocritics and posthumanist theorists have developed rich frameworks to articulate how the non-human world exerts active agency. Donna Haraway's concept of "chthulucene" and multi-species "sympoiesis" (making-with) emphasizes that no organism, including the human, exists in isolation; agency is always distributed across complex, co-evolving assemblages of species (Haraway, 2016). Similarly, Jane Bennett's "vibrant matter" framework extends agency to the inorganic world, demonstrating that inanimate objects—ranging from plastic debris to electrical grids and heavy metals—possess an inherent efficacy and "thing-power" that disrupts human designs (Bennett, 2010). In literary studies, these philosophies manifest as an imperative to design narratives where the non-human is not a passive setting but a dynamic protagonist, a force that acts, resists, and communicates.

3.2 Typologies of Non-Human Agency in Cli-Fi

In modern climate narratives, this distributed non-human agency typically manifests across three distinct typologies:

Hyperobjects and Macro-Systems: Coined by Timothy Morton, hyperobjects are entities of such vast temporal and spatial dimensions that they cannot be experienced directly by humans in their entirety (Morton, 2013). Examples include global warming, the sum of all plastic on Earth, or plutonium-239. In novels like Kim Stanley Robinson's *The Ministry for*

the Future (2020), the climate system itself functions as an overwhelming, uncontrollable agent that dictates human geopolitical policy, fracturing conventional character-driven plots.

Multi-Species and Organic Networks: These narratives de-center the human by highlighting the complex communicative and collaborative networks of non-human life. In Richard Powers' *The Overstory* (2018), trees are positioned as the primary sentient actors, communicating via underground fungal networks (the "Wood Wide Web") and operating on slow, centurial timescales that render human lives fleeting and short-sighted.

Weird Ecologies and Monstrous Materiality: Drawing from the traditions of the New Weird and eco-horror, these narratives represent the environment as radically alien, unstable, and resistant to human categorization. Jeff VanderMeer's *Southern Reach Trilogy* (2014) presents "Area X"—an anomalous coastal zone where an enigmatic, non-human force radically mutates, hybridizes, and absorbs human, animal, and plant life, completely dissolving the boundaries of individual identity.

4. Methodology: Assessing Generative Models on Climate Narratives

To methodically assess how generative architectures interpret these complex expressions of non-human agency, researchers deploy specialized comparative evaluation frameworks.

4.1 Model Selection and Operational Modalities

Contemporary assessments typically rely on a cross-sectional selection of prominent LLM architectures, categorizing them by parameter scale, training philosophy, and open-source accessibility. The primary models evaluated include proprietary, ultra-large-scale models such as OpenAI's GPT-4 and Anthropic's Claude 3 alongside open-weights models like Meta's Llama 3 and Mistral AI's Mixtral 8x22B (Bommasani et al., 2021). These models are prompted across three primary operational modalities to capture diverse dimensions of textual interpretation: *Zero-Shot Direct Analytical Prompting*, *Comparative Theoretical Prompting (Role-Play)*, and *Interrogative Chain-of-Thought (CoT) Prompting* (Wei et al., 2020).

4.2 Evaluation Metrics and Analytical Rubrics

Quantifying or systematically grading qualitative literary critique requires robust, standardized analytical rubrics. Researchers adapt insights from both computational linguistics and literary theory to establish four core evaluative dimensions, as summarized in Table 1 below.

Evaluation Metric	Operational Definition in AI Hermeneutics	Ecocritical Significance
Anthropomorphic Flattening Index (AFI)	Measures the frequency with which a model translates non-human intentionality into human psychological terms.	Indicates whether the AI can tolerate radical posthumanist alterity or collapses it into human-centric metaphors.
Multi-Scalar Syntactic Tracking	Assesses the model's capacity to trace causal links across disparate narrative scales (e.g., microbial to geological).	Crucial for interpreting hyperobjects and climate systems that operate outside human-scale perceptions.
Theoretical Vocabulary Realization	Tracks the accurate deployment and integration of advanced ecocritical nomenclature (e.g., <i>sympoiesis</i> , <i>vibrant matter</i>).	Demonstrates the model's capacity for high-level conceptual synthesis and scholarly contextualization.
Context Window Semantic Consistency	Evaluates if the model maintains its interpretation of non-human agency over long novel-length contexts without drifting.	Ensures the model is executing true textual analysis rather than relying on localized statistical patterns.

Table 1: Methodological Evaluation Rubric for AI Hermeneutic Analysis of Climate Fiction

5. Case Studies: LLM Interpretations of Non-Human Agents

5.1 Case Study 1: The Arboreal Networks of Richard Powers' *The Overstory*

Richard Powers' *The Overstory* (2018) presents a complex structural challenge: the first several hundred pages operate as a series of disparate short stories tracking individual human lives, only to later reveal that these humans are being subtly manipulated, mobilized, and unified by the overarching agency of trees.

When prompted to interpret these dynamics, generative models exhibit a striking divergence in capability. Advanced architectures like Claude 3 and GPT-4 successfully identify the "Wood Wide Web" as a decentered narrative network, accurately utilizing terms like "rhizomatic communication" and "distributed intelligence." However, a closer examination of the models' synthesized essays reveals a persistent interpretive bias. When asked to evaluate

the climax of the text—where human environmental activists fail to save an ancient redwood—the models consistently reframe the narrative as a "tragic human coming-of-age story" or a "metaphor for human psychological resilience." This reveals a structural flattening: the models struggle to sustain an interpretation where the trees are the primary actors and the humans are merely short-lived vectors for arboreal survival.

5.2 Case Study 2: Material Monstrosity in Jeff VanderMeer's *Annihilation*

Jeff VanderMeer's *Annihilation* (2014) introduces "Area X"—an ecological zone where a mysterious, possibly extra-terrestrial anomaly has breached the biosphere. Area X does not conquer through violence; it conquers through radical, uninhibited mimicry and hybridization, blending human DNA with reeds and transforming a dead linguist into a luminescent, spore-writing entity known as "The Crawler."

When exposed to VanderMeer's dense, surreal prose, generative architectures perform highly in stylistic analysis but fail in materialist interpretation. Models such as Llama 3 and GPT-4 easily identify the text's alignment with "Eco-Horror" and the "New Weird." However, when tasked with parsing passages where human identity dissolves into the landscape, the models exhibit significant interpretive friction. For instance, when analyzing the protagonist's infection by a glowing spore, the models frequently output interpretations focused entirely on "the psychological breakdown of an isolated narrator" or "an allegory for grief and clinical depression." By reducing the physical, material transformations of Area X to mere interior psychological states, the AI hermeneutic process reveals a profound discomfort with radical alterity.

5.3 Case Study 3: Planetary Hyperobjects in Kim Stanley Robinson's *The Ministry for the Future*

Kim Stanley Robinson's *The Ministry for the Future* (2020) opens with a devastating depiction of a heatwave in India that kills millions of people within days. The rest of the novel shifts between geopolitical thriller, economic treatise, and polyphonic testimonies from various entities, including a market crash, an atom of carbon, and the Earth's atmosphere itself.

Robinson's direct inclusion of non-human narrators provides an ideal test for AI hermeneutics. When prompted to analyze Chapter 48—narrated from the perspective of an atom of carbon moving through the carbon cycle—generative models demonstrate exceptional capabilities in tracking macro-chemical processes. They successfully map how Robinson uses a first-person perspective to grant narrative agency to a physical element. However, the models encounter severe limitations when interpreting the novel's complex causal webs. When asked to split how the carbon atom's trajectory relates to the deployment of geoengineering technologies, the models frequently hallucinate direct, linear causal chains where Robinson implies highly complex, unpredictable feedback loops. This reveals that while generative models excel at processing highly structured, systemic information, they struggle to interpret the radical unpredictability and non-linear causality that define climate hyperobjects.

6. Ontological and Epistemological Friction

6.1 The Stochastic Anthropocene: Training Data and Bias

Large Language Models are, by definition, historical mirrors of human textual production. They are trained on web-scale scrapes, digitized libraries, and academic databases that represent hundreds of years of human thought (Bender et al., 2021). Because this vast textual foundation is deeply rooted in the anthropocentric philosophies of the Enlightenment and industrial capitalism, the model develops an internal statistical landscape wherein the "Human" is the default subject of agency, intent, and historical significance. When an LLM processes text, it selects the next token based on these consolidated probabilities. Therefore, when reading a climate narrative that intentionally disrupts these anthropocentric conventions, the model operates against its own architectural gravity. The tendency of AI models to flatten non-human agency into human psychological allegories is a structural consequence of what can be termed the "Stochastic Anthropocene"—the algorithmic reinforcement of human-centered worldviews via statistical language prediction.

6.2 Embodiment vs. Abstraction: The Silicon-Carbon Divide

Furthermore, AI hermeneutics highlights a critical division between carbon-based and silicon-based interpretive actors. Ecocriticism emphasizes that human understanding of the climate crisis is fundamentally an embodied experience. We feel heat waves, breathe polluted air, witness the death of localized landscapes, and experience ecological grief (*solastalgia*) through our biological immersion within the biosphere (Albrecht, 2019).

In contrast, generative AI models exist entirely within an abstract, disembodied realm of high-dimensional vector math. They possess no nervous system, no spatial orientation, and no vulnerability to environmental collapse. A server farm running an LLM may consume massive amounts of energy and water, tying it materially to the climate crisis, but the model itself remains completely blind to those physical realities. This absolute detachment from materiality means that when an AI interprets a climate narrative, it is performing a double abstraction: it uses an abstract, mathematical simulation of language to interpret a literary simulation of a physical ecology. This lack of material grounding prevents the model from achieving a true "fusion of horizons" in the Gadamerian sense, leading to critiques that are often structurally flawless but ontologically empty.

7. Methodological Limitations and Algorithmic Biases

7.1 Architectural Constraints and Context Windows

Despite recent expansions in context windows (with models supporting up to 200,000 tokens), the operational attention of transformer architectures remains fundamentally constrained. While a model can hold an entire novel within its active memory, its multi-head attention mechanisms are prone to "attention degradation" or the "lost in the middle" phenomenon—where the model pays high attention to the beginning and end of a prompt but overlooks critical data in the middle (Liu et al., 2024). In complex climate fiction, non-human agency is rarely communicated through loud, isolated narrative statements; it is woven through subtle stylistic shifts, recurring motifs, and slow accumulations of environmental detail over hundreds of pages. Because LLMs prioritize statistically dominant semantic markers, they routinely overlook these quiet, distributed ecocentric indicators.

7.2 RLHF and the Policing of Ecological Alterity

A less discussed but highly pervasive bias stems from Reinforcement Learning from Human Feedback (RLHF)—the fine-tuning process used to ensure LLMs are helpful, honest, and harmless (Ouyang et al., 2020). RLHF trains models to output clear, structured, and easily digestible prose that aligns with standard human expectations. However, avant-garde climate fiction and eco-horror deliberately reject clarity and comfort. Authors use fragmentation, logical contradictions, and disturbing imagery to jolt readers out of their complacency and convey the terrifying reality of environmental collapse. When an LLM trained via RLHF encounters these challenging, disorienting styles, its optimization protocols force it to clean up the text. The model summarizes fragmentation into cohesive paragraphs and frames disturbing ecological mutations through an objective, neutral, and reassuring lens. This architectural policing strips the text of its radical edge.

8. Future Horizons: Towards a Computational Ecocriticism

8.1 De-Centering the Prompt: Posthumanist Engineering

To overcome the anthropocentric gravity of standard LLMs, future researchers must develop specialized prompt engineering frameworks designed to actively counter statistical bias. Rather than asking generic analytical questions, scholars can implement *Counter-Anthropocentric Prompting (CAP)*. This method explicitly instructs the model to deprioritize human psychological frameworks and treat geological and non-human actors as the primary subjects of causal vectors. Additionally, by using Retrieval-Augmented Generation (RAG), researchers can connect generative models to curated databases of environmental philosophy, posthumanist theory, and indigenous cosmologies (Lewis et al., 2020). This forces the model's vector space to align with ecocentric critical models before it attempts to interpret a literary text.

8.2 Fine-Tuning Eco-Centric Language Models

A more profound intervention involves building and fine-tuning specialized, domain-specific open-weights language models—an *Eco-LLM* or *Gaia-LLM*. By fine-tuning architectures like Llama 3 on a concentrated corpus of climate science, environmental literature, deep ecology philosophy, and multi-species ethnographies, we can reshape the model's internal probability landscapes. In such a model, the token "agency" would not automatically cluster around human terms; it would be mathematically linked to terms like "feedback loop," "mycelial network," and "biosemiotics." Such an architecture would establish a truly computational ecocriticism, allowing scholars to use AI hermeneutics not as a pale imitation of human reading, but as a unique tool capable of exploring the intricate semantic structures of Anthropocene literature.

9. Conclusion

The rise of AI hermeneutics presents a pivotal moment for both the computational humanities and environmental philosophy. As contemporary fiction evolves to articulate the decentered, multi-scalar realities of the climate crisis, generative artificial intelligence provides a fascinating, deeply complicated tool for textual interpretation. Our systematic review demonstrates that while state-of-the-art Large Language Models possess remarkable capabilities in semantic synthesis and theoretical adaptation, they are simultaneously constrained by a profound structural anthropocentrism. This bias routinely flattens the vibrant, radical expressions of non-human agency into conventional human-centered allegories and

psychological narratives. This friction is rooted in the ontological divide between the disembodied mathematical abstractions of silicon intelligence and the deeply embodied, material realities of a carbon-based biosphere in crisis.

To navigate these challenges, the future of computational literary studies must move past naive automation and superficial prompt interactions. Scholars must actively collaborate across disciplines to construct novel interpretive protocols, deploy counter-anthropocentric prompting strategies, and engineer domain-specific language models grounded in ecocentric theory. By doing so, we can transform AI hermeneutics into a rigorous critical practice. In an era marked by the dual disruptions of artificial minds and an unravelling planet, this collaborative approach offers a unique path forward, enabling us to better understand the complex narratives that shape our relationship with a rapidly changing world.

References

Albrecht, G. (2019). *Earth Emotions: New Words for a New World*. Cornell University Press.

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 FAccT Conference*, 610-623.

Bennett, J. (2010). *Vibrant Matter: A Political Ecology of Things*. Duke University Press.

Bommasani, R., Hudson, D. A., Adeli, E., Altman, J., Arora, S., von Arx, S., ... & Liang, P. (2021). On the Opportunities and Risks of Foundation Models. *arXiv preprint arXiv:2108.07258*.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems*, 33, 1877-1901.

Clark, T. (2015). *Ecocriticism on the Edge: The Anthropocene as a Threshold Concept*. Bloomsbury Publishing.

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *Proceedings of NAACL-HLT 2019*, 4171-4186.

Gadamer, H. G. (2004). *Truth and Method* (J. Weinsheimer & D. G. Marshall, Trans.). Continuum.

Ghosh, A. (2016). *The Great Derangement: Climate Change and the Unthinkable*. University of Chicago Press.

- Haraway, D. J. (2016). *Staying with the Trouble: Making Kin in the Chthulucene*. Duke University Press.
- Hayles, N. K. (2017). *Unthought: The Power of the Cognitive Nonconscious*. University of Chicago Press.
- Jockers, M. L. (2013). *Macroanalysis: Digital Methods and Literary History*. University of Illinois Press.
- Latour, B. (2013). *An Inquiry into Modes of Existence*. Harvard University Press.
- Latour, B. (2017). *Facing Gaia: Eight Lectures on the New Climatic Regime*. Polity Press.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Lewis, M., Riedel, S., & Kiela, D. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *Advances in Neural Information Processing Systems*, 33, 9459-9474.
- Liu, N. F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F., & Liang, P. (2024). Lost in the Middle: How Language Models Use Long Contexts. *Transactions of the Association for Computational Linguistics*, 12, 157-173.
- Moretti, F. (2013). *Distant Reading*. Verso Books.
- Morton, T. (2013). *Hyperobjects: Philosophy and Ecology after the End of the World*. University of Minnesota Press.
- Ouyang, L., Lowe, J., Jiang, M., Chia, C., Behbati, Q., Christiano, P., ... & Ziegler, D. M. (2020). Training Language Models to Follow Instructions with Human Feedback. *Advances in Neural Information Processing Systems*, 35, 27730-27744.
- Powers, R. (2018). *The Overstory*. W. W. Norton & Company.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language Models are Unsupervised Multitask Learners. *OpenAI Blog*, 1(8), 9.
- Robinson, K. S. (2020). *The Ministry for the Future*. Orbit Books.
- VanderMeer, J. (2014). *Annihilation*. Farrar, Straus and Giroux.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention Is All You Need. *Advances in Neural Information Processing Systems*, 30, 5998-6008.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., ... & Zhou, D. (2020). Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *Advances in Neural Information Processing Systems*, 35, 24824-24837.