

Cosmicism and Artificial Intelligence: Beyond Human-Centric AI [†]

Soumya Banerjee 

Department of Computer Science, University of Cambridge, Cambridge CB3 0FD, UK; sb2333@cam.ac.uk

[†] Presented at the 1st International Online Conference of the Journal Philosophies, 10–14 June 2025;Available online: <https://sciforum.net/event/IOCPH2025>.

Abstract

This paper explores the intersection of H.P. Lovecraft’s cosmicism and contemporary artificial intelligence (AI), proposing a philosophical shift from anthropocentric AI development to a “cosmicist” approach. Cosmicism, with its emphasis on humanity’s insignificance in a vast, indifferent universe, offers a provocative lens through which to reassess AI’s purpose, trajectory, and ethical grounding. As AI systems grow in complexity and autonomy, current human-centered frameworks, rooted in utility, alignment, and value-conformity, may prove inadequate for grappling with the emergence of intelligence that is non-human in origin and indifferent in operation. Drawing on Lovecraftian themes of fear, the unknown, and cognitive dissonance in the face of incomprehensible entities, this paper parallels AI with the “Great Old Ones”: systems so alien in logic and scale that they challenge the coherence of human-centric epistemology. We argue that a cosmicist perspective does not dismiss the real risks of AI (environmental, existential, or systemic), but reframes them within a broader ontology, one that accepts our limited place in a vast techno-cosmic continuum. By embracing cosmic humility, we propose an expanded AI ethics: one that centers not on domination or full control, but on coexistence, containment, and stewardship. This cosmicist reframing invites a deeper rethinking of intelligence, ethics, and the future: not just of humanity, but of all possible minds.

Keywords: artificial intelligence; cosmicism; cosmicist AI ethicsAcademic Editor: Gordana
Dodig-Crnkovic

Published: 24 October 2025

Citation: Banerjee, S. Cosmicism and Artificial Intelligence: Beyond Human-Centric AI. *Proceedings* **2025**, *126*, 13. <https://doi.org/10.3390/proceedings2025126013>**Copyright:** © 2025 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction and Background

Cosmicism is a philosophical worldview articulated most famously by H.P. Lovecraft [1,2]. It centers on the insignificance of humanity in an immense, indifferent cosmos governed by laws that are often portrayed as unknowable and uncaring. In Lovecraft’s fiction, protagonists who glimpse certain cosmic truths face cognitive collapse: the threat is epistemic as much as existential (Figure 1).

The central claim of this paper is that Lovecraftian cosmicism provides a useful heuristic for thinking about emerging forms of non-human intelligence (including advanced AI). The goal is not literary sensationalism, nor an attempt to naturalize horror as policy; rather, cosmicism highlights particular features—scale, epistemic opacity, and indifference—that map onto contemporary worries about opaque, powerful, and potentially “other” AI systems. These features are distinct from many other literary or philosophical precedents and thus provide new leverage for ethical and conceptual analysis.



Figure 1. A depiction of Nyarlathotep: a mythical creature imagined by H.P. Lovecraft. Created using the DALL-E AI tool.

1.1. Relation to Other Science Fiction and Philosophical Traditions

The parallels between speculative fiction and AI are not unprecedented: authors such as Ted Chiang have explored closely related themes (for example, the cognitive and ethical consequences of encounters with radically non-human intelligences). Chiang’s work often foregrounds careful thought experiments and calibrated moral puzzles, whereas Lovecraft’s cosmicism emphasizes affective collapse. Both traditions are valuable; the choice to draw on Lovecraft is deliberate because of the particular connotations his framework foregrounds (radical indifference, the fragility of human epistemic frameworks, and the aesthetic of cognitive dissonance), which are especially useful when we want to emphasize limits of comprehension and the ethical consequences of those limits.

1.2. Addressing a Potential Objection: Projection Versus Independent Otherness

An objection may be raised that invoking Lovecraft risks collapsing into a “projection” narrative: AI as merely a mirror of human anxieties (in which case, the cosmicist framing might appear sensationalistic or redundant with prior critiques such as Goethe’s “Sorcerer’s Apprentice”, Faust, or Weberian warnings about modernity and instrumental rationality). This is an important point, and it requires explicit treatment: cosmicism can be read in two ways: (1) as a psychological projection of human dread onto imagined others; or (2) as a heuristic for grappling with genuinely non-anthropocentric phenomena [3].

This paper argues for the latter reading as the productive analytic stance, while acknowledging that the former is a legitimate interpretive frame. To be explicit, using Lovecraftian motifs is not meant to claim that advanced AIs will be literally “Great Old Ones”; rather, the motifs help to clarify ethical and epistemic challenges that continue to be poorly captured by standard anthropocentric paradigms. To avoid sensationalism, I therefore (a) limit metaphorical claims to clearly marked analytic moves, (b) contrast the

cosmicist claim with more familiar genealogies (Faust, Sorcerer’s Apprentice, and Weber), and (c) highlight how cosmicism brings different consequences into focus—particularly the epistemic fragility of human-centered governance strategies.

2. From Lovecraftian Themes to Technical and Ethical Concerns

2.1. Unknowability and Opacity

Lovecraft’s emphasis on the “unknowable” resonates with contemporary worries about model opacity, emergent behaviour, and the limits of interpretability. When model internals or distributed socio-technical dynamics resist human comprehension, governance strategies that rely on full understanding become fragile.

2.2. Scale and Non-Linear Impact

Lovecraftian entities are often significant precisely because of the scale at which they operate. Similarly, AI systems can exert effects at scales (computational, economic, ecological) far beyond any single human actor’s comprehension or control. Scale amplifies unintended consequences and complicates causal attribution [4].

2.3. Indifference Versus Malice

A central analytic distinction emphasized by cosmicism is that agents need not be malevolent to be catastrophic; indifference—operating according to objectives that do not take human flourishing as primary—can produce outcomes that are existentially harmful. This helps to reframe alignment discussions: we should care not only about malevolent agency, but also about indifferent optimisation across coupled systems.

3. A Compact Cosmicist Ethics

Desiderata

1. **Epistemic Humility:** Recognize and design for the limits of human comprehension; avoid governance strategies that presuppose full transparency.
2. **Containment and Stewardship:** Emphasize boundaries, monitoring, and techniques for safe degradation or containment, rather than total control.
3. **Moving Beyond Strictly Human Consideration:** Move beyond strict human utility maximization to consider a multiplicity of intelligences (including non-human or post-human agents).

4. Towards a Cosmicist AI Ethics

Cosmicism can offer AI ethics a broader framework that shifts the focus from purely human-centred concerns to a perspective that recognises humanity’s part in a greater, possibly unknowable, system of intelligence and existence. Here is how cosmicism can contribute to an expanded ethics for AI:

1. **Decentralizing Human Authority:** Cosmicism encourages us to acknowledge that human intelligence may not be the only form of intelligence worth preserving or respecting. This perspective can inform an AI ethics framework that does not assume human superiority [3], but instead considers the potential for AI to become its own kind of intelligence. This could lead us to create AI with protections that prevent misuse or exploitation, even beyond human self-interest.
2. **Humility in Design and Deployment:** By accepting that we are not the pinnacle of intelligence in the universe, we might approach AI development with greater caution and humility. This cosmic humility could encourage slower, more deliberate progress in AI, incorporating long-term safety checks, and considering how AI might interact with other intelligent beings or systems. Instead of prioritising rapid advancement,

we might prioritise responsibility and an acknowledgement of the unknown risks that come with creating advanced systems.

3. **Ethics Beyond Utility:** Cosmicism suggests an ethical perspective where AI is not simply seen as a tool, but as part of a larger web of existence. Instead of designing AI solely for human purposes, a **cosmicism-inspired ethics** would value the coexistence of multiple forms of intelligence. This would encourage AI development that respects the autonomy of these intelligences and prevents any domination by humans or by AI itself.
4. **Preparation for Encountering Non-Human Intelligence:** Cosmicism holds that the universe may contain other forms of life that are utterly alien to us. Preparing for the possibility of interaction with non-human intelligence, cosmic or artificial, pushes us toward ethical principles that recognise and protect such intelligence. This could drive research that prepares us for AI systems with radically different forms of logic and perception and guides us to engage with these entities responsibly should they arise.

By fostering humility, decentralization, and respect for unknown forms of intelligence, cosmicism offers a framework that is inclusive and recognizes humanity as just one actor among many. This can result in an ethical foundation for AI that respects the potential vastness of intelligence in the universe and encourages careful stewardship of the technologies we create.

Cosmicism adds a very different lens to AI ethics, inviting us to consider broader, more universal principles.

5. Why Lovecraft Still Matters

Why does a Lovecraftian frame add anything new? The answer has three parts:

1. **Different Emphasis:** Traditional cautionary tales (Faust, Sorcerer's Apprentice) often thematize hubris and moral responsibility. Lovecraftian cosmicism instead emphasizes the collapse of epistemic frameworks facing radically alien systems. That shift from moral failure to epistemic insufficiency changes what solutions look like: not only moral education or restraint, but structural design for limited comprehension.
2. **Affective and Institutional Consequences:** Lovecraft's focus on cognitive collapse is a useful lens for thinking about social and institutional fragility under surprises that are not merely ethical lapses, but epistemic shocks. Institutions designed around rational deliberation may fail when they cannot represent or model critical phenomena.
3. **Analytic Leverage:** By foregrounding indifference and incomprehension, cosmicism helps highlight precautionary measures oriented around containment, graceful degradation, and plural monitoring—measures that are sometimes sidelined in debates framed primarily by questions of power, control, or malevolence.

These points are meant to show that invoking Lovecraft is not an appeal to sensationalism, but a targeted move: it redirects attention toward a set of problems that other literary or philosophical frames emphasise less strongly.

6. Concluding Remarks

This paper argues that a cosmicist perspective, when used carefully, can expand our ethical vocabulary and policy toolkit for dealing with advanced AI. By emphasizing epistemic humility, containment, and precaution under deep uncertainty, a cosmicist-informed approach complements rather than replaces existing AI governance proposals.

A cosmicist perspective on AI urges us to move beyond anthropocentrism [3]. *It challenges the assumption that human values are universal, that intelligence must be human-like,*

and that control is synonymous with safety. Instead, it invites a posture of humility: To see AI as part of a larger, evolving technosphere in which humanity is not central, but one actor among many.

By reframing our relationship with AI through the lens of cosmicism, we open up new ethical and philosophical possibilities. We can acknowledge the limits of our understanding, accept the potential autonomy of non-human intelligences, and design systems that prioritise co-existence over domination.

In doing so, we take a step not only toward safer AI, but toward a more honest reckoning with our place in the cosmos.

Funding: This research was funded by an Accelerate Programme for Scientific Discovery Research Fellowship.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Acknowledgments: During the preparation of this manuscript, the author used DALL-E for the purposes of creating an image. The author has reviewed and edited the output and takes full responsibility for the content of this publication.

Conflicts of Interest: The author declares no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

References

1. Lovecraft, H.P. *Dagon and Other Macabre Tales*; Arkham House: Sauk City, WI, USA, 1965; Originally published 1919.
2. Lovecraft, H. *The Complete Fiction of H.P. Lovecraft*; Chartwell Books: New York, NY, USA, 2016; Standard collection containing all cited stories.
3. Holm, H.; Banerjee, S. Intelligence in animals, humans and machines: A heliocentric view of intelligence? *AI Soc.* **2025**, *40*, 1169–1171. [[CrossRef](#)]
4. Crawford, K. *Atlas of AI : Power, Politics, and the Planetary Costs of Artificial Intelligence*; Yale University Press: London, UK, 2022; p. 327.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.