

Book Review

Governance Under Pressure: *Empire of AI* as an Institutional Case Study

Kshitij Taluja

Undergraduate Researcher, Bachelor of Computer Science, Golden Gate University

Dr. Smrite Goudhaman

Global Professor of Practice, Golden Gate University

Author: Karen Hao

Title: *Empire of AI: Dreams and Nightmares in Sam Altman's OpenAI*

Publisher: Penguin Press

Year: 2025

ISBN: 978-0-593-65750-8

Pages: 496

Abstract

Karen Hao's *Empire of AI* claims that the most influential artificial intelligence (AI) labs function as empires, using their size to extract labor, energy, land and data from people who reap very few benefits. This article reviews *Empire of AI* as a case study of how institutions govern missions when they are under pressure from capital, instead of looking at the empire model that it uses. The review's strongest contribution is its attention to the documentary evidence presented in the book. In particular, the book reconstructs the governance structure of OpenAI in detail through interviews with approximately 260 different people and a review of internal documents, and reports on the outsourcing of data labor in Kenya and water and energy usage in Chile and Uruguay. However, the principal conceptual limitation is that the imperial model locates where damage occurs more precisely than it explains how that damage might be addressed. As such, the main contribution of the book is documentation of the authority structures, dependency on capital and outsourced risks.

Keywords: artificial intelligence governance; nonprofit accountability; data labor; environmental externalities; Global South

Karen Hao's *Empire of AI* aims to provide a material accounting of what developing generative AI costs and who pays for it, replacing the technology industry's perpetual narrative of inevitable advancement. The author's organizing claim is that dominant AI laboratories such as OpenAI operate as empires, returning very little to the populations whose resources they draw upon. This review will evaluate *Empire of AI* as an institutional case study of mission-driven governance under capital pressure, and will not evaluate the imperial framework as its

primary analytical focus. We are stating our perspective openly since it will influence every part of this review.

Hao (2025) drew upon more than three hundred interviews conducted with approximately 260 individuals, including more than ninety current or former OpenAI employees. She set out a verification standard uncommon in trade publishing in her author's note: scenes, figures and technical details, verified through a minimum of two sources or documented evidence, and using paraphrase in lieu of quotations whenever she lacked access to the original. The importance of assessing the book lies in how its contested claims can be evaluated as reporting versus assertions.

Governance is where the book does best. Hao has documented OpenAI's transition from a nonprofit structure to a capped-profit limited partnership under nonprofit governance, an arrangement intended to keep investors' returns secondary to the organization's mission. She shows that the capital requirements for large-scale training eventually became so high that the nonprofit model could no longer support them. In addition, she reported on the events surrounding the November 2023 firing and subsequent reinstatement of CEO Sam Altman. As a reporter, Hao found out that the directors had been told by at least seven senior employees why he should be removed. This is what she has discovered through her reporting, and these accounts are treated in this review as reported rather than independently verified. She documented that there was a scope-restricted independent review of how the board reached this decision. However, the results of this review were never published. Hao concluded that the organizational authority granted to the nonprofit had effectively been hollowed out through commercial dependency. In our reading, this example illustrates a separation between formal and operational power. Those formally empowered to direct an organization may be unable to exercise that authority when they do not control the resources on which the organization depends.

As Hao describes, Kenyan employees working for Sama were paid on average less than two dollars an hour to manually classify violent or sexual content so their clients' AI models would recognize them as such. In both Chile and Uruguay, she traces water and electrical power consumption by AI data centers. She states that prior to rationing residential supplies due to drought, Uruguay was the first country in the world to have enshrined the human right to access water within its constitution. Her reportage regarding water usage relies upon independent estimates of AI's global freshwater footprint (Li et al., 2023), developed independently of her

reporting. These examples are illustrative of this broader issue of distributive justice, but should not necessarily be considered indicative of how all countries of the Global South operate labor market systems, regulate industries and manage resources. When viewed in this light, they can also be associated with the decent work and clean water and sanitation aspects of Sustainable Development Goals 8 and 6 (United Nations, 2015). The SDG framing is useful here because it locates these cases within commitments the institutions deploying these systems have already endorsed, rather than treating them as isolated harms.

Although Hao provides good evidence to support her argument, the way she frames that argument, the imperial approach, has a limitation. That limitation is not that the book fails to supply institutional remedies, since it never sets out to do so. It lies instead in the conceptual framework itself. Hao is cautious in how she uses the framework, stating in the introduction that AI companies do not engage in the same type of violent behavior associated with historical colonialism, and she draws upon existing research in developing her framework (Mohamed et al., 2020). One of the most successful ways that Hao's use of the framework illustrates continuity is how data, labor and resource extraction all follow familiar historical patterns. Gray (2023) argues that while these types of frameworks may illustrate continuity of pattern, they can give insufficient attention to differences in mechanism. Historical colonial rule operated in part through force, law and claims of sovereignty, though its forms varied considerably across periods and territories. The firms Hao discusses, by contrast, exercise influence over people located far from the decision makers largely through contracts, procurement and jurisdictional arbitrage. Any attempt at providing a remedy would need to occur through these mechanisms. A framework that emphasizes continuity of pattern over difference of mechanism therefore identifies where harm occurs more readily than how harm may be remedied.

For those interested in organizations as opposed to AI specifically, this book has relevance. The book demonstrates the way that organizations suffer from mission drift as a result of being financially dependent on capital; the limitations of board oversight when directors have control over none of the resources needed by an organization; and how the use of outsourcing can transfer both human and reputational risk to jurisdictions with less stringent labor protections. There are several practical implications for practitioners and action researchers. Practitioners need to test governance structures using resources they actually control as opposed to the nominal authority that a structure may have. To hold stakeholders accountable, practitioners need to understand where the labor behind a system was performed and under what conditions.

Procurement processes seldom bring these issues to the fore. Adopting responsible technologies also requires understanding what dependencies an organization creates when one supplier controls a capability the organization cannot build itself. Finally, there are implications for a reflective institutional process informed by action research principles. This process involves identifying how an organizational mission is changed by the resources necessary to achieve that mission.

These issues relate specifically to universities that currently source, distribute and normalize generative AI systems; therefore they also take on the governance problems Hao identifies. The language or languages in which a system operates determine which student populations are able to use it. Hao noted that while there are many thousands of the world's living languages, very few have been incorporated into the current generation of these tools, drawing on testing conducted by the developers themselves. In addition, the decision about how and where an institution stores its data, and which country or state has authority over that information, affects an institution's freedom from interference or regulation. Institutions may lose control of the research materials they hold once those materials cross international borders. Universities may become dependent upon vendors for services that they could not provide themselves, thereby creating a similar dependency to the one described by Hao at OpenAI.

The documentary accomplishment of this book as a whole transcends the controversy over its framing. Hao sets down a record of how an organization governed by a mission was altered by the requirements of its own mission, and who bore the cost of those alterations. The imperial argument will likely continue to be debated; it should be. The book will be particularly valuable for researchers of AI governance, organizational leaders, policymakers, university leaders, procurement professionals, and practitioners concerned with responsible technology adoption. The institutional evidence upon which this argument is based will be consulted by boards, procurement personnel, and researchers long after the debate is resolved.

References

- Gray, C. (2023). More than extraction: Rethinking data's colonial political economy. *International Political Sociology*, 17(2), olad007. <https://doi.org/10.1093/ips/olad007>
- Hao, K. (2025). *Empire of AI: Dreams and nightmares in Sam Altman's OpenAI*. Penguin Press.

- Li, P., Yang, J., Islam, M. A., & Ren, S. (2023). *Making AI less “thirsty”: Uncovering and addressing the secret water footprint of AI models* (arXiv:2304.03271). arXiv. <https://doi.org/10.48550/arXiv.2304.03271>
- Mohamed, S., Png, M.-T., & Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33(4), 659–684. <https://doi.org/10.1007/s13347-020-00405-8>
- United Nations. (2015). *Transforming our world: The 2030 agenda for sustainable development* (A/RES/70/1). <https://sdgs.un.org/2030agenda>