

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 book review examines *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. This review therefore concentrates on the documentary evidence the book presents. 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. One conceptual limitation of the imperial model is that it locates where harm occurs more precisely than it explains how that harm 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 physical cost accounting of the price of developing generative AI and who will pay for it, replacing the technology industry's perpetual narrative of inevitable advancement. The author explains how dominant AI laboratories such as OpenAI operate as empires that extract data, labor, land, and energy from the masses while returning very little to the people from whom these resources are drawn. Her major contribution is documentary; she describes OpenAI's governance structure, its labor outsourcing arrangements, and the environmental impact of large-scale training through interviews and internal documents. This review will evaluate *Empire of AI* as an institutional case study of

mission-driven governance when 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 documents the transition of OpenAI from a nonprofit status to a capped-profit limited partnership under nonprofit governance with the intent that investor profits would remain subordinate to the organization's mission. She further illustrates that the capital required for large-scale training became so high that the nonprofit model was unable to support it. In addition, she reported on the events surrounding the November 2023 firing and subsequent reinstatement of CEO Sam Altman. The book provides evidence that the directors were consistently told by at least seven senior members of staff about various reasons why they should remove him. Furthermore, Hao shows that some board members were aware of concerns surrounding Altman's family situation, but their focus ultimately shifted toward issues they could address institutionally, particularly governance processes and structures. 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 averaging less than two dollars an hour to manually classify violent or sexual content so their clients' AI models would recognize such content. She follows the bulk of the water and electricity consumed by AI data centers in both Chile and Uruguay. She notes that, two decades after Uruguay became the first country in the world to recognize water as a human right in its constitution, the country

experienced severe cutbacks to drinking water supplies during a drought, while industry continued to use water unabated. Her reportage regarding water usage relies upon independent estimates of AI's global freshwater footprint (Li et al., 2023). 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 context, they can also be related to Sustainable Development Goals (SDG) eight (decent work) and SDG six (clean water and sanitation) (United Nations, 2015).

Although Hao provides strong evidence to support her argument, the imperial framework has a conceptual limitation. The book is not expected to provide institutional solutions; rather, the limitation identified here arises from the nature of the 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) indicates that although such models identify patterns that exist in historical colonialism, they do not capture the differences in mechanisms. Colonial powers utilized force, laws, and sovereign rights to control peoples living under them. Applying this distinction to the companies discussed by Hao, we argue that they instead utilize contractual agreements, procurement practices, and jurisdictional arbitration to exert control over individuals who reside at a considerable distance from those making decisions. Any attempt at providing a remedy would need to occur through these mechanisms. Therefore, a model that focuses on the continuity of patterns versus difference in mechanisms clearly defines where damage is created much better than how to mitigate or repair that damage.

For those interested in organizations as opposed to AI specifically, this book has relevance. The book illustrates how an organization can experience “mission drift” due to financial dependency on capital; the limits of board oversight when directors have no influence on resources provided to an organization; and how the utilization of outsourcing can transfer both human and reputational risks to jurisdictions with less stringent labor protections. There are several practical implications for practitioners and action researchers. Practitioners must test governance structures using resources that they have actual control over rather than simply testing against nominal authority a governance structure may possess. To ensure accountability

for outsourced labor, 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. In addition, adopting responsible technologies requires understanding what dependencies an organization creates when one supplier controls a capability an organization cannot build itself. Additionally, there are the implications of institutional reflection, which is central to action research. Institutional reflection is the process of determining how an organization changes its mission because of the resources needed to achieve that mission. Institutional leaders could work through this as a cycle: review the AI procurement decisions already made, identify the dependencies and consequences those decisions created but did not anticipate, change practice where the review indicates it is needed, and then evaluate what those changes actually altered. Approached this way, the questions Hao raises about OpenAI become questions an institution can ask about itself.

These are specific concerns for universities currently sourcing, distributing, and normalizing generative AI systems. They therefore represent similar governance challenges identified by Hao. What languages a system can be used in shapes which students may use it to learn, and Hao states that only a relatively small number of the world's living languages are being adequately supported. Procurement decisions about where institutional data are processed, under what jurisdictions and on what contractual arrangements have significant implications for academic freedom. 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, where those with formal authority did not control the resources the organization depended upon.

The documentary accomplishment of this book as a whole transcends the controversy over its framing. Hao documents how an organization operating under a mission was affected by the requirements of pursuing that mission and who ultimately bore the costs of those effects. The imperial argument will likely continue to be debated; it should be. The institutional evidence upon which this argument is based is likely to 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>