

Review of Feeding the Machine: The Hidden Labour Powering AI

Joel Blechinger

Mount Royal University Library

ORCID: 0000-0002-4907-3601

ABSTRACT: Review of Muldoon, J., Graham, M., & Cant, C. (2024). *Feeding the Machine: The Hidden Human Labour Powering A.I.* Canongate Books.

Keywords: artificial intelligence, labour



This is an Open Access article distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Radical Librarianship, Vol. 10 (2024) pp.223-226. Published 8 Dec 2024.

Feeding the Machine: The Hidden Human Labour Powering A.I. by James Muldoon, Mark Graham, and Callum Cant is the product of over a decade of the authors' fieldwork research on the different labourers that comprise artificial intelligence's (AI) global workforce. As generative AI (GenAI) functionality gets rapidly integrated into legacy library systems and software—such as the Research Assistants currently being pushed out by Clarivate (in Primo, Web of Science, and ProQuest), not to mention many new GenAI-powered discovery tools attempting to disrupt the market like Scite, Elicit, and Consensus—it is imperative that library workers educate themselves about the real harms and risks associated with widespread AI adoption.

James Muldoon is an Associate Professor of Management at the University of Essex, Mark Graham is a Professor of Internet Geography at the Oxford Internet Institute, and Callum Cant is a Senior Lecturer in Management at Essex Business School. Graham is also Director of the Fairwork project, which uses fair work principles to score companies in the platform economy with the aim of improving working conditions. For Muldoon et al. (2024), AI can be conceived of as an “extraction machine”: “The extraction machine draws in critical inputs of capital, power, natural resources, human labour, data, and collective intelligence and transforms these into statistical predictions, which AI companies, in turn, transform into profits. To understand AI as a machine is to unmask its pretensions to objectivity and neutrality” (p. 8).

The book's prologue opens with a shocking story about Mercy, a Meta content moderator in Nairobi. Mercy is tasked with reviewing footage taken at a fatal car crash scene for the platform, only to realize with dawning horror that the crash's victim is her own grandfather. Deeply traumatized by the incident, she is given the next day off by her manager, but still encouraged to remain at work to make her daily ticket targets. Though the psychological toll of content moderation and data annotation labour has been revealed in other sources—such as the illuminating work of Roberts (2019) and Gray & Suri (2019), and, more recently, Perrigo (2023) and Dzieza's (2023) crucial reporting—the simultaneously banal yet cruel nature of Mercy's story grounds us in the material harm wrought by AI tools that often appear dematerialized and that can feel “like magic” to an end user (Herrman, 2023).

From this harrowing account of Mercy, Muldoon et al. (2024) move across vast geographic terrain in *Feeding the Machine*'s eight chapters. Chapter one focuses on Anita, a data annotator for a business process outsourcing company in Gulu, Uganda, who works long hours for meager pay reviewing data that is used to train artificial intelligence for an autonomous vehicle company. Chapter two shifts to Li, a machine learning engineer in Shoreditch, London, employed by a tech company developing its own large language model (LLM). Einar, the subject of chapter three, is a facility manager at a new data centre in Blönduós, a small town in northern Iceland. Chapter four then introduces us to Laura, a voice actor who loses her standard voice acting gigs due to having unknowingly signed away her own voice's rights to a big tech company in perpetuity. *Feeding the Machine*'s fifth chapter demonstrates the real-world effects of algorithmic management as it details the working conditions of Alex, an Amazon fulfillment centre employee in Coventry, England. Muldoon et al. (2024) then draw a stark contrast in the book's next chapter by moving to Tyler, a partner at a Silicon Valley venture capital firm in San Francisco. The book's penultimate chapter features several organizers in content moderation and annotation work, beginning with Paul, who agitates for change at the same company as Mercy in Nairobi. This seventh chapter also delves into the history of the African Content Moderators

Union, as told by three of its founders, Rose, Mohammed, and Brian. Finally, *Feeding the Machine*'s last chapter, "Rewiring the Machine," offers "five concrete steps to guide us in the right direction towards a world in which all AI workers are afforded dignity and respect" (p. 190): build worker power, hold companies accountable, government intervention, worker cooperatives, and dismantle the extraction machine.

Though seven of *Feeding the Machine*'s eight chapters initially focus on the aforementioned individuals, each chapter departs from the specifics of that person's working conditions to explore the broader issues and debates in AI that surround that job. For example, chapter three moves from Einar's work as an Icelandic data centre technician out to the political economy of hyperscale cloud computing providers and what the authors term Big AI's "infrastructural power": "a capacity to deploy large-scale computational resources for achieving strategic goals" (Muldoon et al., 2024, p. 79). Similarly, the chapter on Tyler, the Silicon Valley venture capitalist, includes a dense history of Californian venture capital that spans the genocide of the area's Indigenous population through to Richard Barbrook and Andy Cameron's "Californian Ideology" and the anti-democratic nature of what Muldoon et al. (2024) term the "founder's mentality." This is one of *Feeding the Machine*'s marked strengths: its authors have both the interdisciplinary knowledge and the political vocabulary to properly contextualize AI across diverse work. In doing so, Muldoon et al. (2024) avoid falling back on some of the tedious ideas that one sees uncritically repeated in mainstream library AI discourse—such as the Value-Neutrality Thesis, which asserts that "technology is morally and politically neutral, neither good nor bad" (Miller, 2021, p. 54). Where mainstream library discourse treats such ideas as though they are capital-K Knowledge, universally held by all library workers, Muldoon et al. demonstrate that they are, in fact, subjects of hotly contested debate in disciplines like science and technology studies.

Libraries and librarians receive no mention in *Feeding the Machine*, but this does not mean that the "AI extraction machine," as Muldoon et al. describe it, is not relevant to library work. In the September-October issue of the Canadian Association of University Teachers' *Bulletin*, Ribaric & Farnum (2024) make the compelling case for academic librarians to advocate for clear AI protections for their work in collective agreements. Drawing on the example of the 2023 Screen Actors Guild–American Federation of Television and Radio Artists (SAG–AFTRA) and Writers Guild of America (WGA) strike actions, Ribaric & Farnum (2024) remind librarians that we are also vulnerable to AI automation in a way strikingly similar to Laura, the voice actor. As Ribaric & Farnum (2024) ask, "[c]ould a chat-bot be devised that would replace the time-honoured librarian skill of the reference interview? We're guessing no, but that doesn't mean some service provider won't attempt it" (p. 27). They continue, "[w]hat recourse would there be if such a [chat-bot reference interview] service were shopped around by a vendor? We're going to guess not that much. New norms would be established around these tools to guide expectations of what they reliably could do, but that misses the bigger point, and the battle would already be lost" (p. 27). To my mind, this is an especially trenchant critique lodged by Ribaric & Farnum (2024), because it highlights how academic librarians' unwavering belief in the service ethic—in our ability to teach about a phenomenon, to make a LibGuide about a phenomenon, etc.—combined with our inability to self-conceive as workers in solidarity with other workers—particularly at institutions where librarians have faculty status and, therefore, effectively function as a labour aristocracy (Popowich, 2020)—could obscure the actual threats posed by AI tools. In other

words, if we focus *solely on the outputs* of AI tools—their accuracy, reliability, and teachability—we risk overlooking the more material effects that such tools could have on our labour conditions: acceleration and degradation of work, deskilling, outright role elimination, etc.

Zvyagintseva and I (2023) similarly warned academic librarians about becoming too invested in the then nascent discursive project of “AI literacy”: “[l]ibrary thought leaders will ask us to puzzle out the minute differences between digital literacy, media literacy, information literacy, and AI literacy. Due to the urgency of this situation, however, we caution librarians from becoming too discursively focused on this ideological proposition to the neglect of investing in labour organizing” (p. 19). Reading Muldoon et al.’s (2024) *Feeding the Machine* gives librarians a strong conceptual foundation to use to critically examine AI and to identify where our work is being increasingly drawn into AI as an extraction machine in our own localized contexts. That the book does not feature librarians should not be seen as a barrier to entry for library audiences, rather it can be productively reframed as a challenge to those same audiences to think beyond workplace siloization and discrete role definition to recognize commonalities with other workers in often remote parts of the world whose lives are and will be materially altered by AI as extraction machine, just as our labour as librarians will be.

References

- Dzieza, J. (2023, June 20). *AI is a lot of work*. The Verge.
<https://www.theverge.com/features/23764584/ai-artificial-intelligence-data-notation-labor-scale-surge-remotasks-openai-chatbots>
- Gray, M. L., & Suri, S. (2019). *Ghost work: How to stop Silicon Valley from building a new global underclass*. Harper Business.
- Herrman, J. (2023, January 18). *The AI magic show*. Intelligencer.
<https://nymag.com/intelligencer/2023/01/why-artificial-intelligence-often-feels-like-magic.html>
- Miller, B. (2021). Is technology value-neutral? *Science, Technology, & Human Values*, 46(1), 53–80. <https://doi.org/10.1177/0162243919900965>
- Muldoon, J., Graham, M., & Cant, C. (2024). *Feeding the machine: The hidden human labour powering A.I.* Canongate Books.
- Perrigo, B. (2023, January 18). OpenAI used Kenyan workers on less than \$2 per hour to make ChatGPT less toxic. *Time*. <https://time.com/6247678/openai-chatgpt-kenya-workers/>
- Popowich, S. (2020, February 29). *Precarity and tenure*. Red Librarian Blog.
<https://redlibrarian.github.io/article/2020/02/29/precarity-and-tenure.html>
- Ribaric, T., & Farnum, C. (2024, September-October). AI protections for librarian work: put it in the collective agreement? *CAUT Bulletin*, 71(3), 27–28.
- Journal of Radical Librarianship*, Vol. 10 (2024) pp.223-226

Roberts, S. T. (2019). *Behind the screen: Content moderation in the shadows of social media*. Yale University Press.

Zvyagintseva, L., & Blechinger, J. (2023). “The perils of library instruction”: Contextualizing EdTech's encroachment on teaching through labour theory of value in academic libraries. *Canadian Journal of Academic Librarianship*, 9, 1–25. <https://doi.org/10.33137/cjal-rcbu.v9.40958>