

Do We Need Computing Power Ethics? Participation Conditions in the AI Era

Yunqi Jiang

School of Philosophy, Renmin University of China, Beijing, China
vincijiang@ruc.edu.cn

Abstract

Computing power distribution is often framed in terms of markets, environmental costs, security, or geopolitics, treating computing power as a commodity or strategic resource. These approaches overlook why its distribution matters even before direct harm occurs. I argue for a distinct computing power ethics. Through public-health modelling and marginalised-language cases, I expose the limits of the Market-Permission Principle and propose a condition-oriented account centred on meaningful participation in AI-mediated social life.

Computing Power as an Ethical Object

Computing power comprises the chips, GPU clusters, data centres, cloud services, and related resources used to train, run, and deploy AI models. It determines who can train, reproduce, audit, deploy, govern, and contest AI systems (?). The philosophical question is whether computing power should become a direct object of ethical inquiry.

The existing frameworks address some important issues. Business ethics examines pricing, platform lock-in, and corporate responsibility of computing power (?). Environmental ethics concerns the energy, water, land, minerals, and infrastructure on which computing power depends (?). Technology-risk governance treats computing power as a measurable lever for identifying and restricting dangerous capabilities (?). Global justice examines the concentration of advanced chips and cloud infrastructure, unequal technological capacity, and the division of computing power between firms, elite universities, and less-resourced institutions (??). These approaches are indispensable, but they still typically translate computing power into another object, such as a commodity, environmental burden, security lever, industrial asset, or geopolitical resource.

The resulting gap in fact exists ahead of the appearance of any discriminatory output, privacy violation, or security incident. Mainstream AI ethics often evaluates consequences after an AI system has emerged (??). It overlooks the prior structure that determines who is able to produce, verify, deploy, or challenge such systems. Computing power

ethics addresses this structure, by asking how the availability and allocation of computing power enable or obstruct effective participation in AI-mediated practices. When computing power becomes a practical prerequisite for participation in a significant practice, factors such as ownership, commercial return or national competition can no more by themselves justify its distribution. Computing power then acquires normative significance as a condition of availability.

The Market-Permission Principle with Fault

The allocation currently reflects a commonly held intuition: if computing resources are obtained through lawful purchase, lease, or investment, and their use neither infringes rights nor causes identifiable direct harm, allocation may normally be left to market choice. I call this the Market-Permission Principle. Its weakness is mistakenly treating formal permission to enter the market as equivalent to having the practical capacity to participate.

Let us imagine a public-university team that cannot afford sustained computing power for a local disease-prediction model and commercial actors purchase similar resources for more profitable applications. No transaction is unlawful, and commercial use is not inherently immoral. Repeated exclusion yet may crowd out public knowledge production and decrease local public-health capacity. The team remains formally free to buy computing power while lacking the conditions needed to enter AI practices. Market permission does not register this loss because the logic of public value and that of purchasing power diverge.

A parallel problem arises when a marginalised-language community seeks models for education, healthcare, law, or cultural archives. Those languages require data collection, cleaning, annotation, training, and evaluation, yet their smaller markets make sustained computing power difficult to obtain. General-purpose models may provide nominal coverage without reliability or cultural appropriateness (?). The community is not legally prohibited from developing its language, but it cannot effectively represent itself in AI practices. Also low willingness to pay does not embody low cultural value or weak eligibility to participate.

These cases distinguish *access to market* from *access to participation*. Formal permission is the absence of prohibition. Effective participation requires practical, stable, afford-

able, and sustainable conditions for entering, maintaining, and shaping a practice. When computing power mediates public health, research, education, governance, or cultural expression, its allocation simultaneously affects the distribution of social opportunities. Lawful transactions remain normatively important, but legality and the absence of direct harm may fail to recognize the allocations that systematically undermine public welfare or cultural representation.

Objections and Scope Conditions

One objection holds that competition law, environmental ethics, and global justice, especially infrastructure ethics, already cover the issue. Those approaches can regulate ecological burdens, monopolies, and transnational inequality. The remaining question is who can participate in AI systems themselves. Computing power both supports downstream activity and shapes systems that reorganise education, research, public administration, language visibility, and institutional decision-making. Control over computing power affects cognitive and political participation, not only access to an ordinary material input.

A second objection is that the argument over-moralises resource competition. Funding, expertise, laboratories, land, and conference access also affect participation. The reply is to distinguish an advantage condition from a threshold condition. An advantage makes participation easier or more competitive. A threshold determines whether an actor can enter or sustain the practice at all. The distribution of computing power warrants special urgency when referring to the latter function.

More specifically, three initiative conditions should be met. First, computing power must be a genuine threshold for a significant AI practice rather than a general convenience. Second, the practice must have public significance, such as scientific research, public governance, cultural preservation. Third, allocation must systematically affect actors' effective position rather than cause an isolated competitive loss. If these conditions are absent, computing power may remain an ordinary commercial resource. Computing power ethics neither requires justice review for every GPU lease nor treats every unmet request as an injustice.

Conclusion

Computing power ethics is a condition-oriented field concerned with whether actors can meaningfully participate in an AI-mediated social world. Its distinct contribution is to evaluate allocation before particular model harms arise and to reveal exclusions that lawful exchange alone cannot identify. When computing power becomes a threshold for socially significant practices, its availability can no longer be treated as merely a technical or commercial question. It must enter ethical evaluation as a condition of meaningful participation.

Acknowledgments

I thank Long Li, Zhekai Zhang and Shengwei Yin for their careful reading and constructive feedback.

References

- Ahmed, N.; and Wahed, M. 2020. The De-democratization of AI: Deep Learning and the Compute Divide in Artificial Intelligence Research. arXiv preprint. arXiv:2010.15581.
- Bang, Y.; Cahyawijaya, S.; Lee, N.; et al. 2023. A Multi-task, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning, Hallucination, and Interactivity. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, 675–718. Association for Computational Linguistics. doi.org/10.18653/v1/2023.ijcnlp-main.45.
- Besiroglu, T.; Bergerson, S. A.; Michael, A.; Heim, L.; Luo, X.; and Thompson, N. 2024. The Compute Divide in Machine Learning: A Threat to Academic Contribution and Scrutiny? arXiv preprint. arXiv:2401.02452.
- Donaldson, T.; and Walsh, J. P. 2015. Toward a Theory of Business. *Research in Organizational Behavior*, 35: 181–207. doi.org/10.1016/j.riob.2015.10.002.
- Floridi, L.; Cows, J.; Beltrametti, M.; et al. 2018. AI4People: An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations. *Minds and Machines*, 28(4): 689–707.
- Mittelstadt, B. 2019. Principles Alone Cannot Guarantee Ethical AI. *Nature Machine Intelligence*, 1: 501–507. doi.org/10.1038/s42256-019-0114-4.
- O'Neill, J.; Holland, A.; and Light, A. 2008. *Environmental Values*. London: Routledge. doi.org/10.4324/9780203495452.
- Sastry, G.; Heim, L.; Belfield, H.; et al. 2024. Computing Power and the Governance of Artificial Intelligence. arXiv preprint. doi.org/10.48550/arXiv.2402.08797, arXiv:2402.08797.
- Vipra, J.; Rodrigues, C.; Bioni, B.; et al. 2024. Governing Computational Infrastructure for Strong and Just AI Economies. T20 Brasil and IT for Change.