Restraint by Design: Rethinking Artificial Intelligence and Environmental Governance in an Age of Ecological Crisis

Dr. Harkirat Singh Bal[1*], Mr. Aditya Shekhar[1], Mr. Amit Kumar[1], Mr. Amritanshu Shekhar[1], Mr. Amit Kumar[1], Dr. Anand Singh Prakash[1], Ms. Elisha Lakra[1]

[1*]ICFAI University, Jharkhand, India

[1]Assistant Professor (All Authors) | Email: singhkirat487@gmail.com | Email: aadityasfj@gmail.com | Email: amitsingh49@gmail.com | Email: amritanshu211@gmail.com | Email: kumaramit990@gmail.com | Email: anandair03@gmail.com | Email: elishalakra0@gmail.com

***Corresponding Author:** Dr. Harkirat Singh Bal, Assistant Professor, ICFAI University, Jharkhand, India | Email: singhkirat487@gmail.com

Abstract

Artificial intelligence is increasingly promoted as indispensable to addressing the climate and biodiversity crises, from optimizing energy grids to modelling deforestation. Yet the same computational infrastructure that promises environmental salvation generates substantial ecological costs of its own: escalating electricity demand, freshwater consumption for data-centre cooling, and resource-intensive hardware supply chains. This paper argues that the dominant "AI-for-climate" paradigm is constrained by an unexamined techno-solutionist logic that treats environmental problems as optimization problems awaiting sufficiently powerful models, while remaining largely silent on AI's own material footprint and its tendency toward centralized, extractive control over natural systems. Engaging recent scholarship on AI's environmental costs alongside classical Daoist ecological philosophy — particularly its account of non-coercive action and spontaneous natural patterning the paper develops an original interpretive framework of "restrained AI governance" that reorients environmental AI deployment away from maximal control toward proportionate, reversible, and ecologically situated intervention. The paper concludes with concrete governance implications for regulators, technologists, and environmental policymakers.

Keywords: *Artificial Intelligence; Environmental Governance; Daoist Philosophy; Non-Coercive Action; Techno-solutionism; Sustainable Technology*

1. Introduction

Two narratives about artificial intelligence and the environment circulate with equal confidence today, and they are rarely allowed to occupy the same page. The first is celebratory: machine-learning systems now forecast extreme weather with unprecedented skill, detect methane leaks from satellite imagery, optimize renewable-energy grids in real time, and track deforestation and illegal fishing at planetary scale.[1][2] International bodies have embraced this narrative enthusiastically — the World Economic Forum frames AI as a "planetary" instrument for sustainability, and the United Nations Environment Programme has incorporated algorithmic foresight tools into its own reporting infrastructure.[3][4] The second narrative is cautionary: the training and operation of large-scale AI systems consume extraordinary quantities of electricity and fresh water, depend on supply chains built on the extraction of lithium, cobalt, and rare-earth minerals, and generate a rapidly growing volume of electronic waste.[5] These two narratives are not, in fact, in tension so much as they are incomplete in

---

[1]David Rolnick et al., "Tackling Climate Change with Machine Learning," ACM Computing Surveys 55, no. 2 (2022): 1-96.

[2]Remi Lam et al., "Learning Skilful Medium-Range Global Weather Forecasting," Science 382, no. 6677 (2023): 1416-1421 (describing DeepMind's GraphCast model).

[3]World Economic Forum, AI for the Planet: Harnessing Artificial Intelligence for Environmental Sustainability (Geneva: WEF Insight Report, 2024), 6-14.

[4]United Nations Environment Programme, Navigating New Horizons: A Global Foresight Report on Planetary Health and Human Wellbeing (Nairobi: UNEP, 2024), 44-49.

[5]Kate Crawford, Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence (New Haven: Yale University Press, 2021), 23-51.

the same way: both treat the environment as a problem domain to be acted upon by AI, rather than as a set of relationships within which AI itself must be situated and governed.

This paper argues that the prevailing discourse on AI and the environment in policy, industry, and much of the academic literature is organized around an inherited assumption that more data, more computation, and more centralized control necessarily yield better environmental outcomes. This assumption, which this paper terms *computational maximalism*, is rarely examined as a philosophical premise rather than a technical necessity. The paper's central and original contribution is to show that classical Daoist ecological thought, not as a mystical alternative to be romanticized, but as a rigorous philosophy of situated, non-coercive action supplies exactly the conceptual resources needed to interrogate and correct this premise. Drawing selectively on Daoist ideas of non-coercive, minimally forceful action, the paper develops a three-principle framework proportionality, reversibility, and situatedness for what it calls *restrained AI governance*: an approach to deploying AI in environmental contexts that treats minimal, well-calibrated intervention, rather than maximal control, as the marker of good governance. The paper's focus remains squarely on the contemporary environmental and technological problem; Daoist philosophy functions throughout as an interpretive and normative lens rather than as the object of inquiry in its own right.

The argument proceeds in five further parts. Section 2 surveys the genuine and growing contribution of AI to environmental monitoring and mitigation. Section 3 critically reviews recent scholarship documenting AI's own environmental costs. Section 4 situates this tension within the broader critique of techno-solutionism in environmental and technology studies. Section 5 introduces Daoist ecological philosophy as a corrective lens, and Section 6 develops the paper's original framework for restrained AI governance, with concrete institutional and regulatory implications.

2. The Contemporary Environmental Crisis and the Turn to AI

The Intergovernmental Panel on Climate Change's Sixth Assessment Synthesis Report leaves little room for complacency: global surface temperature has already risen approximately 1.1°C above pre-industrial levels, current policies place the world on a trajectory well beyond the 1.5°C threshold, and the window for limiting warming to manageable levels is closing rapidly.[6] Biodiversity loss, freshwater stress, and the accumulation of plastic and chemical pollutants compound this picture, producing what the UNEP now describes as a "triple planetary crisis" of climate change, biodiversity loss, and pollution.

Artificial intelligence has been enlisted as a central instrument in the global response to this crisis, and the underlying reasons are not difficult to see. Machine-learning models excel at precisely the kind of high-dimensional pattern recognition that environmental monitoring requires. DeepMind's GraphCast model, for instance, produces ten-day global weather forecasts in under a minute at a fraction of the computational cost of traditional numerical weather prediction, while matching or exceeding its accuracy on the majority of tracked variables.[7] A comprehensive survey by Rolnick and colleagues catalogues dozens of further applications: satellite-based deforestation and illegal-fishing detection, machine-learning-optimized battery chemistries for grid storage, computer-vision systems for wildlife population monitoring, and reinforcement-learning approaches to real-time electricity-demand balancing that reduce reliance on fossil-fuel "peaker" plants.[8] The promise of these applications is genuine, and this paper does not dispute it. What it disputes is the framing within which that promise is typically presented: as an unqualified good, whose costs are either absent or negligible relative to the environmental benefit AI is presumed to deliver.

This framing has become institutionalized. The World Economic Forum's 2024 report on AI and sustainability catalogues use cases across agriculture, energy, and conservation while devoting only a brief, largely aspirational section to AI's own footprint, framed as a problem that further technical innovation more efficient chips, better cooling will resolve.[9] UNEP's foresight reporting adopts a similar structure. The consistent pattern across this literature is an asymmetry: AI's contribution to solving environmental problems is treated as a matter of ongoing empirical demonstration, while AI's contribution to creating environmental problems is treated as a matter of future technical correction. Section 3 addresses why this asymmetry does not withstand scrutiny.

3. The Hidden Ecological Cost of AI: A Critical Review of Recent Scholarship

A rapidly maturing body of scholarship has begun to document what Kate Crawford calls the "planetary costs" of artificial intelligence costs that are distributed across the entire life cycle of an AI system, from the mining of the minerals used in its hardware to the electricity used in its training and the water used to cool the data centres

[6] Intergovernmental Panel on Climate Change, Climate Change 2023: Synthesis Report (Geneva: IPCC, 2023), 4-8.
[7] Lam et al., "Learning Skilful Medium-Range Global Weather Forecasting," 1418-1420.
[8] Rolnick et al., "Tackling Climate Change with Machine Learning," 12-19.
[9] World Economic Forum, AI for the Planet, 41-44.

in which it runs.[10] Crawford traces the material substrate of AI back to lithium extraction in the Nevada desert and Bolivian salt flats, and to the rare-earth mining operations that supply the magnets and semiconductors in every data-centre server, arguing that AI's apparent immateriality its existence as software, as “cloud” obscures an extraction economy as old as industrial mining itself.[11] The International Energy Agency's 2023 review of critical-minerals markets corroborates this account from a policy-analytic angle, projecting that demand for the minerals used in digital and clean-energy infrastructure, including AI hardware, will place unprecedented strain on already fragile extraction and refining supply chains.[12]

The energy question has become still more pressing. Luccioni, Jernite, and Strubell's 2024 study of AI deployment costs finds that inference the ongoing running of trained models to answer everyday queries now constitutes the majority of many AI systems' lifetime energy consumption, not the widely publicized training phase, meaning that the environmental cost of AI scales directly with adoption rather than being a one-time expenditure.[13] This finding is significant precisely because it undercuts a common defence of AI's energy use: that training costs are amortized across a model's useful life and therefore diminish in relative significance over time. If inference dominates lifetime cost, then every new deployment of an AI-based environmental tool a new agricultural monitoring app, a new grid-optimization dashboard adds a durable, ongoing energy burden rather than a one-off cost.

The water dimension of this footprint has received comparatively less attention but is no less significant. Li, Yang, Islam, and Ren's widely cited study estimates that training a single large language model of the scale common by 2023 could consume several hundred thousand litres of fresh water for cooling, with additional, harder-to-quantify “indirect” water costs embedded in the electricity generation that powers data centres.[14] The International Energy Agency's most recent electricity market analysis projects that global data-centre electricity demand will more than double by 2027, with AI-optimized facilities identified as the single fastest-growing category of demand a trajectory that, absent intervention, runs directly counter to the decarbonization pathways that AI is simultaneously being deployed to help achieve.[15] Strubell, Ganesh, and McCallum's earlier but still-influential analysis of natural-language-processing models found that training a single large transformer model could produce carbon emissions comparable to the lifetime emissions of several automobiles an early warning that the field's subsequent trajectory toward ever-larger models has, if anything, intensified rather than resolved.[16]

None of this scholarship argues that AI should not be used for environmental purposes. Its cumulative import is narrower and more precise: that the environmental benefit of any given AI application cannot be assumed but must be assessed net of its own material and energetic cost, and that the current governance discourse exemplified by the WEF and UNEP reports discussed in Section 2 systematically fails to conduct or demand this accounting. This is not primarily a technical failure. It is, as Section 4 argues, a symptom of a deeper and more general pattern in how environmental problems have come to be framed within the AI era.

4. Techno-Solutionism and Its Critics: Situating the Debate

Evgeny Morozov's critique of what he terms “technological solutionism” the tendency to recast complex social and ecological problems as tractable engineering problems with elegant technical fixes — provides a useful vocabulary for the pattern identified in Section 3.[17] Morozov's central worry is not that technology cannot help; it is that solutionist framing forecloses inquiry into whether a problem has been correctly diagnosed in the first place, and tends to favour interventions that are legible to engineers and investors over interventions that might be more effective but less computationally interesting.[18] Applied to the environmental domain, this critique suggests that the enthusiasm for AI-driven monitoring, optimization, and forecasting may reflect not a considered judgment about the most effective response to ecological crisis, but the availability of a technology whose proponents have a strong institutional and commercial interest in its adoption.

[10] Crawford, Atlas of AI, 23-30.
[11] Crawford, Atlas of AI, 92-95.
[12] International Energy Agency, Critical Minerals Market Review 2023 (Paris: IEA, 2023), 18-27.
[13] Alexandra Sasha Luccioni, Yacine Jernite, and Emma Strubell, “Power Hungry Processing: Watts Driving the Cost of AI Deployment?” in Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (New York: ACM, 2024), 85-93.
[14] Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren, “Making AI Less ‘Thirsty’: Uncovering and Addressing the Secret Water Footprint of AI Models,” arXiv preprint arXiv:2304.03271 (2023), 1-15.
[15] International Energy Agency, Electricity 2025: Analysis and Forecast to 2027 (Paris: IEA, 2025), 32-40.
[16] Emma Strubell, Ananya Ganesh, and Andrew McCallum, “Energy and Policy Considerations for Deep Learning in NLP,” in Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (Stroudsburg, PA: Association for Computational Linguistics, 2019), 3645-3650.
[17] Evgeny Morozov, To Save Everything, Click Here: The Folly of Technological Solutionism (New York: PublicAffairs, 2013), 5-18.
[18] Morozov, To Save Everything, Click Here, 291-304.

Jathan Sadowski extends this critique with the concept of digital "smartness" as a mode of governance that treats data extraction and algorithmic control as inherently superior to other forms of coordination a logic he traces through smart cities, precision agriculture, and algorithmic resource management, and which he argues tends toward centralization of control in the hands of the firms and states that own the underlying computational infrastructure.[19] Sadowski's analysis is particularly apt for environmental AI: a machine-learning system for optimizing irrigation or grid dispatch does not merely process information about a natural system; it constitutes a new locus of control over that system, one whose environmental logic is frequently subordinated to the commercial logic of the firm operating it.[20] Crawford's account of AI's material infrastructure makes a structurally similar point from a different angle: the same computational apparatus celebrated for its capacity to model planetary systems is embedded in extractive relationships with the very ecosystems and communities it purports to protect.[21]

Environmental philosophy has long anticipated this critique from a different disciplinary direction. Val Plumwood's analysis of what she calls the logic of "mastery" identifies a recurring pattern in Western environmental thought and practice, in which nature is construed as an object to be known, predicted, and controlled by a rational, disembedded subject a pattern she traces from Baconian science through to contemporary resource-management regimes.[22] The AI-for-climate paradigm, on this reading, is not a departure from this mastery logic but its most technologically sophisticated expression to date: environmental problems are reframed as prediction and optimization problems, soluble in principle by a sufficiently comprehensive model, with the model's own material embeddedness in the world it is modelling treated as incidental rather than constitutive.

What the literature reviewed in this section lacks, however, is a well-developed alternative normative vocabulary a positive account of what a non-mastery-oriented relationship between computational systems and natural systems would actually look like in practice, rather than simply a critique of the mastery orientation. It is at this point that Daoist ecological philosophy becomes analytically useful, not as an exotic supplement to Western critique but as an independently developed tradition of thought that has, for over two millennia, theorized non-coercive modes of engaging complex, self-organizing systems.

5. Daoist Ecological Thought as a Corrective Lens

Two ideas from the classical Daoist corpus are of particular relevance here. The first, commonly rendered "naturalness" or "self-so-ness," designates the capacity of things to unfold according to their own internal patterns of coherence rather than according to an externally imposed design.[23] The Daodejing's twenty-fifth chapter places this idea of naturalness at the summit of a cosmological hierarchy in which "humanity models itself on earth, earth models itself on heaven, heaven models itself on dao, and dao models itself on that which is so of itself" a formulation that resists any reading of this naturalness as mere passivity, and instead positions it as the deepest ordering principle to which even the most fundamental cosmic processes are answerable.[24]

The second, related idea is frequently mistranslated as simple inaction. Contemporary scholarship, following Ames and Hall, instead reads it as *non-coercive* action: action that achieves its ends by working with the grain of a system's own dynamics rather than against them, minimizing forceful, unilateral intervention in favour of subtle, well-timed, and proportionate engagement.[25] The Daodejing's sixty-fourth chapter captures this orientation in its counsel to "attend to things in their nascent state" and to "act without acting on" that is, to intervene early, lightly, and reversibly, rather than to wait for a problem to compound before applying maximal force to correct it.[26]

The Zhuangzi supplies the tradition's most vivid illustration of this orientation in the parable of Cook Ding, whose knife remains perpetually sharp not because he cuts through the ox's structure by force but because he has learned to follow the natural openings and joints of the animal's body, finding "where there is space" rather than imposing an external logic of division.[27] The scholarly consensus represented in the landmark Harvard volume Daoism and Ecology reads this parable, and the broader pairing of naturalness and non-coercive action, as

---

[19] Jathan Sadowski, Too Smart: How Digital Capitalism Is Extracting Data, Controlling Our Lives, and Taking Over the World (Cambridge, MA: MIT Press, 2020), 61-80.
[20] Sadowski, Too Smart, 145-162.
[21] Crawford, Atlas of AI, 211-224.
[22] Val Plumwood, Feminism and the Mastery of Nature (London: Routledge, 1993), 41-68.
[23] Roger T. Ames, "Taoism and the Nature of Nature," Environmental Ethics 8, no. 4 (1986): 317-322.
[24] Laozi, Daodejing, ch. 25, in Roger T. Ames and David L. Hall, trans., Dao De Jing: Making This Life Significant — A Philosophical Translation (New York: Ballantine Books, 2003), 116-118.
[25] Ames and Hall, Dao De Jing, 27-35.
[26] Laozi, Daodejing, ch. 64, in Ames and Hall, Dao De Jing, 175-177.
[27] Zhuangzi, "The Secret of Caring for Life," in Burton Watson, trans., The Complete Works of Zhuangzi (New York: Columbia University Press, 2013), 20-22.

articulating a coherent ecological ethic: one grounded not in the renunciation of human agency but in its recalibration, from a mode of domination to a mode of skilful, minimally disruptive participation in larger, self-organizing systems.[28] James Miller's more recent work extends this reading explicitly into the register of contemporary sustainability, arguing that Daoist cosmology offers a template for “green” institutional design premised on cultivating the resilience of existing ecological and social patterns rather than replacing them with engineered alternatives.[29]

It is important to be precise about the argumentative role this philosophy plays in the present paper. The claim is not that classical Daoist thinkers anticipated or could speak directly to twenty-first-century computational governance such a claim would be historically anachronistic and analytically unhelpful. The claim is narrower: that this framework articulates, with unusual precision, a set of design values proportionate rather than maximal intervention, attentiveness to a system's own emergent patterns, and reversibility of action that are strikingly well suited to correcting the computational-maximalist premise identified in Sections 2 through 4.[30] The Daodejing's second chapter offers a further, structurally important resource here: its account of the mutual arising of opposites (being and non-being, difficulty and ease) as a caution against the belief that any single agent, human or algorithmic, can achieve unilateral mastery over a complex system without generating unanticipated countervailing effects.[31] This is, in substance, a philosophical anticipation of what systems theorists now call the problem of second-order effects in complex adaptive systems a problem to which large-scale AI environmental interventions, precisely because of their scale and speed, are especially exposed.

6. Toward a Framework for Restrained AI Governance

Building on the preceding analysis, this paper proposes a three-principle framework proportionality, reversibility, and situatedness for governing the deployment of AI in environmental contexts. The framework is offered as an original contribution to the emerging literature on AI environmental governance, developed specifically through, though not confined to, the interpretive resources of Daoist ecological thought.

The first principle, proportionality, holds that the scale of an AI system deployed for an environmental purpose should be calibrated to the marginal environmental benefit it delivers, rather than to the maximum computational capacity available. The Daodejing's eightieth chapter, in its counsel toward small, self-sufficient communities, is not read here as a literal political prescription but as a heuristic: bigger is not, by default, better.[32] Empirically, this principle has direct support: Rolnick and colleagues note that many high-value climate applications methane-leak detection, crop-disease identification, grid-frequency regulation can be achieved with comparatively small, efficient models, and that the pursuit of ever-larger general-purpose models for tasks that do not require that scale represents a significant and often unexamined source of unnecessary environmental cost.[33] A proportionality principle would require environmental-AI procurement, whether by governments, NGOs, or firms, to justify model scale against demonstrated marginal benefit, rather than defaulting to the largest available system.

The second principle, reversibility, holds that AI-based environmental interventions should be designed, where possible, to permit correction or withdrawal without irreversible lock-in an institutional analogue to the Daoist preference for early, light intervention over late, forceful correction.[34] This has concrete implications for infrastructure design: an AI-optimized irrigation system, for instance, should retain manual override and revert-to-baseline capability, and its optimization targets should be periodically revisited rather than locked in at deployment. Miller's account of the cyclical, non-linear temporal logic embedded in Daoist ritual practice offers a useful institutional analogy here: governance structures modelled on cyclical review and renewal, rather than one-time deployment followed by passive monitoring, better reflect the underlying emphasis on the ongoing, self-renewing character of natural systems.[35]

The third principle, situatedness, holds that AI environmental tools should be calibrated to local ecological and social conditions rather than optimized against universal, globally standardized targets a direct application of the Daoist emphasis on the internal patterns proper to a particular system, as opposed to patterns imposed uniformly from outside.[36] This principle cuts against a strong tendency in current AI development toward

---

[28] N.J. Girardot, James Miller, and Liu Xiaogan, eds., Daoism and Ecology: Ways within a Cosmic Landscape (Cambridge, MA: Harvard University Center for the Study of World Religions, 2001), xxxv-liii.

[29] James Miller, China's Green Religion: Daoism and the Quest for a Sustainable Future (New York: Columbia University Press, 2017), 3-22, 88-104.

[30] Ames, “Taoism and the Nature of Nature,” 331-336.

[31] Laozi, Daodejing, ch. 2, in Ames and Hall, Dao De Jing, 77-79.

[32] Laozi, Daodejing, ch. 80, in Ames and Hall, Dao De Jing, 192-194.

[33] Rolnick et al., “Tackling Climate Change with Machine Learning,” 44-47, 78-82.

[34] Laozi, Daodejing, ch. 64, in Ames and Hall, Dao De Jing, 175-177; Ames and Hall, Dao De Jing, 27-35.

[35] Miller, China's Green Religion, 141-158, 197-201.

[36] Ames, “Taoism and the Nature of Nature,” 317-322; Girardot, Miller, and Liu, Daoism and Ecology, xlv-l.

universal, transferable models trained on aggregated global data and deployed with minimal local recalibration a tendency Sadowski links to the broader extractive logic of “smart” governance.[37]

These three principles are not merely aspirational; they map onto concrete, currently available regulatory instruments. The European Union's Artificial Intelligence Act already requires energy and resource-use reporting for high-impact general-purpose AI models, a mechanism that could straightforwardly be extended into a proportionality-style requirement specifically for environmental-AI procurement, obliging public bodies to document the marginal environmental benefit of a proposed AI system against its documented resource cost before deployment.[38] India's recently issued AI governance guidelines likewise contemplate sector-specific risk and impact assessments, a framework capable of accommodating situatedness-style requirements for local ecological calibration of AI environmental tools deployed in agriculture and water management.[39] A reversibility requirement, meanwhile, could be implemented through relatively modest procurement conditions: mandatory manual-override capability, periodic mandatory re-justification of deployed systems, and default sunset clauses requiring active renewal rather than passive continuation of AI-based environmental interventions.

None of these instruments requires invoking Daoist philosophy explicitly in legislative text; the framework's purpose is interpretive and heuristic, not doctrinal. Its contribution is to supply regulators, technologists, and environmental policymakers with a coherent underlying rationale restraint as a positive governance value in its own right, rather than merely a residual constraint on an otherwise unqualified good for measures that might otherwise appear as an arbitrary patchwork of energy-reporting rules, override requirements, and localization mandates.

7. Conclusion

This paper has argued that the dominant discourse on artificial intelligence and environmental crisis rests on an unexamined premise of computational maximalism: the assumption that more data, more computation, and more centralized algorithmic control necessarily produce better environmental outcomes. Recent scholarship on AI's own material and energetic footprint its electricity demand, its water consumption, its dependence on extractive mineral supply chains demonstrates that this premise cannot be taken for granted, and that the celebratory framing typical of institutional reports on “AI for climate” systematically understates the environmental cost side of the ledger.

Building on the broader critique of techno-solutionism in environmental and technology studies, this paper has proposed that classical Daoist ecological philosophy specifically its account of non-coercive action and natural patterning offers precise and underused conceptual resources for developing an alternative governance orientation. The resulting framework of restrained AI governance, organized around the principles of proportionality, reversibility, and situatedness, reframes environmental AI deployment not as a race toward ever-greater computational control over natural systems, but as a disciplined practice of minimally sufficient, correctable, and locally attuned intervention.

This reframing carries a broader implication for the study of Daoist thought itself. Rather than treating Daoist ecological philosophy as a historical curiosity to be catalogued alongside other pre-modern environmental sensibilities, this paper suggests that its core conceptual apparatus remains directly generative for one of the most consequential governance problems of the present moment. Future research might extend this framework empirically, testing whether AI systems designed around proportionality, reversibility, and situatedness in fact deliver superior long-run environmental outcomes relative to maximal-scale alternatives, and might further develop the comparative dimension of this inquiry by placing Daoist ecological philosophy in sustained dialogue with other non-Western traditions of situated ecological knowledge that face structurally similar pressures from the expansion of computational governance.[40] What the tradition offers, at minimum, is a reminder that the most sophisticated response to ecological crisis may not be the most powerful algorithm, but the most disciplined restraint in how, when, and at what scale that algorithm is permitted to act.

Bibliography

- Ames, Roger T. “Taoism and the Nature of Nature.” Environmental Ethics 8, no. 4 (1986): 317-350.
- Ames, Roger T., and David L. Hall, trans. Dao De Jing: Making This Life Significant — A Philosophical Translation. New York: Ballantine Books, 2003.

[37]Sadowski, Too Smart, 61-80.

[38]European Union, Regulation (EU) 2024/1689 of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act), art. 40 and Annex XI (2024).

[39]Ministry of Electronics and Information Technology, Government of India, India AI Governance Guidelines (New Delhi, 2025), pt. III, paras 3.4-3.9.

[40]Miller, China's Green Religion, 219-225.

- Crawford, Kate. Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence. New Haven: Yale University Press, 2021.
- European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act). 2024.
- Girardot, N.J., James Miller, and Liu Xiaogan, eds. Daoism and Ecology: Ways within a Cosmic Landscape. Cambridge, MA: Harvard University Center for the Study of World Religions, 2001.
- Intergovernmental Panel on Climate Change. Climate Change 2023: Synthesis Report. Geneva: IPCC, 2023.
- International Energy Agency. Critical Minerals Market Review 2023. Paris: IEA, 2023.
- International Energy Agency. Electricity 2025: Analysis and Forecast to 2027. Paris: IEA, 2025.
- Lam, Remi, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Ferran Alet, Suman Ravuri, et al. “Learning Skilful Medium-Range Global Weather Forecasting.” Science 382, no. 6677 (2023): 1416-1421.
- Li, Pengfei, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren. “Making AI Less ‘Thirsty’: Uncovering and Addressing the Secret Water Footprint of AI Models.” arXiv preprint arXiv:2304.03271 (2023).
- Luccioni, Alexandra Sasha, Yacine Jernite, and Emma Strubell. “Power Hungry Processing: Watts Driving the Cost of AI Deployment?” In Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency, 85-99. New York: ACM, 2024.
- Ministry of Electronics and Information Technology, Government of India. India AI Governance Guidelines. New Delhi, 2025.
- Miller, James. China's Green Religion: Daoism and the Quest for a Sustainable Future. New York: Columbia University Press, 2017.
- Morozov, Evgeny. To Save Everything, Click Here: The Folly of Technological Solutionism. New York: PublicAffairs, 2013.
- Plumwood, Val. Feminism and the Mastery of Nature. London: Routledge, 1993.
- Rolnick, David, Priya L. Donti, Lynn H. Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, et al. “Tackling Climate Change with Machine Learning.” ACM Computing Surveys 55, no. 2 (2022): 1-96.
- Sadowski, Jathan. Too Smart: How Digital Capitalism Is Extracting Data, Controlling Our Lives, and Taking Over the World. Cambridge, MA: MIT Press, 2020.
- Strubell, Emma, Ananya Ganesh, and Andrew McCallum. “Energy and Policy Considerations for Deep Learning in NLP.” In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 3645-3650. Stroudsburg, PA: Association for Computational Linguistics, 2019.
- United Nations Environment Programme. Navigating New Horizons: A Global Foresight Report on Planetary Health and Human Wellbeing. Nairobi: UNEP, 2024.
- Watson, Burton, trans. The Complete Works of Zhuangzi. New York: Columbia University Press, 2013.
- World Economic Forum. AI for the Planet: Harnessing Artificial Intelligence for Environmental Sustainability. Geneva: WEF Insight Report, 2024.