Sustainable AI Research? Three Approaches Within High-Performance Computing Research Communities and a Critique from the Margins

Anne Pasek

Trent University, annepasek@trentu.ca

Brian Sutherland

University of Toronto

Abstract

High-performance computing research, particularly concerning so-called artificial intelligence, is increasingly scrutinised for its environmental impacts. Many researchers are actively debating how to respond. Surveying recent English language research community publications, this paper identifies three distinct but overlapping proposed approaches to greening compute-heavy machine learning research: 1) efficiency, 2) the net benefit of downstream impacts, and 3) virtuous competition through environmental metrics. Alternative, and much more radical, proposals can be found in a small but growing contingent of scholars who seek to 4) subordinate research questions to planetary limits. Each approach is more than merely technical, with significant considerations given to researchers' social incentives and relational commitments to one another. Nevertheless, all approaches are constrained by limitations in enumerating compute and mobilising social incentives. This article assesses tensions across these approaches and concludes by connecting them to wider discussions on the constitutive and increasingly planetary character of method in STS.

Keywords: AI, Research Methods, Sustainability, Environmentalism, Climate Change, Epistemic Virtues

Introduction

The climate impacts of research are a growing site of social concern. While early iterations of these debates focused primarily on the outsized role of air travel in the circulation of scientific knowledge (Berné et al., 2022; Katz-Rosene & Pasek, 2024), the rapid rise of big data and machine learning research (often packaged as 'artificial intelligence' or AI)[1] has brought the environmental implications of high-performance computing under new scrutiny. The training stages of cutting-edge neural network models, for instance, typically require hundreds or thousands of hours of energy-intensive computational work, located in data centres hosting unprecedentedly large datasets, and potentially spread across multiple geographic regions (Bian et al., 2024). This results in significant climate and water impacts. Simply training OpenAI's now-outdated GPT 3 model, with its estimated 175 billion parameters, may have produced 552 tonnes of CO_2e emissions (Luccioni et al., 2023), the equivalent of driving more than 100-gasoline powered cars for a year. A conversa-

tion with ChatGPT, a recent study concludes, consumes roughly the equivalent of a 500 ml bottle of water (Li et al., 2025).[2]

While these figures have been highly circulated by news media, posing public relations challenges for tech companies, the problems that they index have also generated significant and unresolved debates within compute-heavy research communities, on their own terms. Although current practices in the field are generally understood to be unsustainable by practitioners (Michael et al., 2023), there is as yet no consensus on what sustainable compute-heavy machine learning research would look like, nor how it should be achieved. At times, scholars have advanced this project through personal reporting, with some researchers including the carbon emissions of their computationally-demanding modelling work in their academic papers as a matter of ethical disclosure (Wu et al., 2022; Zhang et al., 2022). Other researchers have drafted preliminary green guidelines that might scale up to more collective or structural changes, including early open source guides (The Turing Way, 2022), sets of principles and best practices (Lannelongue et al., 2023; Nazaré et al., 2023; Schwartz et al., 2019), and calls for formal guidance from funders and regulators (Bashir et al., 2024). In so doing, these interlinked communities are beginning to articulate early questions and approaches that may come to define a new epistemic virtue around climate-informed modes of computational work and scientific practice: one where 'good' research is 'green' research. This paper offers a map of emerging and at times conflicting approaches to achieving this goal.

'Epistemic virtues' are a key means to this end. They describe normative moral commitments that are internalised and practiced by individual researchers within scholarly communities. They are a shared sense of what exemplary knowledge production should look like, and how knowledge production articulates across both inner dispositions and external conduct as well as individual values and incentives versus collective outcomes (Uygun Tunç & Pritchard, 2026). They are a key part of wider epistemic cultures that characterise global research fields or regional communities beyond simple disciplinary identifications (Knorr Cetina, 1999). Epistemic virtues have histories, intersect and travel across different disciplinary fields in different ways, and tend to accrete and overlap rather than sweep each other aside in grand paradigmatic reversals (Daston & Galison, 2007: 18). Within and beyond scientific communities, they must be cultivated, argued, taught, and internalised (Van Dongen & Paul, 2017). The 'how' of the research is as important as the research outputs to this end because method is understood as constitutive to the membership and character of a field.

These themes are deeply germane to both long-standing and recent questions in science and technology studies. In addition to investigations of how politics are socially admitted or bracketed from the domain of scientific practice, and the relatedly social implications of given technical modes of measuring the world (Porter, 2020[1995]; Shapin & Schaffer, 1985), the intersection of epistemic virtues and sociotechnical accounts may also extend to more contemporary theoretical investigations of the role of infrastructures (Jensen & Morita, 2015) and environments (Goldstein & Nost, 2022). In the context of constraining energy, water, and mineral supply chains (Cooper, 2023; Hogan, 2015), shifting orientations towards the indexicality and uses of data (Alvarado, 2023; Bender et al., 2021), and a deepening climate crisis (Mollen et al., 2026; Pasek, 2019), data infrastructures provide an increasingly important foothold into questions of technoscientific power and practice – and thus potential debates around in-group norms and ideals.

As this article demonstrates, high-performance computing practitioners are themselves actively debating such questions. The distinctions and boundaries of potential green machine learning virtues are not yet fixed, nor is it clear how such things might coexist with previous models of 'good science.' Participants in these conversations call on often contradictory and incongruent theories of change and moral conduct when articulating what sustainable machine learning research methods might be and how to bring them about. Moreover, the empirical fact of the emissions intensity of a given action in a digital network – such as the carbon footprint of storing a terabyte of data for a year, or that of a training run for a neural net – is often up for debate, owing to unsettled methodological tensions in how

the carbon footprint of a discrete action within a digital network should be disaggregated and assessed (Pasek et al., 2023). This makes it more difficult to calibrate degrees of concern and urgency across scholarly (and more than scholarly) communities, and to cohere emerging ideals into practice. It also brings the ostensibly distant figure of the dispassionate scientist into new proximity with contestations around the political economy of the cloud, as well as the climate trajectories of the wider information and communications technology (ICT) sector and struggles for social equity measures within and beyond the university.

This article assesses the state of this internal conversation by surveying recent academic and grey literature publications within compute-heavy machine learning research communities forwarding collective environmental ideals and strategies. Building on previous analyses of 'sustainability' as an underdetermined and contingent site of meaning making and political struggle (Alaimo, 2012; Vaughan et al., 2023), we explore how these efforts are framed as both technical problems and questions of more-than-technical social coordination. In other words, to dismiss these efforts as mere technosolutionism, even if one finds limitations in the scope or likelihood of their climate mitigation potential, would be to miss the important analysis this community is offering about itself as a community, and the group relations and pressures therein. Such a gesture would also overlook the interesting ways in which these social analytics are entangled in ever-more planetary scales of scholarly actors – both human and non-human – and thus decisions about whether and how to admit or obfuscate these networks in proposals for environmental action.

We describe three common and overlapping strategies in the mainstream machine learning research literature – relative efficiency, net benefits, and peer competition through green metrics – and analyse the strengths and limitations to each, both in terms of technical feasibility and the social affordances they point to within their fields of study. We then contrast these approaches with a degrowth ethos that can be found in an emerging 'radical flank': a small but growing set of researchers interested in subordinating the question of computational design to that of planetary limits. In contrast with accounts of past epistemic virtues (Daston & Galison, 2007), we show how emerging strategies for forwarding green high-performance computing are largely predicated on the social composition of the research community as a group, driven by relational rewards and constraints, rather than as a matter of individual dispositions. This analysis leads to a concluding discussion of the constitutive character (and carbon intensity) of research methods more broadly, bridging insights from science and technology studies to better explore how emerging digital environmental practices might forge both different kinds of social research cultures and different kinds of planetary configurations made in and through research.

Literature review

Epistemic virtues can provide a particularly useful analytic foothold for science and technology studies (STS) scholars looking to parse the interrelations between technical and social systems. While the field takes the presence and influence of these connections as both a foundational given and an enduringly important object of study, many scientists themselves have at times been reluctant to acknowledge such influences, for fear that doing so would threaten the dispassionate and apolitical character of an imagined objectivity (Harding, 2015). This tendency can be particularly acute in computer science, where formal structures of abstraction (Golumbia, 2009; Malazita & Resetar, 2019) and the bracketing of ethics as a tertiary concern in computer science education (Darling-Wolf & Patitsas, 2024) frequently serve to deepen an imagined disconnect between the design of digital systems and their socio-material outcomes – especially in the domain of 'pure' research. Yet, as STS scholarship has long affirmed, such boundaries are harder to maintain during periods of shifting epistemic foundations (Shapin & Schaffer, 1985). Debate admits internal advocacy within expert communities, self-reflection about how a discipline's political and material orders are organised, and theories about how to change them.

Contemporary environmental concerns broached within the research community, and the

ways they cut across previously assumed spatial, temporal, and social boundaries that separate the researcher and the researched, may serve to accentuate this tendency. When the techniques of knowledge production contribute to the environmental problems researchers are trying to isolate and understand, scholars will find themselves – often uncomfortably – within the frame of what they are studying. In this way, considerations of the climate impacts of the research process can provoke ethical and epistemic challenges to the straightforward pursuit and dissemination of knowledge from a position of remove, and through any means (Knox, 2020). The result can be both new forms of responsive activism within expert communities (Anderson, 2013) as well as efforts to refuse these demands and maintain previous community norms and practices (Bendell, 2016). No given outcome is guaranteed from these moments of destabilisation, but they do provide provocative calls to reassess the boundaries between researchers, research infrastructures, and the material benefits and harms they distribute across the world.

Researchers in the social sciences and humanities are not exempt from these dynamics. A growing body of work on Anthropocene questions, the energy humanities, and climate history has also provoked positional reflections in a range of fields, including STS (Latour, 2004; Puig de la Bellacasa, 2017), history (Chakrabarty, 2009), literature (LeMenager, 2014), media studies (Peters, 2015), and anthropology (Knox, 2014). This work's redoubled emphasis on the causal importance of macro-scale material forces—whether oil, climate systems, or geology—frequently point to the role environments can play as a third axis modifying accounts of socio-material inter- and intra-actions. In some schools of thought, the figure of the 'planetary' helpfully encapsulates this analytic (Clark & Szerszynski, 2021). As argued elsewhere in this issue (Çelik et al., 2026), planetary analyses of knowledge formation challenge the sufficiency of both ontologically flat actor-networks and socio-technical dyads. Instead, taking planetary processes seriously requires a multidimensional accounting of entangled relations across inequitable social groups, global locations, and geophysical potentials and harms. It is to grapple with the fact that all researchers are enmeshed within the worlds they seek to describe, with shared (if never equal) relational responsibilities and consequences across the work of knowledge formation (and effects) that are simultaneously social, material, and planetary.

Computer Science as a whole has yet to fulsomely grapple with the planetary, having (understandably) first taken global data infrastructures as a collective accomplishment rather than a cause for ethical self-interrogation (Edwards, 2010; Ensmenger, 2013). The development of abundant and personal compute resources is a story of increasing accessibility, with career stages playing less and less of a role in mediating the finite use of research machinery (Hu, 2015). This access story is not complete, however, with enduring Global North/South divides characterising the availability and performance of both vernacular digital tools (Shahjahan & Niloy, 2026) and specialised compute infrastructures (Lehdonvirta et al., 2024). With the growing importance of scale and cutting-edge graphical processing units (GPUs) in machine learning research, moreover, progress on research infrastructure accessibility may be backsliding along with the climate targets and efficiency measures that characterised previous periods of digital industry self-governance (Pasek, 2023).

A planetary analytic, taking stock of local material constraints, meso-scale social group formations, and macro-scale environmental pressures, therefore has increasing salience to accounts of compute-heavy research. As we describe below, there are indeed moments where debates in the field bend in this direction. In the face of spiralling compute costs, carbon emissions, and public critiques, machine learning researchers have good cause to think about human and non-human planetary relations and the role of both in-group incentives and external environmental constraints in adjudicating new styles and sensibilities of being a high-performance compute researcher. Yet, as previous examples stress, these openings also come with anxieties, sacrifices, and reactionary efforts to prevent hierarchical destabilisations or even minor headwinds in the pursuit of knowledge. As such, this paper seeks to parse how different voices within the field's centre, or alter-

natively commenting from its margins, propose new settlements in the figure of green research along both social, material, and planetary lines. Reading with voices from the field, we ask: what kind of problem should climate change be to the field, and what new kinds of epistemic virtues and relational commitments could help resolve that challenge?

Four approaches to sustainable high-performance computing

This paper's assessment is distilled from reviewing a corpus of over 60 articles and white papers, published between 2018-2024, that address the question of green epistemic virtues from within English language machine learning (or so-called 'AI') communities. These texts were assembled through a combination of digital literature searches, informal recommendations from practitioners and co-panellists at public events and conferences, and tracing relevant citations in bibliographies. Based on the centrality of these texts within citation networks and major research institutions, the corpus represents a significant portion of the influential work published on this topic, but does not attempt to be exhaustive in scope or totalising in its claims. Instead, the analyses that follow were formed through interpretive content analysis with the goal of identifying common heuristic patterns and exemplary arguments about unsettled empirical and ethical questions in the field. It thus offers more of a social cartography (de Oliveira Andreotti et al., 2016) than a meta-analysis.

For each section, we begin by describing the proposed approach through one to two exemplary cases, and then discuss the mechanisms by which the approach seeks to achieve its technical (questions of definition, measurement, infrastructural design and utilisation, etc.) and social (questions of attention, prestige, opportunity, and training, etc.) ends. By splitting these different means and motivations apart in our analysis, we aim to better demonstrate overlapping appeals and pressures in differing socio-material accounts of digital sustainability—how each approach to green epistemic virtue articulates simultaneously across technical and community relations. We also aim to stress how planetary concerns threaten to complicate both sides of this sociotechnical dyad (for instance, in the differential global distribution of renewable energy resources, or enduring Global North/South prestige asymmetries) and are variably rhetorically mobilised or bracketed from the framing of solutions. Each section concludes with our assessment of internal limitations to these proposals and external barriers to their enactment.

Environmentalism through relative efficiency

Often the first approach to mobilising around new problems of practice is to draw clear distinctions between favoured and disfavoured conduct. Schwartz et al.'s highly influential intervention (2019) does so with admirable simplicity by splitting so-called artificial intelligence research into two basic categories: Red AI and Green AI. This intervention well characterises the broader emerging uses of efficiency as a potential green epistemic virtue for the machine learning research community.

Red AI is defined as any research effort that privileges the accuracy of deep learning models as the sole metric of success. Researchers today typically pursue this goal by means of exponentially increasing the compute resources (in both datasets and training time) they use in model construction – in effect, a brute force approach to get better results. Schwartz et al.'s survey of deep learning papers published in central journals to the field from 2013-2019 demonstrates a worrying dominance of Red AI trends through a 300,000 fold increase in compute over this time – a rate that exceeds that of Moore's Law (Schwartz et al., 2019: 56). Because Moore's Law provides a somewhat predictable pace for energy efficiency improvements in wider ICT industry infrastructures (Pasek, 2023), rapid leaps forward in research infrastructure requirements represent a distinct and worrying environmental trajectory: the field is getting dirtier faster than the industry is getting cleaner.

Green AI, conversely, would be any form of machine learning research that breaks this trend, and so does not require exponentially greater amounts of compute resources to produce novel results (Schwartz et al., 2019: 59). Ideally, this environmental approach would go so far as to reduce, rather than just maintain, current computational

requirement rates, but definitionally Green AI is simply understood as any project that uncouples the strict correlation between academic success and resource expenditure – that eschews a brute force approach. It is thus an orientation towards conservation through process efficiency, rather than a commitment to meeting absolute targets or an articulation of particular conditions of sufficiency, that separates Red AI from Green AI. It suggests that sustainability can ultimately be determined through relative measurement.

Understanding sustainability through efficiency has its social advantages. The Red vs. Green designation provides a usefully clear conceptual intervention and provocation (the paper has been cited more than 3,000 times at time of writing). It also allows for a moral defence of machine learning research's past, in tandem with a critique of its future. Red AI, Schwartz et al. (2019: 54) are at pains to emphasise, "has been yielding valuable scientific contributions to the field"; the only problem is that "it has been overly dominant." By arguing simultaneously for a relational boundary between sustainable and unsustainable research practices, and by seeking only to adjudicate the relative distribution of each, it both brokers space for action and accommodates a plurality of approaches. Projects pursuing benchmark improvements at any cost are still potentially welcome, just not as the sole direction of the field. In this way, past epistemic virtues can coexist with emerging green ones, in a range of favoured scientific practices. This may pre-emptively defuse potential critics who remain committed to Red AI approaches, or are at least apt to object to a condemnation of both their extant research methods and research products. Similarly, by allowing researchers to respond to climate imperatives with quantitative precision, via a range of mathematically-modelled efficiency algorithms (see, for example, Bian et al., 2024; Li et al., 2025; Silva, 2021), this approach allows for familiar tools to be brought to bear (unlike the rather more qualitative issues with algorithmic gender, racial, and disability bias that the field continues to struggle to navigate).[3]

The relative efficiency approach also allows for proponents of environmental reforms to make arguments about the social impacts of exponentially increasing resource needs within the composition of the research community. This too can be strategic; if environmental and social concerns are aligned, then a greater global coalition of interested parties might be rallied to the cause of greener research norms. In the account of Schwartz et al. (2019: 56), Red AI represents such a threat in so far as it has come to define the field, "through the use of massive computational power while disregarding the cost – essentially 'buying' stronger results". The Red AI, brute-force approach can produce diminishing returns – a 10-fold increase in deep learning model performance may require as much as a 10,000-fold increase in compute resources (Thompson et al., 2021) – which effectively prices out junior and Global South researchers who lack access to institutional or grant-subsidised compute clusters (see also Bender et al., 2021; Strubell et al., 2019). It also potentially diminishes the role of creativity in research, rendering research advancement, to a degree, as another 'pay to publish' game with a computational processing twist.[4] Moving against this trend, conversely, would potentially lower barriers to entry and better value the cognitive ingenuity of researchers. As the authors write, "we want to shift the balance towards the Green AI option—to ensure any inspired undergraduate with a laptop has the opportunity to write high-quality papers" (Schwartz et al., 2019: 54). This appeal to the future of the discipline, cast in the face of a young and deserving student, can thus be read in jointly moral and environmental terms (even if said student is themselves not primarily motivated by the environmental call to action). The epistemic virtue here therefore adheres to group norms and outcomes, not individual moral values or judgements.

Shared social coordination, as well as flexible horizons for the achievement of green ends, are thus both rhetorical and practical benefits that a relative efficiency approach – rather than restrictive targets or caps – offers in adjudicating the definition and trajectory of sustainable research methods. It suggests that the problem of climate change for machine learning researchers is one that demands collective coordination and that is simultaneously implicated questions of academic equity yet can ultimately be addressed through technical means that mitigate these parallel relational and planetary pressures. Infrastructural

expenditures drive these questions forward, but in ways that may ultimately be resolved at the scientist's desk, if enough coordinate to do so.

There are, however, practical downsides to using relative efficiency to guide this discussion. Framing the question around the ratio between computational inputs and outputs risks ignoring material complexities that extend beyond the discrete performance of a given compute cluster at a given moment in time. By some estimates, 66-75% of the energy needs of a model come from the sundry work of inference and data management and storage, not use or iterative refinement, bringing a singular emphasis on model training efficiency into question (Luccioni et al., 2023; Wu et al., 2022). In other instances, being strategically selective about the kinds of metrics (and best-in-class infrastructures) can seemingly counterbalance the exponential consumption concerns of Red AI (Patterson et al., 2022). Further questions have been raised around the wider sectoral build out of data centres, and the extent to which the promise of data-intensive applications and research agendas both outpaces, and induces, demand (Hogan, 2021). There are also growing concerns about potential rebound effects, whereby increases in the efficiency of a resource makes it less precious, and thus apt to be used more in aggregate (Lannelongue et al., 2023). The politics of infrastructure, especially at the fenceline or through regional grids that host data centres, is not especially legible in this approach. A future full of machine learning tools and datasets that no longer outpace Moore's Law could still be one with growing aggregate energy demand and worsening environmental impacts. In such instances, models of sustainability exclusively developed through relative resource consumption rates internal to the research community may prove inadequate to the task of assessing and leading a way through the planetary problem. They may, to the contrary, obfuscate it (Cellard et al., 2025). Green AI, after all, isn't actually green; it's just less red.

Being green, on balance

The ethos of efficiency that characterises Green AI combines with an emphasis of social impact accounting in the second approach that we highlight. This school of thought draws different boundaries around where the adjudication of sustainability can be said to begin or end. By considering the wider effects of compute-heavy research when deployed in the world, researchers and advocates can more effectively assess its net effects. As such, instead of creating a conceptual cut between favoured and unfavoured scientific conduct, the net effects approach advocates that all research, overall and on balance, should be green. Importantly, this provides a way of sanctioning research efforts that are themselves highly resource intensive and even relatively inefficient, but which have overwhelmingly positive downstream environmental benefits.

The Royal Society, an organisation representing UK scientific interests, exemplifies this approach in a report that was developed out of several workshops about digital knowledge production and emerging technologies facilitated between researchers, government, and industry (The Royal Society, 2020). Highlighting both the potential benefits and harms of compute-heavy data products and programs of inquiry, the report urges considerations of "energy proportionality" in the moral and administrative judgements of digital research processes and outputs, and further calls for regulators to develop guidance to these ends (The Royal Society, 2020: 80). This echoes a larger trope in the literature (Kaack et al., 2022; Luers et al., 2024) and public discussions (Lohr, 2024) that sees machine learning tools as a double-edged sword (Cellard et al., 2025): a catalysing set of technologies that may result in both hazardous and utopian outcomes, and as such requires sober and skilful oversight.

Proportionality implies a measured consideration of the costs of training, constructing, and using digital resources against the potential energy or carbon savings that an application might enable. For example, an algorithmically-controlled building heating and cooling management system may be very resource-intensive to train and run, but could potentially enable considerably more emissions reductions over the course of its lifetime. Alternatively, a high-performance computer system may not reduce emissions, but its development could nevertheless still be considered sustainable and ethical through a combination of compensatory carbon offset purchases, green energy procurements, and perhaps even

social equity measures. A tacit net zero calculus is afoot here, with the positive impacts of a technology's application all roughly imagined to count against the climate (or social) harms of the research and technical development process. This calculus, however, remains largely figurative; the Royal Society report declines to name specific examples in which energy proportionality has not been achieved, with the sole exception of Bitcoin (The Royal Society, 2020: 81).[5]

A figurative accounting approach is also a social proposition. By highlighting intended long-term positive outcomes alongside harmful intermediate means, the net balance proposition emphasises researchers' best intentions and future redemptive possibilities while potentially palliating anxieties about the unsustainability of compute-heavy technology as a whole. It also reduces the potential for individual researchers themselves to be judged harshly for their work, positioning them merely as the developers of 'tools' whose effects will be largely determined by "the context in which they are developed and deployed" (The Royal Society, 2020: 114). Researchers, in other words, can exercise responsibility but need not be held wholly responsible. This in turn sanctions the recommendation for governments to increase the funding of "high-risk, high-reward" data science and machine learning inquiry, based on anticipated if not wholly detailed environmental and social benefits (The Royal Society, 2020: 111). Here, again, a larger collective guarantees the production of epistemic virtues – one that includes more than just the research community itself.

This net impact approach thus has further implications for the relationship between researchers and the wider cloud industry. The work of sustainability, in effect, can be largely outsourced or externalised to the providers of compute resources and the remediators of e-waste rather than the designers and users of machine learning systems. Google researchers Patterson et al.'s (2022) set of sustainability best practices, for instance, emphasises the use of best-in-class infrastructures like Google's cloud network and the renewable energy and carbon offsets purchased by Google, Facebook, and Microsoft.[6] Google's environmental practices in its data centre design and load shifting work are further highlighted in the Royal Society report (2020: 78) along with calls for the industry to "lead by example" (The Royal Society, 2020: 5). In both instances, efficiency gains from researchers, combined with the green leadership of private industry, combine to potentially negate the problem of escalating digital climate impacts. This all serves to reframe the problem of climate change for the field as one of good and bad infrastructures, rather than good and bad research efforts. As a result, machine learning research methods can themselves be seen as largely neutral, rather than a key site of struggle towards different planetary outcomes. Put otherwise, green virtues can be deferred to compute providers and technological entrepreneurs; they perhaps need not be epistemic at all.

As in debates about relative efficiency, the net accounting approach also presents many practical and political limitations. For one, industry self-regulation is likely an insufficient strategy to get to truly sustainable compute infrastructures, however this might be defined. Flagship environmental data science projects are often subsidised by cloud companies as a means to win social license for broader data-intensive business strategies, rather than as a credible means of performing environmental governance within and beyond the sector (Espinoza & Aronczyk, 2021). A focus on providing advanced models and compute infrastructures to environmental and conservation scientists can also obfuscate more fundamental scientific needs for basic data collection in the field (Lukacz, 2022), especially in the Global South (Karlsson et al., 2007). What's more, under the semantic confusion of the category of 'artificial intelligence,' benefits specific to more narrowly scoped machine learning techniques are frequently used to justify the costs and impacts of more general purpose technologies (Joshi, 2026). Not unrelatedly, the ICT sector's absolute emissions continue to grow because of accelerating generative "AI" market bets, even within the ledgers of the industry's seemingly most exemplary corporate climate actors (Saul et al., 2024). State governance, in line with these aims, remains largely absent. Additionally, efforts to monitor and manage global climate and environmental impacts often ignore the continuity – or creation – of local environmental burdens and

frictions (Pasek, 2019; Sefyrin & Velkova, 2026), sometimes replicating familiar environmental racism dynamics (Bender et al., 2021: 613). Net benefits *for whom* remains an essential, though generally underexamined, question in such debates.

Additionally, while the emissions savings potentials of emerging big data technologies are highlighted extensively in defence of proportionality judgements, the field's parallel role as an accelerant of climate change is often not discussed to the same degree. Given the growing centrality of machine learning in new oil and gas exploration (Galaz, 2025: 42; Zero Cool, 2019), these are not trivial questions. They are also potentially incalculable—especially on the timescales that matter in weighing precautionary regulatory interventions. A recent attempt to outline a framework for assessing the present-day climate impacts of generative machine learning algorithms describes this work as a "seemingly intractable task," the difficulty of which "cannot be overstated" (Bashir et al., 2024). Speculative benefits and harms in the future represent yet a further order of difficulty – one that may be far more qualitatively than quantitatively thinkable, and thus will not reconcile with existing life cycle assessment frameworks (Kaack et al., 2022). Finding a usefully descriptive and expansive net impact number may ultimately not be possible. In this context, calls for proportionality guidance could risk deferring questions of research norms and regulatory action to further and future stakeholders that may not be equal to the task. As critics have argued, a focus on AI ethics in insular terms can ignore the wider consequences of the mass adoption of these technologies in terms of how power is distributed across companies, people, and the planet (Crawford, 2021).

Eco-competition

A third approach to green epistemic virtues is imagined to proceed through the promotion of best practices that all researchers should follow to ensure that they fall on the general side of sustainability. However, instead of setting general efficiency or net benefit criteria, this school of thought seeks to create new metrics and areas for comparison that grant prestige to a range of greener research approaches. Measurement, in this context, is both technically and socially relational. It implies that one can be ahead or behind the group (rather than merely in or out of a 'green' category), and that making these distinctions visible within communities of practice is key to effecting change. It thus suggests that the problem of climate change for the field is one that requires not only technical innovations but also forms of social competition.

Researchers imagine that new forms of awareness and incentives can be cultivated through multiple communicative channels internal to the scientific community. These include academic publications; several scientists (Lannelongue & Inouye, 2023; Rafat et al., 2023) have argued that carbon footprint estimates of machine learning research processes should be published along with new findings and models, reinforcing the salience of the climate issue within the research community and stoking healthy competition. Patterson et al. (2021) suggest that conferences might favour the publication of papers on efficient or small models in specific green tracks. Perhaps most famously, Henderson et al. (2020) argue for the creation of leaderboards to better track and evaluate progress in a research group, classroom, or wider field, and have released code on Github to facilitate this. Leaderboards are already a common feature of machine learning research, tracking performance across speed and size concerns. Adding carbon intensity to these metrics could be a simple technical means to stimulate a different kind of race.

This highly social model of environmental change seeks to harness peer pressure and scholarly recognition within the research community in order to better value more sustainable forms of compute-heavy inquiry, creating a "virtuous cycle where ML [machine learning] practitioners compete to… bend the curve of ML carbon footprint growth," (Patterson et al., 2021: 2). The eco-competition approach allows for flexible frameworks and comparisons. For example, given the range of relative differences between the carbon and financial costs of compute resources in different parts of the world, researchers may struggle to contract with the most efficient or environmentally well-provisioned set of cloud

providers. In such cases, they could choose to collaborate with researchers in other areas of the world who do enjoy such access (Lannelongue et al., 2023) or simply pick from alternative actions on a list of green tips.[7] Alternatively, researchers could simply publish their results, with the high carbon intensity of their local compute resources given as a caveat. In this way, even simply participating in the creation and publishing of carbon metrics can be read as an action – contributing towards social coordination and collectively altered norms (if not yet a direct carbon mitigation effect).

Like the previous two approaches, this eco-competitive outlook on sustainable heavy-compute research may also create grounds for inertia and contradictory outcomes. For one, even if the mass reporting of the energy and carbon costs of developing and/or running high-performance models became common practice, it doesn't immediately follow that accurate comparisons could be easily drawn across the heterogeneous contexts of different research architectures, nor that social pressure would be effectively mobilised on the basis of such accounts. There is no commonly agreed methodology for such calculations. There are, instead, many different ways of muddling through such questions, with significant trade-offs between accuracy and ease of use (Lannelongue et al., 2023). While the accounting protocols of some cloud compute providers, having more standardised data centres across different global locations, make it relatively simple to call the run time, and even energy use, of processors used to train a model, the carbon intensity of real-time, local grid power can often still be difficult to correlate (Lannelongue & Inouye, 2023). The Henderson group's leaderboard code, for example, is only compatible with a narrow range of systems and chips, and draws on data from Amazon Web Services that have not been independently verified. It is also the case that such narrowly cast assessments often ignore the larger lifecycle costs of manufacturing and maintaining equipment outside their window of use by a given researcher. In some accounts, the carbon footprint of model training can double when attempting to include these 'external' factors, inferred as they are from imperfect manufacturer disclosures (Luccioni et al., 2023). Further complicating matters, efforts to spatially and geographically shift computational work to minimise carbon emissions (in effect, aligning compute with solar energy peaks) can inversely increase water usage from cooling (Li et al., 2024). Optimising for one environmental outcome can thus come at the cost of another. Making measurements into targets, as researchers in diverse fields have proven, often renders those targets ineffective as governance measures, and vice versa (Archer, 2024; Strathern, 1997).

It's also the case that a fair competition in a privileged track for environmental performance may not be guaranteed. For one, the ease of achieving green benchmark performances will vary across the planet, given the unequal geographic distribution of green(er) energy infrastructures and the price premiums of working with best-in-class data centres. Absent a common standard, moreover, there is an incentive for researchers who do choose to report their carbon footprint to select the most favourable means of determining this figure, or to merely report on the best of many runs of a model (Schwartz et al., 2019: 62). An honest competition is not guaranteed and, as other scientific fields remind, the replication of results – even by the same research team – can be an enduring challenge (Baker, 2016). Machine learning research, with its substantially sized training datasets, code bases, and variable factors impacting model training, may raise the gambit of the challenge considerably (Desai et al., 2025).

A final, and perhaps more interminable, complication lies in the vexed relation between information and action. Despite growing awareness around green computing principles, a large gap between ideals and practice still characterises the machine learning field today (Nazaré et al., 2023). The virtuous cycle has yet to take off: the publication of carbon footprints is not presently a widely adopted norm, nor have energy-reduced model tracks revolutionised publishing incentives within or beyond machine learning conferences. This dynamic is not specific to high-performance computing research; the degree to which awareness compels, or is even terribly correlated to, behavioural changes is frequently overestimated (Moser, 2015). Yet, as an additive and voluntary metric, it is fair to say that environmental

considerations still lack the social and institutional pull of other modes of evaluating success. At least for now, the wider research community continues to value (and socially recognise) accuracy and speed more highly than sustainability.

Critique from the margins

While these three prior approaches to developing a green epistemic virtue for machine learning come from scholars working in the field's core, there are also additional perspectives at the margins, looking in. A growing body of critical scholars, digital artists, and computing practitioners are questioning dominant directions and articulating a radical alternative: an ethos not of efficacy, net positive impacts, or peer competition, but of degrowth: of subordinating the design and function of digital systems to external ecological limits.[8] In so doing, this group forwards directions and alternatives that are otherwise unthinkable within the mainstream of resource-intensive computing. They generally do not offer complete solutions for researchers seeking to green their extant research programs, but rather aim to question whether or not such dominant directions are inevitable, desirable, or justifiable. This research proceeds through a range of methods, including speculative design, design fiction, and multispecies design (de Winkel et al., 2020; Nilsson et al., 2024; Santos et al., 2021; Soden et al., 2022; Tanenbaum et al., 2016), testing what is technically possible and exploring what is socially emergent from drastically different computing arrangements. While dialog between this group and those we explored previously is currently minimal, we highlight them here as a useful foil, looking to the margins in order to bring the omissions and common-sense perspectives of the dominant sustainability conversations within high-performance computing into sharper relief.

Participants in discussions about degrowth computing come from a diverse set of research contexts. This includes subdisciplinary considerations (human-computer interaction, media and information studies, and integrated design are especially common) as well as sites for research exchange. Two conferences in particular – the Workshop on Computing within Limits (LIMITS) and Information and Computing Technology for Sustainability (ICT4S) – are important sites for this discussion. Across a plurality of approaches, this community articulates a politics of ecological limits for the future of computing in a warming world, marking a break from other approaches to 'green ICT' (Nardi et al., 2018). Rather than struggling to engineer a way around environmental constraints or otherwise justify the sector's impacts, these thinkers start from the assumption that research agendas and social expectations for digital systems – especially those associated with machine learning – should be oriented towards a future that will have stricter environmental restrictions.

This has significant implications for the 'cornucopian model of computing' that characterises contemporary digital systems (Abbing, 2021). Instead of creating network, storage, and processing capacities that are overbuilt so as to be available in all circumstances, regardless of resource costs, degrowth computing proponents strive to prototype, imagine, and analyse infrastructures of computational sufficiency – contexts where bandwidth and server demand has plateaued, energy inputs are largely intermittent and renewable, and equipment is maintained and recycled over much longer time periods across decentralised networks that more equitably serve a larger global community of users with heterogeneous hardware and bandwidth capacities.

Two case studies exemplify the characteristics of ecological limits for compute infrastructures and the technical and social analyses that follow. The first is Solar Protocol, a network of small, raspberrypi servers, hosted by an international group of stewards and powered by local, off-grid solar energy. Developed by Tega Brain, Benedetta Piantella, and Alex Nathanson, these servers are connected as a shared digital art web platform through a custom networking protocol that delivers content to an Internet audience based on whichever server currently has the highest battery charge (and thus is located in the sunniest location on the planet at the given time) (Piantella et al., 2020). This differs significantly from both mainstream internet protocols, which route content from globally distributed cloud servers based on network speed, and the dynamic carbon-responsive compute distribution algorithms proposed by mainstream researchers and industry, which

seek to factor energy prices, carbon budgets, and latency concerns into the flexible scheduling of large data processing tasks (Bian et al., 2024; Ren et al., 2012). Against the impulse to design additional layers of algorithmic complexity in pursuit of energy efficiency, Solar Protocol instead subordinates the routing and performance of its system to the availability of the sun. Servers either are, or are not, online. They either do, or do not, have the highest battery charge in the network. The sun, accordingly, picks the server. Planetarily distributed computing, in the form of a necessarily global volunteer network, makes this dynamic relay possible. The infrastructure, rather than the infrastructure's users, enact the moral argument.

Solar Protocol's developers provocatively describe this approach as one that follows the "natural intelligence" of environmental systems, in contradistinction to "artificial intelligence," and speak to the virtues of contouring human behaviour and system design around its "logics" (Brain et al., 2022: 5). For the platform, this means designing with intermittence and energy-constraints in mind: server-side static site generation with fairly low-resolution graphics that can fit on a device that can be powered by a relatively small residential solar panel. As a byproduct, the platform involves no adware, blockchains, or trackers. The developers further note, how "transitioning towards the types of computational systems that renewable energy regimes can support… [can provide] rich opportunities for reimagining and redesigning the ways we want computation to be in our lives" (Brain et al., 2022: 4). The fact that many data-intensive tools are not easily reconciled with the energy and hardware constraints of Solar Protocol is presented as a useful opportunity to shed attachments to environmentally and socially harmful directions in digital practice and research. The sole machine learning project hosted on the platform is also provocatively small: a collection of poems generated from hand-transcribed public domain datasets, generated with a mere 32 KB of RAM, a basic Markov chain algorithm, and a direct drive solar panel circuit (Parrish, 2022). Its size, careful data curation, and energetics are a deliberate refusal of machine learning research's dominant tendencies.

The project's platform logics point to the ways that degrowth computing considerations are more than just eco-technical. In addition to the privacy benefits that the developer team highlights by jettisoning user tracking and dynamic site design, the process of building and maintaining a globally distributed network has also involved navigating complex local entanglements with community infrastructures, care work, and skill asymmetries with the stewards who host the servers (Piantella et al., 2023), as well as shifting institutional partnerships and state restrictions on internet traffic. Absent a for-profit company or government agency to manage the infrastructures of digital research, researchers themselves must do the social work of building and sharing resources across distinct and varied local contexts. As a result, the findings of this interrogative design project have been perhaps equally legal, social, and economic as they have been energetic (Brain et al., 2022: 4). Epistemic virtues, figuratively and infrastructurally, take a village to build and maintain.

Beyond the community-building dimensions in its direct network, Solar Protocol's wider approach is politically prefigurative: by building a system that reflects design constraints and energetic norms that are not yet practiced by dominant society, it seeks to emphasise the generative work of the means of socio-technical organisation, rather than to argue only for a given end. As a design provocation, Solar Protocol has circulated widely in media arts circles and been nominated for, or won, several distinctions. It has also hosted exhibitions on its own platform, hosting artists and writers that are further exploring the role of energy-constrained design. By building community through building an infrastructural commons, it wends technical and social alternatives together in a mutually constitutive way.

The politics of degrowth, however, are about more than just cutting back extraneous consumption; they also explore opportunities for synchronicities and mutualism between sociotechnical systems. This is the focus of our second case study: 'information batteries' that intermittently use excess power from renewable energy generation systems for flexible computational work in ways that enhance the viability of increased wind and

solar generation on wider electrical grids (Switzer & Raghavan, 2021).

Renewable energy production is often characterised by periods of intense surpluses, in excess of what grid demand can support. This energy can be diverted to conventional or novel energy storage systems, such as pumped hydro, thermal, or lithium-ion batteries, which add additional financial and ecological costs to energy systems. Alternatively, it can be wasted. However, Switzer and Raghavan's (2021) exercise in 'speculative execution' involves 'storing' surplus renewable production as information products, engaging predictable compute jobs, typical of many machine learning research tasks, as a flexibly-dispatchable energy load. This is an explicit recognition of the fact that computational work, materially, is electricity – and as such, can be mobilised to address dissonant patterns of planetary and social energetics. The researchers calculate that an information-battery-enabled data centre at 90% efficiency would be commensurate with 150,000 kWh of energy storage (Switzer & Raghavan, 2021: 9), smoothing over moments of mismatched supply and demand and thereby providing additional stability and revenue for suppliers to an increasingly fossil-free energy systems. They thus demonstrate how informational work and infrastructures can be integrated into grid management, to the ultimate benefit of the grid.

Compared to Solar Protocol's prefigurative politics, information batteries may at first appear to be far less radical. Switzer and Raghavan (2021) examine existing data centres as they are currently configured, and largely through the social language of market conditions, pragmatically assessing the theoretical effectiveness of the energy costs, compute times, and caching speeds that would make their intermittent, surplus energy-contingent use economically viable to both users and providers of energy resources. In this proposal, research work and payment systems could be leveraged as a means to reduce the costs of wider energy transitions, without necessarily requiring substantial reforms to energy markets or heavy-compute research trajectories. Information batteries aim more modestly to only mobilise existing infrastructures and to avoid building expensive energy storage projects.

Yet, in so doing, this proposal also implies a shift in the relational roles of machine learning researchers, cloud companies, and the grid. Meaningfully, it places environmental and engineering challenges ahead of the immediate needs and interests of the wider ICT sector, finding ways for some computational research methods to support, rather than only increase, the labour of decarbonising electrical systems. Such a shift extends to the researcher's social position; instead of a customer, entitled to on-demand services and guaranteed run times, they are instead integrated into an infrastructural network that prioritises the grid's temporal concerns above their own. Additionally, while the authors acknowledge that many compute-intensive projects will not easily fit into this use case, it does demonstrate how those that do could be financially rewarded. This could in turn provide a way to recognise forms of more sustainable high-performance computing in more than symbolic terms, helping to address some of the extant inter- and intranational resource inequities within the research field discussed earlier in this paper.

These altered roles, and the wider convergence of energy informatics and energy transitions in research and industry, also point to the ways that many computational researchers are increasingly thinking of themselves in terms of 'energy citizenship' (Kumar et al., 2023). This is underlined by the paper's efforts to encourage researchers and companies to see energy and information as largely interchangeable – in Raghavan's words, "the core concept here is that information has an embodied energy to it" (Simon, 2022). In the information battery approach, citizenship and energy are articulated not through personal consumer austerity, but through community practices that make a researcher's compute tasks into a part of a utility's support systems. It also attunes participating researchers into the temporal and spatial dynamics of wind and sun in specific planetary locations, asking researchers to align their efforts with the extant movement of nonhuman forces instead of capturing those forces and holding them entirely in reserve.

This suggests that knowledge work might have a time and a place in which it is particularly favoured, rather than always being an *a priori* good. Along with Solar Protocol, the information battery approach proposes that the problem of climate change for the field of machine learning requires relational situatedness more than metrics, and a willingness to negotiate trade-offs between scientific goals and environmental ones. Green epistemic virtues, in other words, may depend at times more on the humbleness of the scientist than their technical triumphs.

At the same time, the Solar Protocol and information batteries concepts both demonstrate obstacles and challenges to an ethos of computing within limits. Most notably, these provocations from the margins have failed to travel terribly far into the centre of high-performance computing. Solar Protocol's suggestion that there are forms of computer science that may be simply incompatible with ecologically concordant energy systems is not likely to persuade scholars with years of professional experience and existing research programs that require heavy compute loads. Nor can it effectively compete with industry influence in the agenda-setting of computer science education (Marshall & Grooms, 2022) and in adjudicating questions of prestige and pay in career trajectories (Kirdani-Ryan et al., 2024). Prefigurative technical politics may ultimately be better at inspiring new minor traditions than challenging existing major ones. There are also, in the context of global digital networks, practical limits to how far these self-hosted experiments can go; Solar Protocol's servers still operate through corporate internet service providers and across subsea cables and internet exchange points that are designed to maximise up-time and informational speed through largely fossil fuel grids and backup generators. Interrogative design, in order to still participate in extant sociotechnical systems, cannot interrogate everything at once.

Similarly, the information batteries approach, while a technical provocation to reorganise computational expectations and work rhythms around moments of energetic surplus, ignores some of the complexities in the political economy of the cloud. Some large data centre operators, in pursuit of internal environmental goals, are already starting to distribute and schedule some workloads to align with periods of high-renewable energy production on the grid (Koningstein, 2021) and are experimenting with curtailing loads during periods of extreme energy demand (Mehra & Hasegawa, 2023). However, this pursuit has not reduced the expansion of the data centre industry, with its continued base loads, nor has it become a key practice grid managers utilise to negotiate the economics of renewable energy peaks. It has, instead, strengthened the economic role of some demand response programs in the siting and continuity of data centres; witness the rise of cryptocurrency mining businesses that make money both from drawing grid energy into their servers to generate financial instruments, and that make money (at times, tens of millions of dollars) in demand response programs by turning their servers off when energy is scarce (Sigalos & Smith, 2023). More narrowly, it is also the case that, given the heightened expense and limited working life of the classes of GPUs most favoured in cutting-edge neural net research, model training and inference companies are more economically incentivised to run this class of chips inflexibly, at all possible hours (Lacey et al., 2026). Energy responsive computing practices, in other words, may not necessarily limit the kinds or continuities of aggregate informational work. Actually decarbonising these infrastructures will require a more fundamental challenge to free market principles and the social licence of so-called AI.

This points to the problems of academic ownership and control over compute resources –something that is growing rarer as the scale of processing demands within machine learning rises – as well as the embodied carbon and environmental resources in the proliferating infrastructures of the cloud. If ICT companies are unwilling to pursue information batteries as a business case, who will, and how can this not ultimately result in ever-more landfilled silicon? Elsewhere the information battery concept's authors, along with other LIMITS thinkers, propose the construction of mobile data centres, composed of retired server racks in shipping containers than can be flexibly dispatched to areas of seasonal electrical surplus without contributing further to the embodied emissions of the ever-growing, ever-renewing

data centre industry (Switzer et al., 2021). The challenges of building such a system, according to the authors, are threefold: working with geographically variable power systems, squeezing more computational work out of end-of-life units, and coordinating across multiple scales of carbon measurement. The questions of who will pay for these infrastructures and how they will compete in a buyer's market of hyperscale compute, conversely, remain unaddressed. In Thomas Hughes' (2012) terms, the mass computation consumption problem is a 'reverse salient' where a potentially significant advance is being inhibited along a specific dimension: ICT's otherwise exponential demand for material and energy flows.

Conclusion

Within machine learning, discussions around green epistemic virtues have been defined in reaction to current exponential trends in resource use, and characterised by a range of overlapping social values, technical conditions, and forms of favoured conduct – the interrelations of which are frequently at least partially acknowledged and mobilised by proponents. Researchers at the margins have further expanded the horizons and potential ends of this transition by providing critiques and strategies predicated on degrowth thinking. As these discussions prove, the problem of climate change could come to mean many things to the research community, from proposals to better attend to in-group inequities, to externalise responsibility to parties upstream and downstream of research, to form new measures of competitive excellence, or to subordinate and enmesh the practice of science within dynamic and global energetic infrastructures.

While there yet is no consensus on what a positive definition of sustainable computing might be, there is emerging agreement among (at least a small but growing group of) proponents that some sort of green epistemic virtue should be cultivated. Across these examples, changes to in-group social recognition, along with technical improvements to cloud computing infrastructures, are still the primary means by which such a transformation is imagined to take place. Whether a question of efficiency, net benefits, competitive achievement, or computing within limits, sustainable compute-heavy research is imagined as a direction of travel for the community via many overlapping and noncoercive means, reflecting the importance of social coordination in such a transition. However, these emergent values and strategies have yet to effectively mobilise awareness into large-scale action and ideals into persuasive collective incentives for planetarily-responsive research methods.

While we have examined the barriers and limitations to each intersecting approach, we don't intend to presume that such challenges are totalising or intractable. It remains to be seen what the future progress of these efforts may win, or where unexpected openings in the sociotechnical fabric may occur. Researcher and corporate engagements with the issue may be heightened by continued media attention, personal experiences with climate disasters, or changing expectations and innovative ideas from peers. Institutional pressures, in the form of university sustainability plans, government regulations, or revised research ethics requirements may add further momentum to this cause. Economic fallout from the eventual end of over-hyped large language model tools will doubtlessly provide another. As these debates progress, further scholarship – particularly by ethnographers – will be highly valuable for assessing how administrators, scholars, and technicians in heterogeneous contexts approach the computational and interpersonal challenges ahead through shared deliberations, individual creativity, and the many forms of social coordination that exceed those captured by the print sources we have analysed here.

Finally, regardless of the ultimate efficacy of the varied approaches we analyse when it comes to actually reducing compute-heavy research's environmental impacts, and the epistemic virtues that may be required to these ends, we want to highlight their value as an analytic resource in the study of political deliberations within scientific communities. As recent research shows, climate change is exerting new kinds of pressures on the traditional boundaries drawn around academic freedoms – whether in the individual academic's pursuit of mobility and prestige goods, now cast in newly rivalrous terms (Katz-Rosene & Pasek,

2024; Sheller, 2022), or in different disciplinary and collective responsibilities to self-regulate (Hardy, 2024). As this study demonstrates, many within machine learning communities do not imagine the environmental impacts of their work to be exempt from public or internal scrutiny – explained away as a necessary cost of curiosity-driven research. In fact, these communities' professional and infrastructural entanglements with digital tech sector companies, whether as direct employers or indirect data centre providers, likely intensify these questions, even as these relationships often delimit what is imagined as obvious and actionable in the pursuit of more sustainable futures. What's more, the grounds for future collective action may be particularly good; in a recent survey of natural language processing researchers, a majority agreed that the carbon emissions of model training should be a major concern for the field (Michael et al., 2023). At the same time, only 17% agreed that scaling up machine learning models would likely help address real-world problems, and a further 67% asserted that the majority of work being published in their field was of questionable scientific value (Michael et al., 2023). The meaning of 'good' knowledge and 'green' knowledge are legibly poised to align here. This is, potentially, an eminently organisable group.

One implication of these factors for STS thinkers and environmental advocates engaging on the issue is the need to take stock of the internal analyses happening within the field. The goal cannot simply be to demand that machine learning scientists become ecologically educated, moving from ignorance to action. This work is already proceeding from within. Similarly, despite potential limitations, the proposals in play all exceed mere technosolutionism, and include complex accounts of social relations, both within proximate research institutions and across a wider global academy. It is therefore essential to track how machine learning researchers develop their own internal theories of change, how they increasingly understand social, infrastructural, methodological, and planetary outcomes to be mutually entangled, and how they seek to expand or limit moral and material imaginaries that might better respond to such a conjuncture. This knowledge can provide further points of leverage. For one, and breaking from previous periods of epistemic virtue formation that were predicated on the formation of new scientific selves (Daston & Galison, 2007), the texts analysed in this article suggest that this community is taking the group, rather than the individual, as the central agent in this political story. External advocates, seeking to engage researchers, might do well to do the same.

Finally, these discussions point to the increased importance of method in struggles over the internal composition of the research community as well as the wider worlds it helps enact. Methods, as John Law (2004) argues, are constitutive. They help make the worlds they describe. Global warming ups this ante. How machine learning knowledge is valued and sought – especially when formalised in shared methods and standards of inquiry – has at least three kinds of constitutive impacts: it shapes both the technical futures it brings into being, the social and demographic character of research communities, and the ways that heavy-compute inquiries themselves act on the energy and climate systems, warming the planet bit by digital bit. A high-carbon future for these fields is likely to be one with worsening distributional injustices in both the differential global opportunities to create new knowledge and in the negative impacts of climate change. Lower-carbon research norms and shifted epistemic virtues, conversely, could provide social, environmental, and creative dividends. A coalitional politics to reform research methods for both climate and equity reasons lies ahead (Pasek, 2026). Perhaps the task for science and technology scholars today lies in determining how best to enter into these budding coalitions and to help them overcome the current impasse.

Acknowledgements

This work was supported by the Canada Research Chairs Program and Internet Society Foundation (grant number 950-233016) and a fellowship at the Centre for Advanced Internet Studies. Our thanks to our reviewers and the special issue editors for their very helpful suggestions.

References

Abbing RR (2021) 'This is a solar-powered website, which means it sometimes goes offline': a design inquiry into degrowth and ICT. *LIMITS: 7th Annual Workshop on Computing Within Limits,* 14 June 2021. https://doi.org/10.21428/bf6fb269.e78d19f6.

Alaimo S (2012) Sustainable This, Sustainable That: New Materialisms, Posthumanism, and Unknown Futures. *PMLA* 127(3): 558–564.

Alvarado R (2023) AI as an Epistemic Technology. *Science and Engineering Ethics* 29:32.

Anderson K (2013) Hypocrites in the air: should climate change academics lead by example? *KevinAnderson.info*, 12 April. Available at: http://kevinanderson.info/blog/hypocrites-in-the-air-should-climate-change-academics-lead-by-example/ (Accessed 14 December 2017).

Archer M (2024) *Unsustainable: Measurement, Reporting, and the Limits of Corporate Sustainability*. New York: New York University Press.

Baker M (2016) 1,500 Scientists Lift the Lid on Reproducibility. *Nature* 533(7604): 452–454.

Bashir N, Donti P, Cuff J, et al. (2024) The Climate and Sustainability Implications of Generative AI. *MIT: An MIT Exploration of Generative AI*. 27 March 2024. Available at: https://doi.org/10.21428/e4baedd9.9070dfe7 (Accessed 16 September 2024).

Bendell J (2016) Carry on flying: why activists should take to the skies. *Open Democracy,* 30 May. Available at: https://www.opendemocracy.net/en/transformation/carry-on-flying-why-activists-should-take-to-skies/ (Accessed 29 May 2020).

Bender EM, Gebru T, McMillan-Major A, et al. (2021) On the dangers of stochastic parrots: can language models be too big? . *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 1 March 2021: 610–623.

Berné O, Agier L, Hardy A, et al. (2022) The Carbon Footprint of Scientific Visibility. *Environmental Research Letters* 17(12): 124008.

Bian J, Wang L, Ren S, et al. (2024) CAFE: Carbon-aware federated learning in geographically distributed data centers. *e-Energy '24: Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems,* June 4-7, 2024: 347–360.

Brain T, Nathanson A and Piantella B (2022) Solar protocol: exploring energy-centered design. *LIMITS: 18th Annual Workshop on Computing Within Limits,* Online, https://computingwithinlimits.org/2022/papers/limits22-final-Brain.pdf.

Cellard L, Parker C and Haines F (2025) Beyond AI as an Environmental Pharmakon: Principles for Reopening the Problem-Space of Machine Learning's Carbon Footprint. *Environment and Planning E: Nature and Space* 8(3): 1020–1045.

Chakrabarty D (2009) The Climate of History: Four Theses. *Critical Inquiry* 35(2): 197–222.

Clark N and Szerszynski B (2021) *Planetary Social Thought: The Anthropocene Challenge to the Social Sciences*. Cambridge & Medford: Polity.

Cooper ZGT (2023) That Which Escapes: Thinking Through Heat in Proof-of-Work Systems. In: Pasek A, Lin CK, Cooper ZGT, and Kinder JB *Digital Energetics*. Minneapolis and Lüneburg: University of Minnesota Press / meson press, pp. 43–70.

Crawford K (2021) *The Atlas of AI*. New Haven: Yale University Press.

Darling-Wolf H and Patitsas E (2024) "Not my Priority:" Ethics and the Boundaries of Computer Science Identities in Undergraduate CS Education. *Proceedings of the ACM on Human-Computer Interaction* 8(CSCW1): 174:1–174:28.

Daston L and Galison P (2007) *Objectivity*. New York: Zone Books.

de Oliveira Andreotti V, Stein S, Pashby K, et al. (2016) Social Cartographies as Performative Devices in Research on Higher Education. *Higher Education Research & Development* 35(1): 84–99.

de Winkel J, Kortbeek V, Hester J, et al. (2020) Battery-Free Game Boy. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4(3): 111:1–111:34.

Desai A, Abdelhamid M and Padalkar NR (2025) What is reproducibility in Artificial Intelligence and machine learning research? *AI Magazine* 46(2): e70004.

Edwards PN (2010) *A Vast Machine: Computer Models, Climate Data, and the Politics of Global Warming*. Cambridge: MIT Press.

Ensmenger N (2013) Computation, Materiality, and the Global Environment. *IEEE Annals of the History of Computing* 35(3): 78–80.

Espinoza MI and Aronczyk M (2021) Big Data for Climate Action or Climate Action for Big Data? *Big Data & Society* 8(1): 2053951720982032.

Galaz V (2025) *Dark Machines: How Artificial Intelligence, Digitalization and Automation Is Changing Our Living Planet*. New York: Routledge.

Goldstein J and Nost E (eds) (2022) *The Nature of Data: Infrastructures, Environments, Politics*. Lincoln: University of Nebraska Press.

Golumbia D (2009) *The Cultural Logic of Computation*. Cambridge: Harvard University Press.

Harding SG (2015) *Objectivity and Diversity: Another Logic of Scientific Research*. Chicago: University of Chicago Press.

Hardy A (2024) Decarbonizing Research Laboratories? The Tensions Associated with the Commensurability of Carbon and How it Opens Up the Boundaries of Responsibility Attribution. *Review of Agricultural, Food and Environmental Studies* 105(1): 179–198.

Henderson P, Hu J, Romoff J, et al. (2020) Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning. *Journal of Machine Learning Research* 21: 1–43.

Hogan M (2015) Data Flows and Water Woes: The Utah Data Center. *Big Data & Society* 2(2): 2053951715592429.

Hogan M (2021) The Data Center Industrial Complex. In: Jue M and Ruiz R (eds) *Saturation: An Elemental Politics*. Durham: Duke University Press, pp. 283–305.

Hu T-H (2015) *A Prehistory of the Cloud*. Cambridge: MIT Press.

Hughes TP (2012) The evolution of large technological systems. In Bijker WE, Hughes TP, and Pinch T (eds) *The Social Construction of Technological Systems.* Cambridge: The MIT Press, pp. 45–76.

Jensen CB and Morita A (2015) Infrastructures as Ontological Experiments. *Engaging Science, Technology, and Society* 1(0): 81–87.

Joshi K (2026) The *AI Climate Hoax: Behind the Curtain of How Big Tech Greenwashes Impacts.* AI for Climate Claims Report, Friends of the Earth International, February 2026. Available at: https://foe.org/wp-content/uploads/2026/02/AI-for-climate-claims-Report_FEB-2026_FINAL.pdf (Accessed 23 February 2026).

Kaack LH, Donti PL, Strubell E, et al. (2022) Aligning Artificial Intelligence with Climate Change Mitigation. *Nature Climate Change* 12(6): 518–527.

Karlsson S, Srebotnjak T and Gonzales P (2007) Understanding the North–South Knowledge Divide and its Implications for Policy: A Quantitative Analysis of the Generation of Scientific Knowledge in the Environmental Sciences. *Environmental Science & Policy* 10(7): 668–684.

Katz-Rosene RM and Pasek A (2024) Spiral-scaling Climate Action: Lessons from and for the Academic Flying Less Movement. *Environmental Politics* 33(2): 259–280.

Kirdani-Ryan M, Ko AJ and Borisova EA (2024) "Taught to be Automata": Examining the Departmental Role in Shaping Initial Career Choices of Computing Students. *Computer Science Education* 34(1): 87–113.

Knorr Cetina, K 1999. *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge: Harvard University Press.

Knox H (2014) Footprints in the City: Models, Materiality, and the Cultural Politics of Climate Change. *Anthropological Quarterly* 87(2): 405–429.

Knox H (2020) *Thinking Like a Climate: Governing a City in Times of Environmental Change*. Durham: Duke University Press.

Koningstein R (2021) We now do more computing where there's cleaner energy. *The Keyword*, 18 May. Available at: https://blog.google/outreach-initiatives/sustainability/carbon-aware-computing-location/ (Accessed 30 September 2024).

Kumar A, Naqvi B and Wolff A (2023) Exploring the Energy Informatics and Energy Citizenship Domains: A Systematic Literature Review. *Energy Informatics* 6(1): 13.

Lacey S, Golin C, and Shah J (2026) Who Controls Power in the AI Era? *Open circuit*, 9 January. [Podcast]. Available at: https://www.latitudemedia.com/news/open-circuit-who-controls-power-in-the-ai-era/ (Accessed 10 February 2026).

Lannelongue L and Inouye M (2023) Carbon Footprint Estimation for Computational Research. *Nature Reviews Methods Primers* 3(1): 1–2.

Lannelongue L, Aronson H-EG, Bateman A, et al. (2023) GREENER Principles for Environmentally Sustainable Computational Science. *Nature Computational Science* 3(6): 514–521.

Latour B (2004) Why Has Critique Run Out of Steam? From Matters of Fact to Matters of Concern. *Critical Inquiry* 30(2): 225–248.

Law J (2004) *After Method: Mess in Social Science Research*. London: Routledge.

Lehdonvirta V, Wú B and Hawkins Z (2024) Compute north vs. compute south: The uneven possibilities of compute-based AI governance around the globe. In: *Proceedings of the 7th AAAI/ACM Conference on AI, Ethics, and Society*, 21–23 October 2024: 828–838. Washington: The AAAI Press.

LeMenager S (2014) *Living Oil: Petroleum Culture in the American Century*. New York: Oxford University Press.

Li P, Yang J, Wierman A, et al. (2024) Towards environmentally equitable AI via geographical load balancing. In: *The 15th ACM International Conference on Future and Sustainable Energy Systems*, 4 June 2024: 291–307. New York: Association for Computing Machinery.

Li P, Yang J, Islam MA, et al. (2025). Making AI Less 'Thirsty': Uncovering and Addressing the Secret Water Footprint of AI Models. *Communications of the ACM* 68, 7(July 2025): 54–61.

Loeb Z (2021) The Lamp and the Lighthouse: Joseph Weizenbaum, Contextualizing the Critic. *Interdisciplinary Science Reviews* 46(1–2): 19–35.

Lohr S (2024) Will A.I. ruin the planet or save the planet? The New York Times, 26 August. Available at: https://www.nytimes.com/2024/08/26/climate/ai-planet-climate-change.html (Accessed 23 September 2024).

Luccioni AS, Viguier S and Ligozat A-L (2023) Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model. *The Journal of Machine Learning Research* 24(1): 11990–12004.

Luers A, Koomey J, Masanet E, et al. (2024) Will AI Accelerate or Delay the Race to Net-Zero Emissions? *Nature* 628(8009): 718–720.

Lukacz PM (2022) Data Capitalism, Microsoft's Planetary Computer, and the Biodiversity Informatics Community. In: Smits M (ed) *Information for a Better World: Shaping the Global Future.* Cham: Springer, pp. 355–369.

Malazita JW and Resetar K (2019) Infrastructures of Abstraction: How Computer Science Education Produces Anti-Political Subjects. *Digital Creativity* 30(4): 300–312.

Marshall SL and Grooms AA (2022) Industry's Push for Computer Science Education: Is Computer Science Really for All? *Policy Futures in Education*. Epub ahead of print, 21 January 2022. DOI: 10.1177/14782103211045601.

Mehra V and Hasegawa R (2023) Supporting power grids with demand response at Google data centers. *Google Cloud Blog,* 3 October. Available at: https://cloud.google.com/blog/products/infrastructure/using-demand-response-to-reduce-data-center-power-consumption (Accessed 30 September 2024).

Michael J, Holtzman A, Parrish A, et al. (2023) What do NLP researchers believe? Results of the NLP community metasurvey. In Rogers A, Boyd-Graber J, and Okazaki N (eds) *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics,* 9–14 July(1): 16334–16368. Kerrville: Association for Computational Linguistics.

Mollen A, Jansen F, Kannengießer S, et al. (2026) *AI Infrastructures and Sustainability: Expanding Perspectives on Automation, Communication and Media*. Cham: Springer Nature.

Moser AK (2015) Thinking Green, Buying Green? Drivers of Pro-Environmental Purchasing Behavior. *Journal of Consumer Marketing* 32(3): 167–175.

Nardi B, Tomlinson B, Patterson DJ, et al. (2018) Computing Within Limits. *Communications of the ACM* 61(10): 86–93.

Nazaré T, Gadelha J, Nepomuceno E, et al. (2023) Green Computing for Energy Transition: A Survey. *IEEE Latin America Transactions* 21(9): 937–948.

Nilsson EM, Jensen RH, Hansen A-MS, et al. (2024) What do we teach when teaching more-than-human perspectives in computing and technology design education? An emerging pedagogical framework. *LIMITS: 10th Annual Workshop on Computing Within Limits*, Online, https://computingwithinlimits.org/2024/papers/limits2024-nilsson-pedagogical-framework.pdf.

Parrish A (2022) Solar powered dawn poems: Progress report. *Allison Posts,* 21 March. Available at: https://posts.decontextualize.com/solar-powered-dawn-poems-progress-report/ (Accessed 12 February 2026).

Pasek A (2019) Managing Carbon and Data Flows: Fungible Forms of Mediation in the Cloud. *Culture Machine* 18(1). https://culturemachine.net/wp-content/uploads/2019/04/PASEK.pdf.

Pasek A (2023) From Atoms to Electrons: An Energy History and Future of Computing. In: Pasek A, Lin CK, Cooper ZGT, and Kinder JB *Digital Energetics*. Minneapolis, Lüneburg: University of Minnesota Press / meson press, pp.15–42.

Pasek A (ed)(2026) *Low Carbon Research Methods: Making Equity and Epistemological Gains Through Decarbonising Academic Work*. London: Goldsmiths Press.

Pasek A, Vaughan H and Starosielski N (2023) The World Wide Web of Carbon: Toward a Relational Footprinting of Information and Communications Technology's Climate Impacts. *Big Data & Society* 10(1): 20539517231158994.

Patterson D, Gonzalez J, Le Q, et al. (2021) Carbon Emissions and Large Neural Network Training. *ArXiv* preprint: 2104.10350v3, 21 April. DOI: 10.48550/arXiv.2104.10350 (Accessed 6 November 2023).

Patterson D, Gonzalez J, Hölzle U, et al. (2022) The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink. *Computer* 55(7): 18–28.

Peters JD (2015) *The Marvelous Clouds: Toward a Philosophy of Elemental Media*. Chicago: University of Chicago Press.

Piantella B, Nathanson A, Brain T, et al. (2020) Solar-powered server: designing for a more energy positive internet. In: *2020 CHI Conference on Human Factors in Computing Systems*, 25 April: 1–4. New York: Association for Computing Machinery.

Piantella B, Nathanson A, Brain T, et al. (2023) Listening for the 'I' in Infrastructure (and Community) Repair. *Journal of Environmental Media* 4(1): 101–106.

Porter TM (2020)[1995] *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton: Princeton University Press.

Puig de la Bellacasa M (2017) *Matters of Care: Speculative Ethics in More than Human Worlds*. Minneapolis and London: University of Minnesota Press.

Qin Y, Havulinna AS, Liu Y, et al. (2022) Combined Effects of Host Genetics and Diet on Human Gut Microbiota and Incident Disease in a Single Population Cohort. *Nature Genetics* 54(2): 134–142.

Rafat K, Islam S, Mahfug AA, et al. (2023) Mitigating Carbon Footprint for Knowledge Distillation Based Deep Learning Model Compression. *PLOS ONE* 18(5): e0285668.

Ren S, He Y and Xu F (2012) Provably-efficient job scheduling for energy and fairness in geographically distributed data centers. In: *32nd International Conference on Distributed Computing Systems*, 18-21 June 2012: 22–31. Washington: IEEE Computer Society.

The Royal Society (2020) *Digital Technology and the Planet: Harnessing Computing to Achieve Net Zero*. December. London, UK: The Royal Society. Available at: https://royalsociety.org/-/media/policy/projects/digital-technology-and-the-planet/digital-technology-and-the-planet-report.pdf (Accessed 26 September 2023).

Santos R dos, Kaczmarek M, Shankar S, et al. (2021) Who are we listening to? The inclusion of other-than-human participants in design. *LIMITS: 7th Annual Workshop on Computing Within Limits*, Online, https://doi.org/10.21428/bf6fb269.09f36751.

Saul J, Nicoletti L, Rai S, et al. (2024) AI is wreaking havoc on global power systems. Bloomberg, 21 June. Available at: https://www.bloomberg.com/graphics/2024-ai-data-centers-power-grids/ (Accessed 16 September 2024).

Schwartz R, Dodge J, Smith NA, et al. (2019) Green AI. *Communications of the ACM* 63(12): 54–63.

Sefyrin J and Velkova J (2026) The Good Infrastructure: Digital Futures, Values and Friction. In: Mollen A, Jansen F, Kannengießer S, et al. (eds) *AI Infrastructures and Sustainability: Expanding Perspectives on Automation, Communication and Media*. Cham: Springer Nature, pp. 341–362.

Shahjahan RA and Niloy N (2026) The Temporal Politics of Decarbonizing Academic Work: Mobilizing Decolonial and Global South Perspectives. In: Pasek A (ed) *Low-Carbon Research Methods: Making Equity and Epistemological Gains through Decarbonising Academic Work*. London: Goldsmiths Press, pp.179–196.

Shapin S and Schaffer S (1985) *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life*. Princeton: Princeton University Press.

Sheller M (2022) The End of Flying: Coronavirus Confinement, Academic (Im)mobilities and Me. In: Bjørkdahl K and Franco Duharte AS (eds) *Academic Flying and the Means of Communication*. Singapore: Springer, pp. 53–77.

Sigalos M and Smith J (2023) Texas paid bitcoin miner Riot $31.7 million to shut down during heat wave in August. CNBC, 6 September. Available at: https://www.cnbc.com/2023/09/06/texas-paid-bitcoin-miner-riot-31point7-million-to-shut-down-in-august.html (Accessed 30 September 2024).

Silva GD (2021) *Performance and energy efficiency analysis of machine learning algorithms towards green AI: A case study of decision tree algorithms*. Instituto Brasileiro de Informação em Ciência e Tecnologia, Brasília, Brazil. Available at: https://bdtd.ibict.br/vufind/Record/LNCC_eb4fcb879712e1aaa2f1a26fdd9bc563 (Accessed 16 September 2024).

Simon, M. (2022). Maybe green energy needs 'information batteries' too. *Wired,* 31 January. Available at: https://www.wired.com/story/maybe-green-energy-needs-information-batteries-too/ (Accessed 3 October 2024).

Soden R, Ratto M, Verhoeven GA, et al. (2022) Photovoltaic imagination: solar strategies for community integrated research and graduate training. In: *LIMITS: 8th Annual Workshop on Computing Within Limits,* Online, https://doi.org/10.21428/bf6fb269.7b608a0d.

Sørensen E, Çelik L and Laser S (2026) Planetary knowledge: Shifting from social to planetary studies of scientific knowledge. *Science & Technology Studies* 39(3): XXX.

Strathern M (1997) 'Improving Ratings': Audit in the British University System. *European Review* 5(3): 305–321.

Strubell E, Ganesh A and McCallum A (2019) Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, July 2019: 3645–3650. Kerrville: Association for Computational Linguistics.

Switzer J and Raghavan B (2021) Information Batteries. *ACM SIG: Energy Energy Informatics Review* 1(1): 1–11.

Switzer J, McGuinness R, Pannuto P, et al. (2021) TerraWatt: Sustaining sustainable computing of containers in containers. *ArXiv* preprint:2102.06614, 12 February. DOI: 10.48550/arXiv.2102.06614 (Accessed 30 September 2024).

Tanenbaum TJ, Pufal M and Tanenbaum K (2016) The limits of our imagination: Design fiction as a strategy for engaging with dystopian futures. *LIMITS: 2nd Annual Workshop on Computing Within Limits*, 8 June 2016. https://dl.acm.org/doi/10.1145/2926676.2926687. New York: Association for Computing Machinery.

Thompson NC, Greenewald K, Lee K, et al. (2021) Deep learning's diminishing returns. *IEEE Spectrum*, 21 September. Available at: https://spectrum.ieee.org/deep-learning-computational-cost (Accessed 28 August 2024).

The Turing Way (2022) The Environmental Impact of Open Research. *GitHub*, issue #2803. Available at: https://github.com/the-turing-way/the-turing-way/issues/2803 (Accessed 26 September 2023).

Uygun Tunç, D and Pritchard, D (2026) Individual and Collective Epistemic Virtue in Science. *Synthese* 207: 21.

Van Dongen J and Paul H (eds) (2017) *Epistemic Virtues in the Sciences and the Humanities*. Cham: Springer International Publishing.

Vaughan H, Pasek A, Silcox NR, et al. (2023) ICT Environmentalism and the Sustainability Game. *Journal of Language and Politics* 22(5): 640–660 .

Watts L and Velkova J (2025) From Flatlands to Liveable Futures: Unflattening Carbon Metrics in the Energy and Data Industries. *Big Data & Society* 12(4): 20539517251396074.

Wu C-J, Raghavendra R, Gupta U, et al. (2022) Sustainable AI: Environmental implications, challenges and opportunities. *ArXiv* preprint: 2111.00364, 9 January. DOI: 10.48550/arXiv.2111.00364.

Zero Cool (2019) *Oil is the new data.* Logic Magazine, 7 December. Available at: https://logicmag.io/nature/oil-is-the-new-data/ (Accessed 25 October 2023).

Zhang S, Roller S, Goyal N, et al. (2022) OPT: Open pre-trained transformer language models. *ArXiv* preprint: 2205.0106, 21 June. DOI: 10.48550/arXiv.2205.01068.

Notes

1 In this article (and elsewhere) we somewhat stubbornly distance ourselves from the AI moniker in an effort to more accurately describe and demystify the computational work and infrastructures of this field. Descriptions of 'intelligence' work against this goal. For more on the long history of such rhetorical lines in the sand, see Loeb (2021).

2 This figure (somewhat unusually) includes water use in electricity generation, not just use in data centres specifically. Some have argued that this leads to unfair comparisons.

3 We thank Dawn Nafus for this point.

4 This is to say, researchers can potentially 'buy' results (and thus a publication) by simply purchasing additional compute and/or data to feed into a model to produce an advance, rather than exploring novel approaches or testing hypotheses.

5 Cryptocurrency makes for a particularly easy villain even though, from an infrastructural perspective, Bitcoin mining and neural net training are not so dissimilar. See Cooper (2023) for a further discussion on this point.

6 Some researchers, aided by reporting metrics from their cloud providers, may also choose to purchase further carbon offsets. See Qin et al. (2022) for an example.

7 This is, of course, a fairly large presumption about the ease of research collaborations across North/South divides, and the harmful effects of power asymmetries that arise therein.

8 It is difficult to assess the total size of this group using our social cartography approach. However, it's fair to assume that degrowth proponents are considerably smaller than those of the more mainstream approaches to sustainability previously covered in the paper. For example, the 2024 LIMITS conference featured only 11 papers. ACM FAccT (The Association for Computing Machinery Conference on Fairness, Accountability, and Transparency), conversely, is a conference associated more directly with a reformist, ethical AI framework, inclusive of environmental efficiency measures. Its 2024 conference included over 175 papers. Though they do not all concern environmental questions, they represent a larger audience for that kind of conversation.