

"As an individual, I suppose you can't really do much": Environmental Sustainability Perceptions of Machine Learning Practitioners

Sinem Görücü
King's College London
London, United Kingdom
sinem.gorucu@kcl.ac.uk

Gabrielle Samuel
King's College London
London, United Kingdom
gabrielle.samuel@kcl.ac.uk

Yuheng Ren
King's College London
London, United Kingdom
yuheng.ren@kcl.ac.uk

Georgia Panagiotidou
King's College London
London, United Kingdom
georgia.panagiotidou@kcl.ac.uk

Abstract

The growing need for computational resources and energy to develop, train, and run Machine Learning models has been gaining increasing attention. Still, while a considerable body of research is focusing on tools and methods to calculate the energy consumption and carbon footprint of ML processes, there is a distinctive lack of research into the perceptions, practices, challenges and behaviors of ML practitioners regarding the environmental sustainability of their work. To address this gap, we conducted a series of semi-structured interviews with ML practitioners from academia and industry to understand how they engage with the topic within their teams. We found that sustainability, even within our relatively motivated interview cohort, was still seen as a 'nice-to-have' with considerations such as speed, market competition, or time management considered as having higher priority. We discuss these findings in-light of prior work on sustainability within household and work settings and highlight the need for individual as well as collective responsibility and agency. This work contributes to the design and development of future ML eco-feedback tools and promotes sustainable awareness and practices within this context.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**; *HCI theory, concepts and models*; • **Social and professional topics**;

Keywords

Environmental Impact of AI, Sustainable HCI, Green AI

ACM Reference Format:

Sinem Görücü, Yuheng Ren, Gabrielle Samuel, and Georgia Panagiotidou. 2025. "As an individual, I suppose you can't really do much": Environmental Sustainability Perceptions of Machine Learning Practitioners. In *The 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*, June 23–26, 2025, Athens, Greece. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3715275.3732088>



This work is licensed under a Creative Commons Attribution 4.0 International License. *FAccT '25, Athens, Greece*

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1482-5/25/06

<https://doi.org/10.1145/3715275.3732088>

1 Introduction

The rapid adoption of machine learning (ML) across industries has brought significant environmental challenges [2, 9, 24, 25, 94]. As ML models increase in scale, complexity and deployment, the resources required to support them—spanning water, minerals, energy, and labor—have grown substantially [2, 21, 24]. Among these, the rising energy demand of large models and the resulting carbon emissions have drawn particular attention [17, 25, 46]. In response, the ML community has been increasingly focused on measuring and mitigating the carbon footprint of ML processes [71, 105]. Existing efforts include a variety of calculation techniques and ML eco-feedback tools such as computational schedulers [5, 102], data visualizations [23, 63, 98], recommendation interfaces [51, 110] and product performance ratings [13]. Despite these advancements, however, the perceptions, observations, experiences of ML practitioners regarding the sustainability of their work have yet to be empirically studied or integrated into sustainable ML initiatives. To address this gap, we conducted interviews with 23 ML practitioners spanning a various industries, roles, application areas, and domains.

Our findings reveal a growing interest in sustainable or Green AI [94] practices, yet many practitioners hold vague or incomplete assumptions about ML's environmental impact. Participants cited various barriers to action, including beliefs that their individual contributions are negligible, disciplinary challenges, and the perceived value of ML outputs. We summarize participants' suggestions for fostering sustainable practices both individually and collectively, and explore two overarching themes that emerged from our interviews: (a) who holds responsibility for environmental sustainability, and (b) who has the power and agency to act. This work provides a nuanced perspective on the discourse around Green AI, grounded in the lived experiences of ML practitioners. It offers practical insights for designing and developing effective eco-feedback tools and promoting sustainability awareness within the ML community.

2 Related Work

2.1 Environmental Impact of Machine Learning

The greenhouse gas (GHG) emissions caused by Information Communication Technology (ICT) account for 2.1%–3.9% of global GHG emissions, having more than doubled over the past decade, and will continue to grow if left without intervention [42]. For instance,

Google reported a 48% rise in GHG emissions over the past five years due to its data centers [46, 76], and training OpenAI's ChatGPT-3 is estimated to have consumed approximately 700,000 liters of clean water for data center cooling [65], along with electricity equivalent to the annual consumption of about 130 US households [17]. Such substantial resource demands have delayed the retirement of coal power plants [49] and intensified conflicts with local communities in regions like the United States and Spain [35, 68, 69], and while some of this energy required can come from renewable or carbon credit-offset resources, such measures remain limited in availability, costly to the environment and may even take away from other potential uses of green energy [4, 9]. This impact is not only significant but also challenging to measure [59, 71], with terms such as “cloud” computing and “artificial” intelligence making this process even more intangible and abstract [24, 25, 52, 75].

To mitigate the issue, various tools are currently being developed that calculate model energy consumption and carbon emissions [5, 61, 63, 98], also referred to as ML eco-feedback tools [50]. Such tools typically work by taking input data related to ML work (e.g., GPU, memory, runtime) through automatic monitoring or manual input by users, and through calculations based on specific formulas, they provide an estimate of energy consumption and carbon emissions as output in various forms. Although these tools vary in how they retrieve data (e.g., through automatic monitoring or manual input) and in how they present results, they are essentially eco-feedback technologies designed to promote eco-friendly behavioral change and reduce environmental impact [44] of ML work.

2.2 Sustainability Awareness in Home and Workplace Settings

Given the significant proportion of emissions attributed to individual or household activities [6, 82] numerous studies have focused on understanding and influencing sustainability awareness and practices within the household context, including sustainable awareness evaluation [16, 60, 80], interventions aimed at raising awareness [48, 55, 90, 99], and challenges associated with implementing such interventions [7, 91, 96]. Domestic energy use has great potential for reduction [27] among others because it is closely linked to personal awareness and behavior [47]. Yet, energy consumption in the workplace also accounts for a considerable part of overall energy usage. For instance, in the US, work settings account for 36% of electrical energy consumption [112]. And across Europe, electricity consumption in non-residential buildings has increased by 74% over the last 20 years [34]. Nonetheless, in contrast to residential energy, research focused on energy consumption and sustainability awareness in the workplace remains relatively limited [86]. This may partly stem from additional barriers in organizational settings, such as the lack of direct financial responsibility and feedback on personal actions, as well as shared facilities and workspace [18, 33]. Nevertheless, some HCI research has adopted approaches similar to those used in domestic settings to enhance sustainability awareness in the workplace, including gamification [45, 83, 100, 103], visualization [12, 87, 111], internet of things (IoT) sensors [1, 58, 88] as well as hybrid approaches [84]. Still several risks including privacy, misinterpretation, and cost in making workplace energy practices more transparent have been documented in work settings [95] together

with a lack of individual visibility of consumption, complexity of workflows, and technological limitations [57]. To mitigate these, research has suggested more personalized [20] and role-specific feedback [40].

Despite existing studies on workplace sustainability, there remains a lack of research specifically addressing ML work settings, which present unique characteristics and challenges. In order to design appropriate energy interventions, workplace characteristics and contexts need to be well understood [8, 15, 20, 40, 95]. Still, most of this research refers to building or facilities energy usage rather than ICT making the issue here especially pertinent for ML workplaces given the computational demands and distributed infrastructure required. While existing research has explored ML practitioners' practices and perceptions in contexts such as ML fairness [29] and responsible AI [73], the environmental impact of ML and practitioners' sustainability awareness remain largely unexamined.

3 Methodology

We conducted semi-structured interviews with 23 ML practitioners spanning various backgrounds (see Table 1), and coded the interviews through a collective and iterative thematic analysis approach [14].

Participant Recruitment. Our recruitment criteria for the interview study involved professionals working on ML pipelines regardless of the country of residence, age, gender, educational background, or specific role. We adopted a broad definition of ‘ML practitioners’ to refer to developers and researchers that apply data-driven artificial intelligence techniques to solve practical, prediction-based problems, regardless of the dataset size, data type, or application domain. We created a pre-interview registration form, with two control questions to check how many years of experience participants had working with ML and a brief explanation of what their role involved. Participants who did not provide sufficient information about their role definition were excluded from the study. This was to ensure that the interviews would be realized with the appropriate target audience and avoid spam.

We used a snowball sampling method for recruitment. Initially, we created a recruitment poster which we shared across relevant professional platforms such as social media (e.g., LinkedIn, X, Reddit), online mailing-list communities for ML professionals and university research subject newsletters. At the same time, we also sent e-mail invitations to eligible contacts from our professional networks, and during interviews, we asked the interviewees to further share the recruitment call within their own networks. After reviewing the responses from the pre-interview registration form, suitable participants were invited to a 45-minute online video-call interview at a time of their convenience. Before the interview, participants were sent a digital information sheet and consent form. This study was approved by our university's ethics board.

We recruited 23 ML practitioners as participants for our interview study. These participants ranged in age from 18 to 40 years old and held various roles, including ML engineer, data scientist, algorithm engineer, manager when coming from the industry, and lecturer/assistant professor, undergraduate/PhD student, Postdoc researcher when coming from academia, with their experiences

ranging from 1 year to more than 10 years. A small number of participants were excluded due to limited or irrelevant experience with ML work according to the responses from the pre-interview form.

Interview Design and Process. The interview plan outlined key topics and potential questions for exploration, comprising of 27 open-ended questions organized into four sections: 1) background information, 2) work practices, 3) sustainable factors in work practices, and 4) sustainability awareness, policies, and industry trends. An indicative interview plan can be found in the supplementary material¹. The interviews were not strictly bound by the interview plan instead, the plan served as a guiding framework for researchers to facilitate in-depth exploration of relevant topics. For instance, in a few cases, the plan was modified to include more detailed questions on the participants' backgrounds and practices. We conducted three online video pilot interviews to evaluate the interview plan, recruiting participants who matched the inclusion criteria. Pilot interviews were excluded from the analysis and were only used to improve the interview guide by rephrasing, adding, or removing certain questions to ensure clarity, flow, and inclusivity.

We conducted the final semi-structured interviews with the 23 participants on MS Teams. We recorded and transcribed all interviews via the platform's embedded features except one, where the participant requested not to have a recording and we took active notes during the interview instead. Interviews averaged 45 minutes, with some extending up to 80 minutes as participants shared additional insights. Each participant received a £15 (GBP) voucher as compensation.

Following our interview plan, we started the interview by asking participants to describe their current role and responsibilities, including how they interact with ML, their work structure and the typical work process. In particular, we paid attention to understanding what the criteria were for them when they were trying to make various decisions during the work and examined if and how sustainability played a role. Then, we probed participants broadly on "sustainability", beginning with the question "What does the term 'sustainability' mean to you (in the context of ML/AI)?" to understand how they interpreted the term "sustainability". Only then did the questions zoom into environmental sustainability of ML, including the sustainability factors in their work practices, if and how their perceptions of sustainability were applied, and if they had experience using any carbon tracking tools. Finally, we queried their opinions regarding sustainability in the wider context of ML, as well as if and how external factors such as policy can affect the perception and application of sustainability. The interviews took place from September to December 2024.

Analysis. We conducted an inductive thematic analysis of the interview data, following Braun and Clarke [14]. We began by revisiting all 23 transcripts to gain familiarity with the content. Then, we systematically coded the interviews, starting with all researchers individually coding the same 3 transcripts to help alignment. The authors held weekly meetings to discuss and reflect on the codes and emerging themes. Once initial themes were generated, each

researcher independently coded a different set of 3 transcripts each time, while still collaboratively reviewing and clustering codes to merge commonalities. After several iterations on the themes and their relationship, we generated the final set of themes. We used the web-based thematic analysis tool Atlas.ti for coding and clustering.

4 Findings

Our thematic analysis produced three core themes. While these are outlined here separately, they have clear intersecting aspects and they form part of the larger narrative the participants brought into our discussions.

4.1 Current Perception and Understanding of Sustainability

4.1.1 Interest in Sustainability of ML. While most participants were somewhere in-between, some participants expressed little interest in the environmental sustainability of ML (P2I, P6A), and others demonstrated deep concern, even reporting experiences of eco-anxiety and carbon guilt related to climate change (P5I, P15I, P18A). These participants reflected on the environmental footprint of their work: "It's something I feel guilty about, and I do try to keep it in mind" (P18A).

There has been a documented increase in interest in the environmental impact of ML as also expressed by participants due to recent academic and popular media articles or even discussions on online platforms (P3A, P5I, P8I, P10A). P10A for instance, vividly recalled their perceptions and understanding during past projects with large language models (LLMs), explaining how their behavior would have been different with their current awareness about the extensive resource consumption of LLMs, stating: "Because it's like very apparent nowadays just how much resource those are using. So if I were to go back, I guess it's something I'd want to consider, but it was something I didn't really understand, I guess at the time" (P10A).

More than a quarter of the participants also critically discussed the lack or inadequacy of action in the field, indicating the emergent interest in the topic does not echo in their work. "I think sustainability was never really a main concern in any of the environments [academia or industry] that I've been in, and I feel like that is the case for a lot of the institutions that are doing ML research. This is unfortunate though, like I don't endorse it. I don't like it." (P5I). Only a small number of participants elaborated on the existence of ML sustainability discussions or measures at their work and those were most often received with surprise. P16IA for instance recalled a recent event they attended on the use of AI in the armed forces: "it was the first time in a very long time, one of the delegates actually asked around sustainability of AI. That is a rare, rare comment to have had, and I've not had that for years before then." (P16IA).

However, it became apparent that even the participants who cared deeply about the issue perceived environmental sustainability as a moral and ethical issue to consider voluntarily, rather than a necessity, describing that they just "try to be as responsible as they can" (P20A). This sentiment was elaborated directly and indirectly throughout the interviews: "It's a good ethos to have and it's kind of like, you know, a morality thing. It's like, 'at least I'm not completely abandoning all hope and care'" (P18A).

¹Full interview plan: https://osf.io/3sbg4/?view_only=1f6f807c215d4e56bf7eaa5fc113b0c3

Participant ID	Domain	Job title	Years of experience	Gender	Age	Place of Residence
P1A	Academia	Assistant Professor	10+	Female	31–40	United Kingdom
P2I	Industry	Algorithm Engineer	3	Male	18–30	China
P3A	Academia	PhD Student	4	Male	18–30	United Kingdom
P4A	Academia	PhD Student	1	Male	18–30	Belgium
P5I	Industry	ML Engineer	10	Woman	18–30	Turkey
P6A	Academia	PhD student	3	Male	31–40	United Kingdom
P7I	Industry	ML Engineer	1	Male	18–30	China
P8I	Industry	Senior Manager	10+	Female	31–40	United Kingdom
P9A	Academia	PhD Student	4	Female	18–30	Australia
P10A	Academia	PhD Student	7	Woman	18–30	United Kingdom
P11IA	Academia & Industry	AI Engineer	3	Male	18–30	Turkey
P12I	Industry	Advanced Analytics & AI Strategy Lead	4	Male	31–40	United Kingdom
P13A	Academia	Undergraduate Student	2	Male	18–30	United Kingdom
P14A	Academia	PhD Student	7	Woman	31–40	Turkey
P15I	Industry	AI Solutions Architect	10+	Male	18–30	United Kingdom
P16IA	Academia & Industry	Assistant Professor	10	Male	31–40	United Kingdom
P17I	Industry	Associate	Prefer not to say	Prefer not to say	Prefer not to say	Prefer not to say
P18A	Academia	PhD Student	4	Male	18–30	United Kingdom
P19I	Industry	Senior Data Scientist	4	Female	18–30	United Kingdom
P20A	Academia	Lecturer	8	Male	31–40	United Kingdom
P21A	Academia	Postdoc Researcher	4	Male	18–30	United Kingdom
P22I	Industry	Responsible AI Lead	10+	Female	31–40	Spain
P23A	Academia	Research Fellow	5	Male	31–40	Ireland

Table 1: Participants' backgrounds as collected from the pre-interview survey. Gender is here as self-described by the participants. For each participant number, a letter ("A", "I", or "IA") is appended to indicate whether the participant's background is from academia, industry, or both.

4.1.2 Assumptions on ML's footprint. While the majority of the participants expressed concern about the environmental footprint of machine learning, their understanding on the level of the energy usage and related environmental harm varied. Even the participants who took measures to lower their computational power suggested that ML's footprint is not as much of a major contributor to carbon emissions as it appeared, making rough assumptions of what percentage ML operations contribute to the global or national energy consumption (P21A, P23A) "I think it will be some effect of machine learning on environmental sustainability. But still, when you look at other subjects it will not be too big. OK. It is important, I agree. But it affects maybe two per cent, three per cent? I don't think it's more." (P21A). P2I also discussed the highly efficient ML industry to support this position: "I think, overall, the efficiency of using resources is pretty high in the ML industry. And, for now, the impact, the emission, is pretty small" (P2I). P23A further described sustainability as something to be addressed down the road rather than a priority compared to other sustainability efforts their institution is focused on, "We're doing a lot of this low-hanging fruit stuff at the moment, you know, like we're insulating the walls, we're replacing the windows, we're replacing the boiler systems. So for someone in to be looking at how much the machine is using [...], it's a long way down the pecking order, I suppose." (P23A)

Echoing this stance, several participants also referred to other industries, either ICT or technologies like Bitcoin (P18A) or even traditional industries like fossil fuels for comparison (P2I), stating that the environmental impact of ML is "pretty small compared to traditional industries" (P2I), and these industries are "obviously not quite comparable" (P18A) to what they are doing. P21A also associated the perceived high energy consumption levels with the

fact that ML is an umbrella term involving operations in a variety of industries.

Moreover, multiple participants voiced their doubts about the media's representation of the topic. For instance, during the interview, P5I questioned the reliability of information shared on the footprint of certain LLM models by news outlets while still expressing their will to believe in it. "I just hear that, but as I said, I didn't research it myself. And news... I always... Like [they're] too sensational. So I'm not sure. I was not sure if it was correct or it was a bit bloated, but I would believe it if it were true" (P5I). Similar skepticism on the data published in the media was shared by other participants as well, as P16IA mentioned: "and again, the media misuses it a little bit," demonstrating the need for direct and more transparent sources of energy documentation.

4.2 Challenges and Reasons for Inaction

Throughout the interviews, participants reflected on a variety of reasons that have prevented them from considering sustainability in their work. These included perceptions around individuals' roles and responsibilities, disciplinary challenges, perceptions related to the value and benefit of ML operations, and practical considerations or limitations including the inherent uncertainty of ML development and of energy consumption.

4.2.1 Individuals' role and responsibility. There was a consistent perception of the minimal role and responsibility individuals can have. This sentiment was discussed through different viewpoints such as the belief that their works' contribution to environmental harm is negligible, and skepticism about who holds responsibility.

Participants believed that their contributions were not a major factor compared to contributions of other models, domains or

phases/roles. They particularly emphasized how models and procedures they have been working on have been “fairly lightweight” (P19I), “quite basic” (P3A), or “just on my local machine without big servers” (P10A), and it is “a very small amount of this unsustainable system because there are much more powerful models” (P11IA). Some of these participants also specifically pointed towards LLMs as the energy-intensive model instances that rely on massive datasets, to direct our attention and efforts towards —especially with regard to the training phase, indicating that not all phases of the ML pipeline contribute to the environmental footprint equally.

Some participants referred to this imbalance among phases to explain how their role does not involve the most computationally heavy procedures. For instance, P19I, a senior data scientist working in healthcare, explained how sustainability just recently started to be a consideration for them as they started working on LLMs but the participant’s role did not involve the training of it. As P5I explicitly mentioned, “What I do is I just so minuscule compared to pre-training. [Although] there might be people who do it in the company” (P5I).

Another participant, P14A, who is a PhD student working in computer science, echoed this view and similarly pointed to NLP as a highly consuming area: “You know, some firms working on NLP. If I was a member of these firms, I would be also aware of it. But in my lonely research, the environmental resources I am wasting is not comparable with this.” (P14A). Another PhD student in life sciences (P18A) reflected on ChatGPT —an LLM based chatbot, claiming that “probably the focus as usual needs to be on the people at the top who are using large amounts of energy. You know, I love ChatGPT. It’s been transformative, like I’m super grateful for it. But on the other hand, it’s using massive amounts of energy, maybe things like that need to be targeted first and then, you know, you see where you can go from there.” (P18A).

Importantly, participants also conveyed a lack of individual agency to act on or cause change, both from academia and industry. As P23A —who works at a research center that bridges industry with academia mentioned, “Some of the figures are a bit shocking, I suppose. You know, you wonder where the line is, but as an individual, I suppose you can’t really do much about it.” (P23A). This and other quotes indicated participants’ overarching view that the responsibility lies with tech service providers and big technology corporations. P23A also mentioned that even if not as small as that of a domestic user, their impact would still not be comparable to bigger organizations: “Like when you think it as a nation, we probably burn about 6 gigawatts a day across the country.[...] So we’re [as a research centre] not even 1 per cent of that. We’re not even 0.1 per cent of it. So like, unless you’re using in megawatts, you’re not gonna change the world. You know, you have to go after the people who are burning megawatts” (P23A). P8I, a senior manager in ML industry discussed that without “big tech” companies “consider(ing) any additional sustainable actions”, industries that are using their off-the-shelf tools would not be able to “take action themselves”.

4.2.2 Disciplinary challenges. Another theme that emerged in the interviews was related to the challenges that stemmed from domain priorities, the nature of the field, and the fundamental ways in which work is conducted within it.

“In my opinion, the hard focus on model performance should change because increasing model performance is the key to this unsustainable way” (P11IA). Echoing Section 4.1.1, we repeatedly heard how sustainability was not a priority in participants’ work. As P18A succinctly explained it: “Environmental sustainability is not one of our results, it’s not a metric of performance.” Participants pointed out several possible reasons for this, such as time, labor, and resource limitations (P5I, P14A), lack of external drivers (e.g., client interest or conference procedures) (P1A, P14A), lack of internal sustainability accountability mechanisms (P9A, P23A) and bureaucracy or legacy issues (P12A, P23A).

The nature of the work was described as ‘fast’ and ‘competitive’. P5I, a machine learning engineer at an LLM startup, emphasized the competitive nature of the field: “We’re startup, and we’re moving super fast because everyone in this field are like, every time you see a state-of-the-art model, doing the best on benchmarks and things like that.” (P5I). Meanwhile, academic participants further discussed the time limitations specific to research work, such as the pressure to publish papers (P14A), the need to complete a research project or a PhD within a set timeframe (P9A, P10A, P14A, P18A, P23A), and the competitive environment around finishing and publishing their latest research (P14A). “I need to do my research and if I was to tell my supervisor no, I’m not going to use the HPC because I feel bad for the penguins in the Antarctic, then that wouldn’t go down so well.” (P18A). This standpoint was also shared in industry where participants discussed that even if they were provided with eco-feedback of their process, “in the end of the day, (they) still need to do the things that (they) need to do” (P19I). At the same time, some participants reflected on time limitations referring to optimization of their models for efficiency, only if “you’re a huge company and have people dedicated to working on it” (P5I), while others were suggesting that the inadequacy of allocated resources and time pre-development, such as “user research or customer (sectors) to nail down the end goal” (P19I) was also a limitation.

While sustainability was not a central concern, other metrics, such as performance or accuracy (P8, P6A, P9A, P11IA, P12I, P13A, P18A, P21A) and metrics like data quality (P7I, P11IA, P21A), model fitting (P17I), and fairness (P10A, P12I, P19I, P22I) or interpretability (P18A) received primary attention in the process: “The model accuracy is more on the front than the sustainability.” (P11IA). Although several participants discussed metrics for efficiency and optimization processes, P21A emphasized the importance of maintaining accuracy, stating; “If my accuracy goes down too much, I directly stop the optimization there.” (P21A). Ethical considerations like bias, interpretability, and fairness were also often prioritized, unlike sustainability. This, however, differed among domains as for instance, participants working in healthcare referenced the sensitivity and high-stakes nature of their field: “I would prefer to have a slightly less accurate model, but that, that is not biased against a certain patient subgroup or demographics” (P19I).

This also emphasized that there are limited internal accountability practices in place. P19I described for instance how there is no need for energy justification in their workplace: “As long as I’m able to justify (you know) the use of those resources. Then it’s generally OK [as long as I do not forget] to shut down the cluster during the weekend.” (P19I).

Similar to the lack of internal sustainability accountability practices, there was also a documented lack of external drivers. Participants repeatedly acknowledged that neither the customers they work with nor the industry partners are in demand of sustainability considerations associated with the ML process, but are rather stressing performance. One of them with a background in the mining industry compared: "In production mine site, it's very common. It's needed, it's the must. But in machine learning they don't. They don't have awareness about how resources are used." (P14A). Similarly, P16IA, who is based in both academia and industry, brought up the lack of consideration from their country's major funding bodies: "[environmental sustainability] is not part of any requirement as part of [Name] Research funding." (P16IA). And P1A, an Assistant Professor in Computer Vision, also discussed how the conferences were not adopting sustainability accountability practices mentioning how "...this is not in the agenda, it's not like a priority. And therefore, even if you do hear about these issues, [they are] not part of your responsibility as a researcher." (P1A).

4.2.3 ML's value and benefit. Participants described that another reason for inaction is the perceived value and benefit of ML. Some of the participants touched on the notion of 'doing the best they can' while acknowledging the revolutionary benefits of ML (P20A) and striving to find a balance between benefiting from ML and minimizing the harm (P5I, P15I). "You want heavily performance model or you want highly sustainable model? So I think, for me, I have to find a sweet spot, a balance between these two" (P15). As participants came from a variety of backgrounds and domains using ML for various purposes, their perception of the societal value of their work differed. For instance, participants working in healthcare communicated that healthcare helped 'save lives' and thus differed from some other fields (P18A, P19I, P21A). "Generally, in my area, people care about accuracy, not energy consumption. Yeah, actually, they are not wrong because in medical cases, if your accuracy is less because of your energy consumption (meaning reduction), people will judge you basically." (P21A). This was contrasted with other more theoretical disciplines: "if I'm a mathematician doing some pure math analysis where I'm dealing with these crazy numbers, I would expect people like that [to] have an obligation to be a bit more energy efficient. [...] Whereas if, if you know, a doctor is developing an algorithm that's gonna save more lives, I am a bit more willing for them to use up a bit more energy." (P18A).

4.2.4 Uncertainty of calculation and measurement. A practical limitation highlighted by participants (P11IA, P18A, P19I, P20A, P21A, P23A) was the exploratory and iterative nature of ML work, which makes it challenging to calculate or estimate the environmental impact until the final product is achieved. Multiple participants reflected on their 'simpler-to-complex', iterative development process, explaining how most often they start working with simpler or smaller models and work their way up. P11IA flagged this challenge, explaining: "It really consumes a lot of energy, and most of the time it's just trying to make a good work. It doesn't even make a perfect solution. It's not even a perfect solution, it's just a try. If it's working or not... If it's a good model or not, then we are trying again to improve our model." (P11IA). Another participant quoted, "When you create your ML model in the beginning, your initial ML model, you are creating it randomly", and then continued

further detailing the uncertainties during the process; "So you don't know its efficiency, accuracy and energy consumption.. Then you start to optimize it. Sometimes when it is smaller it's accuracy become higher, so it becomes more efficiency and more accuracy. And sometimes when it is higher, it becomes high accuracy." (P21A).

Participants also expressed that most of the time, they cannot foresee or track the computational load of their operations, expressing a need for feedback—through tools or procedures—to better understand their energy consumption (P9A, P13A, P14A, P15I, P18A, P21A). However, others were skeptical about such feedback, challenging its uncertainty and inaccuracy due to the difficulty of measurement (P16IA, P19I, P20A). P16IA discussed: "It's [a] stab in the dark of guessing of actually the impact, because unless you know the architecture for which your cloud functions are running on at that time, it's very difficult to attribute energy output." (P16IA). P20A further pointed out the uncertainty of the energy sources in the grid, as a reason for uncertainty; "Ideally would like to be using as much as renewable as possible in order to reduce CO₂ emissions. But the issue that arises at renewable energy is uncertain, you don't know how much wind you will have tomorrow in California, for example." (P20A).

Another challenge of uncertainty is related to the fact that it is not easy to predict the inference load of user-facing products. P15I discussed: "I would say, you can easily calculate or compute the emissions for data engineering, and then the model development or model training, but you never know like how customers will use it, right? [...] I know how much energy to develop it, but I'm not sure how many users will use it in a year or two. [...] So that's the problem here" (P15I).

4.2.5 The abstract nature of energy. While not expressed directly, the abstract nature of energy and the difficulty of experiencing it was also a recurring challenge among participants. This abstraction was evidenced by how participants used proxies to describe their energy usage. The most commonly used proxies were computational time and server costs with speed differences across different programming languages being commonly referred to (P11IA, P12I, P15I, P21A): "I know that Python is really slow and slow in this type means it needs more sources to run. And that's why I think Python is a little more environmentally unfriendly" (P11IA). Interestingly, participants also used physical indicators as proxies to understand the energy consumption of their local hardware (P3A, P9A, P11IA, P21A). For instance, P9A mentioned the fan noise and spikes in electricity bills: "So, when I was trying to run something on my local computer, Oh my gosh, it's really crazy with all these fans running, how my electrical bills go [up], it was like crazy. So I [am] definitely understanding all that." (P9A). P21A similarly referred to the heat generated by the device: "[...] and my laptop was burning, really burning, and it was not good" (P21A). When describing such physical experiences, the participants generally used more emotional and emphatic language, demonstrating to us the potential for more direct and visceral understanding of energy consumption.

4.3 Paths to Initiate Change

During the interviews, several ways forward or potential 'solutions' were brought up or speculated even when discussing challenges or

instances of inaction. We present these and their related discussions organized from technical solutions to culture-wide shifts.

4.3.1 Technical solutions. “I’m aware of all that, but I don’t know how to actually reduce all that, to be honest” (P9A). This sense of helplessness was shared overall. Still, some participants were also able to point out simple technical solutions to reduce their environmental footprint. Such solutions included: reusing existing models or data (P6A, P12I, P14A, P15I, P18A), adopting smaller models (P5I, P10A, P15I, P19I, P21A), using compiled programming languages (P11IA, P12I, P15I, P21A), infrastructural or hardware improvements for better performance (P6A, P9A, P12I, P13A, P15I, P17I), continuous efficiency monitoring (P12I), moving to cloud solutions for resource distribution (P12I, P14A, P16IA, P23A), ensuring clarity of the code (P3A, P5I, P17I).

Furthermore, some participants also naturally suggested existing eco-feedback tools as potential solutions even before they were asked by us about them. In particular, they listed several ways that eco-feedback tools could be transformative, from raising self-awareness to informing behavior change, and even further encouraging colleagues or leading to social responsibility. Some expressed that eco-feedback tools could make them more conscious (P19I) or “increase awareness of the consequences of [your] code” (P3A). Eco-feedback tools might also help inform others: “[...] I could show my friends that this is not a very sustainable way to do it” (P11IA). Yet, the relation between awareness and changing practices was somewhat divided among participants. P9A, a PhD student working on NLP, expressed that they could possibly change behavior if they received such feedback mentioned that they would “really love to change how I actually work after I see all these results from these carbon trackers.” (P9A). Whereas, others further mentioned that tools can only raise awareness and are unlikely to cause any behavior change among ML practitioners. “I mean it’s... I think it’s just gonna make me very conscious of all my work, whereas I mean, I don’t know how much I’d actually use it to improve my efficiency. But yeah, it’s just gonna get me more conscious in general.” (P17I).

4.3.2 Individual action. Raising awareness among ML professionals was perceived as the key across most interviews (P3A, P4A, P8I, P9A, P12I, P13A, P14A, P15I, P18A, P21A, P22I). As P22I explained: “Carbon emissions, like I’ve read some articles but [...] to be honest, I have no idea what is the impact of our work in the planet. So awareness is one of the key problems I think.” (P22I). P12I, an AI strategy lead whose duties include promoting sustainable ML practices within the company, also emphasized the importance of “informally sharing and socializing” sustainable awareness and practices, since “most of the things do not come from any books or anything”, and they just “implement something, learn it, and then socialize it” (P12I).

Participants also highlighted a lack of awareness about metrics for assessing sustainability, particularly in measuring the energy consumption of ML. P22I, who recently began working as a Responsible AI lead, expressed a strong desire to address sustainability in their new role but admitted uncertainty about how to measure it: “Because we’re now debating whether sustainability should be one of the [responsible AI measures], it should be [...]. But we don’t know how to measure that.” (P22I). The need for widely-used metrics was further articulated by others. P15I, an AI solutions architect, pointed

out that ML practitioners “don’t understand how the things that they develop can be translated to energy”, and a metric is therefore necessary. P4A, a PhD student in AI, also mentioned that the main issue is that there is no metric to evaluate.

Participants further exemplified how such metric and knowledge could inform the model selection process while adopting or incorporating available models (P18A, P12I).

4.3.3 Work environment and local policy. Participants repeatedly suggested the need for local workplace policies, discussing the power of administrative requirements for rationalizing energy use (P7I, P9A, P10A, P12I, P17I, P23A). Although this stance was shared across academia and industry, participants from academia further discussed potential practical application areas. For instance, P10A, who was finishing their PhD at the time of the interview, discussed the need for more reporting mechanisms to be applied from the research design stage. “When you’re doing your research and planning your research projects, [...] from early PhD stages, could be really useful. [...] considering [it] from the design point and not just the endpoint...” (P10A). P14A further suggested concrete examples like enforcing the use of carbon tracking tools by publication venues.

The effect of the senior managerial and administrative staff on sustainability practices in ML processes (P5I, P11IA, P12I, P21A) was also repeatedly raised. P21A remarked that: “You need to definitely increase knowledge of the research team leader. This [is] the most important part because if my team leader gives me information, [...] I will reconsider about this.” (P21A).

4.3.4 Regulatory approaches. Many participants asserted the need for regulation to address the environmental footprint of ML operations (P1A, P5I, P8I, P10A, P11IA, P20A) at different scales (P1A), even if they perceived that defining the practicalities of this regulation was beyond their expertise (P5I, P10A). Participants debated a variety of approaches to policy ranging from carbon taxing (P6A, P10A, P13A, P22I), to carbon reporting (P5I, P10A, P14A), limits to data retention (P19I, P23A), monetary incentives (P3A, P5I, P13A, P14A, P23A), penalty structures (P5I, P14A, P23A), limits to data center sizes and facilities (P13A, P23A), ‘forcing for more sustainable sources of energy’ (P3A), licensing/external auditing (P6A) or limits/quota to computation or model size (P3A, P5I, P9A, P11IA, P18A, P20A). Diversity of these approaches highlights that reducing ML’s environmental impact requires interventions at multiple levels of the technological pipeline.

Furthermore, the talk of regulatory solutions also resonates with the perception that the challenges are discipline-based and not (only) individual (Section 4.2). For instance, P5I explained how the cap on energy usage should be domain-specific and even technology-specific (P16IA) as they have different needs.

Others also discussed the governmental, institutional but also intergovernmental actors in policy-making that could lead such a process — “Maybe European Union or the UNDP, or maybe the nation: can make an announcement [...] If you research or if you support these parameters in your research group, we will also give you these foundations.” (P14A), or mentioned how regulation will need to ‘catch-up’ (P3A) with the speed of ML innovation, suggesting a dynamic and adaptive policy approach rather than static regulations. Sustainability policies from other industries like

motor vehicles (P6A, P23A), household energy consumption (P6A) or plastic bottle recycling (P23A) were also brought up by the participants to support their ideas for domain-regulation. While these references were mostly to provide instances of how incentives or taxes worked for sustainability goals, they were also referenced to communicate that there will always be workarounds for such policies. P6A drew parallels with an infamous case of an EU car company. "[...] their car failed to pass to tests, but they put some code piece to the car so that car can understand if it's on the test." (P6A).

Furthermore, there was a widespread belief across participants that such regulations would be hard to be imposed successfully (P16IA, P5I, P21A). Participants expressed similar views, stating: that 'it's really hard for these to happen' (P5I), that 'nothing can be regulated about this' (P21A), or even that they 'will be amazed if I ever see any regulation directly related to AI and environmental impact' (P16IA). This shared skepticism among participants indicates the tension between the recognized need for systemic change and doubts about its feasibility within the current socio-political climate.

While discussing the reasons for this disbelief, some referred to companies' lobbying activities (P10A), while others spoke about inter-governmental competition (P20A), or even that ML is an umbrella field not to be treated as a single entity with clear boundaries (P21A). The difficulty of the various actors using ML from large corporations to individuals was also viewed as a difficulty for regulatory practice (P6A, P14A, P16IA, P18A, P23A): "[They would need to] track companies rather than people or the individual" (P14A). These insights together portray a regulatory landscape that has technical, political, and structural challenges.

4.3.5 Culture shift. Bigger scale or systemic issues and solutions were also raised to point out to need for a culture shift. A common observation among participants (P13A, P14A, P15I, P18A, P19I, P23A) included the over-use/over-reliance on ML, calling it 'a sledgehammer to crack a nut' (P23A). Avoiding this 'misplaced hype' (P19I) on LLMs or ML in general was discussed as part of the solution. "Whenever customers come to me and ask for the solution, I say, if your problem doesn't need AI, then don't go for AI. I always advise if your problem can be solved with a simple Excel bot or like a macro, then why are you going for like a LLM agent? You don't have to build just because you have LLM, and have access to them, that doesn't mean like you can build an agent for all the problems." (P15I). This discussion highlights a call for more appropriate, context-sensitive use of ML and resisting the urge to over-engineer simple solutions. Participants also discussed the need for slowing down (P5I, P11IA, P12I). Some saw regulation and external constraints not as obstacles but as opportunities for fostering more thoughtful and sustainable innovation. P11IA for instance, explicitly mentioned "laws are slowing this [AI growth] down [...]. But this also does mean that you can make it sustainable on the way. [...] it could also be an opportunity for some institutions to create more sustainable solutions just as well because it slows down." (P11IA).

In addition, Some participants spoke about the need for public pressure to drive change (P7I, P10A, P13A, P20A). P13A for instance stated that the public recognition of the problem would be "a big

contributor" for people [practitioners] to realize the harm of their work on the environment. P7I also argued that the issue has more to do with society rather than individual companies, since shareholders are primarily driven by 'need to make money' while making decisions. "If you want the shareholders to care about this, you first need to make [the] public care about this" they discussed (P7I).

While this standpoint may sound like shifting the responsibility onto others, it is undeniable that societal engagement is also a crucial part of promoting more sustainable practices in AI development. A more in-depth discussion on the issue of responsibility will be elaborated in Section 5.1.

5 Discussion and Future Directions

Addressing the lack of research regarding ML practitioners' understandings and perceptions on environmental sustainability of ML, we conducted 23 interviews with practitioners from a wide range of roles and domains. From the rich data we collected we identified that despite the growing interest and knowledge about sustainability of ML among practitioners, sustainability was hard to move up the ladder of priorities in the fast-paced, competitive and exploratory field of ML development and also that individual change is difficult. Summarizing our findings, we discuss two overarching questions that ran through our interviews: (a) whose responsibility environmental sustainability is, (b) who has the power and agency for action and accordingly, how we can move forward from here.

5.1 Who holds responsibility for environmental sustainability?

Participants often perceived sustainability as the responsibility of others, referring to big tech companies, decision-makers, managers at work, other domains or colleagues working on different phases of the pipeline. These participants typically viewed their own work as less environmentally harmful than others, either due to its scale (e.g., using a small model, a small lab environment etc.), the nature of their work/phase (e.g., not doing training but only data processing) or its perceived high value to society (e.g., healthcare applications).

These arguments are consistent with prior work. For instance, past interviews with researchers working in the field of biobanking found similar value-based judgments of giving healthcare a 'free pass' regarding its environmental impact solely because of the domain just as the healthcare ML researchers claimed in our interviews [93]. Previous research also hypothesized that training is over-emphasized as having the main impact [71]. This was clearly visible in our data given how few discussed embodied emissions even when we know that inference together with embodied emissions are equally as impactful [71, 85, 109], especially depending on the task at hand [70]. Such siloed thinking in domains and in stages exclusively also has parallels to the individuality found in early environmental discourse [26, 41].

We thus argue that ML practitioners should be empowered to acknowledge that they (as individuals) also have responsibility regarding their choices. While indeed some large entities may have the majority of emissions, climate change is a collective issue to which we all contribute to and are affected by (even if disproportionately [101]) so that both individual and collective responsibility and action are required [10, 43, 54]. One way to achieve this awareness

is through dedicated data and frameworks that help individuals contextualize their impact. Sustainable HCI research [32] through persuasive technology assumes that some sort of ‘trigger’ [38] or feedback is necessary to raise awareness or even behavior change. Equally in our interviews, participants regularly expressed their need for ‘feedback’, pointing to data proxies, or clear information regarding their actions and decisions with their work not only to change their behaviors individually but also to socialize the information within their teams thus highlighting this need.

While early, there are some promising examples towards this goal. For instance the Green DiSC certification scheme introduced by the Software Sustainability Institute [64], or the draft Web Sustainability Guidelines prepared by the W3C [28], both help clarify and then redistribute responsibility within team environments spanning various roles whose decisions can not be evaluated in silos or independent from each other.

Various eco-feedback tools have also been developed to estimate energy consumption and carbon emissions of computational processes [5, 61, 63, 98] (see also Section 2) even if they have been critiqued for their focus on efficiency [108] and for falling short of moving beyond individual change [50]. Such solutions, together with more open-data and transparent sustainability operations from the ML industry [69] can be a crucial step forward towards individual and collective responsibility.

Furthermore, as the field grows exponentially, and new practitioners are entering the field through increasingly diverse pathways the knowledge gap regarding environmental sustainability of ML work becomes further institutionalized through established conventional training paradigms and norms. While CS and ML education is training practitioners to be highly skilled at optimizing for accuracy they often still lack the integration of awareness, knowledge and tools to properly evaluate environmental responsibility in ML work. Echoing the learnings of literature on integration of ethics into CS education [37], we argue that ethics within ML and/or CS education frameworks [92] should also bring an emphasis on integration of environmental sustainable thinking into the CS and ML training process.

5.2 Who has the power and agency to act?

Participants discussed a range of perceived limitations to take action, often pointing to the work culture of the industry and domain, emphasizing the pressure for speed and competitiveness but also hierarchies and multi-level decision-making structures at work. Research on resistance and organizational communication within organizations [77, 97] has shown that even within these type of constrained contexts certain types of action can still drive change, even though these actions should be critically examined within the limits of employee agency who navigate precarity within larger capitalist systems.

Our participants spoke about the need to do their task no matter what, referring to the limits of their decision-making ability regarding what factors needed to be prioritized at work. Previous literature on pro-environmental behavior studies suggested that the perception of agency for sustainable behavior at workplace majorly differs from households [11, 81]. In workplaces, ‘breaking out of old roles’ and adopting new behaviors appears to be

harder due to dependency on the company deep-grounded norms, values, roles and routines beyond the employees’ personal norms and values [11, 81]. Moreover, while personal opinions can support sustainable behavior change in the household, within workplace settings such change in actions require factual data for justification [81]. This parallels other areas of ethical decision-making at work as well, where the lack of ethical considerations often is a result of bureaucratic structures or workplace culture [11, 19, 56].

Previous research on responsible AI displays similar findings, highlighting the lack of organizational support [3, 53, 89] and misaligned interests among the industry (e.g., towards software product launches [3]) as reasons for the lack of adoption. This echoes our findings on the limitations of ML practitioners’ decision-making autonomy, workplace hierarchy and the ‘fast’ and ‘competitive’ nature of ML work as a significant factor contributing to inaction on environmental sustainability. While considerations like speed, market competition, or time management are assigned a higher priority, environmental sustainability is still seen as ‘nice-to-have’ as expressed by our participants. Liu et al. [67] examining impact statements from NeurIPS conference similarly documented a lack of agency and responsibility in the statements’ tone of writing. Similarly, Widder and Nafus [106] describe how responsible artificial intelligence interventions, like ethics checklists and guidelines, assume panoptic knowledge and control over systems whereas in their daily reality, developers did not find that these are within their agency, or responsibility to address due to the inherent modularity of AI development.

Our participants often referred to regulation as a driver or motivator for action. The environmental impact of ML has been recently catching the attention of regulatory bodies and intergovernmental organizations, some of which started to consider environmental sustainability as part of the fair and responsible AI efforts amid the accelerating climate emergency [36, 74, 104]. While generally in favor, our participants expressed skepticism around regulation pointing to the difficulties of operationalization, tech companies’ lobbying activities as well as the uncertainty and difficulty of measuring ML’s environmental impact. They also mention a need for more interdisciplinarity and dialogue among policy-makers, civil society as well as technical and non-technical ML professionals to ensure better policy implementation. We add that additionally including environmental psychologists and HCI experts in this discourse will also help design for the human (and non-human) actors of this issue and not just focus on its technical dimensions.

Finally, we believe the types of data we argued for in the previous section can also act as a stepping stone to enable a more collective understanding of the issue and ultimately support local policy and collective action. Previous work has shown for instance how public displays of temperature within workplaces has brought about critique and change of local policy [22]. Data have also brought communities together to rally across common issues [30, 31, 107] like food security and housing. We argue that similar socio-technical initiatives are necessary within ML settings to examine the boundaries of perceived and actual agency and can help build *local* policy even before more generalized one. There are early examples of such initiatives. For instance, conferences [79], repositories like HuggingFace [72], and even funding schemes [39, 62, 78] are now encouraging applicants to estimate their environmental impact at

the time of submission. Nevertheless, there is still a lack of broad adoption of these efforts to date —as also evidenced in our interviews and by the lack of environmental impact information in model cards [66], highlighting the need for further efforts.

While our findings align with the earlier responsible AI literature on challenges and perceptions of practitioners, our interviews indicate, that practitioners' perceptions of environmental sustainability and subsequent feelings of agency and responsibility are qualitatively different in two ways. First, responsible AI is a well-established field that has been extensively studied through a range of frameworks, from both technical and social perspectives. In contrast, sustainable ML is a relatively new and emerging area that only recently started gaining traction. While sustainability metrics and frameworks are being increasingly developed and are becoming more common, they are not yet widely adopted or even well-known among practitioners, resulting in a knowledge gap in the industry. Comparing to Widder and Nafus [106], for instance, we see it is not so much the development process modularity as much as a lack of awareness of what the responsibility is – perhaps because the discourse has been centered around the training of large foundational models overlooking many other actions in the process.

Another fundamental distinction lies in the nature of environmental impact more broadly. Responsible AI frameworks address who AI is responsible for, who it may put at risk, and who it has the potential to harm [67], making its ethical concerns more immediate and directly addressing a community. In contrast, the environmental impact of AI has a more diffused and localized effect on individuals (for instance due to communities living near data centers [68]). Unlike fairness, where harm is often visible and its impact is on clearly identifiable (even if marginalized) groups, the consequences of environmental degradation unfold gradually over time, making them harder to grasp and internalize. This distinction is significant because it influences how these responsible AI topics are prioritized.

5.3 Limitations

As all research, this work also has limitations. First, the participants for the most part discussed energy consumption over other environmental impacts such as water use and biodiversity. While we believe that this is shaped by the current state of the discussion in the media and research more broadly, this was also possibly intensified by our questions around eco-feedback energy tools. Accordingly, we encourage future studies to explicitly explore environmental impact beyond energy consumption, such as water use and biodiversity, possibly by broadening the scope of questions or tools used. Second, our participants were self-selected so that there was a clear bias towards individuals who had sustainability sensitivities. While this is a limitation of the work, given the findings and the overall 'disappointment' and lack of awareness even among this population, we believe that there is transferability of our findings. Our participant sample is also not geographically well-distributed, therefore we acknowledge that more work is needed to claim generalisability and global representability of these findings. While we provide our findings and discussion as deep insights into relatively few participants, we expect that given the distributed nature of ML work that many of these ideas would be transferrable. It is crucial

to note that the focus field is fast changing so that it is unclear what the state of events will be in a few years in terms of technology and sustainability perceptions around it. We thus encourage periodic studies to capture how perceptions and practices change over time. Finally, future research can leverage this study's insights as a starting point to design more inclusive and scalable methodologies or develop tools that address the identified gaps.

6 Conclusion

Despite the growing body of research on technical solutions in the field, practitioners' perceptions and experiences with environmental sustainability of ML work in particular remains largely undocumented. We thus conducted 23 interviews with ML practitioners, revealing that despite the growing interest on the topic, sustainability hasn't yet been prioritized in the fast-paced ML settings and that practitioners perceive individual change to be difficult and ineffective to be realized. We discussed our findings through two core questions on whose responsibility environmental sustainability is and who has the power and agency for action and argued for future steps to be taken across various stakeholders. Providing a rich overview of perceptions, practical challenges and future needs of ML practitioners to prioritize sustainability in their workflows, we argue that this work presents a much needed nuanced understanding of user interaction for the design and development of future ML feedback tools and beyond.

Researchers' positionality

As authors, we acknowledge that our professional and personal backgrounds inevitably shaped not only how this qualitative research was designed, but also how it was conducted and analyzed. As researchers grounded in the fields of sustainable HCI and environmental justice, this research was driven by our shared interest and commitment in addressing the climate change and the role of ICT sector and specifically ML work in it. While none of us have a background in ML (second and last authors have a background in CS), our interview design allowed us to comprehend practitioners' experiences and views without filtering them through our own practical experience or technical expertise. Still, we remained vigilant not to impose our assumptions regarding both the scale and scope of the problem but also varying understandings, perceptions and levels of engagement of practitioners with the topic.

Lastly, we also acknowledge that the demographic distribution of the interviewees was partially affected by the authors' personal backgrounds. Considering that all four authors are based in the UK, the majority of our participants were also based in the UK. Furthermore, the first and second authors being from Turkey and China respectively also contributed to the visibility of our participation call among ML communities within those two countries.

Acknowledgments

The authors would like to thank the interview participants for their time and insight. They would also like to thank Jerone Andrews as well as the anonymous reviewers for their close reading and helpful suggestions on the manuscript. This work was supported by the King's Climate & Sustainability Seed Fund scheme.

References

- [1] Ahmad Abusukhon, Bilal Hawashin, and Mohammad Lafi. 2021. An Efficient System for Reducing the Power Consumption in Offices Using the Internet of Things. *International Journal of Advances in Soft Computing & Its Applications* 13, 1 (2021). <https://www.i-csrs.org/Volumes/ijasca/2021.1.1.pdf>
- [2] Open AI. 2018. AI and compute. <https://openai.com/research/ai-and-compute> (Retrieved January 21, 2025).
- [3] Sanna J. Ali, Angèle Christin, Andrew Smart, and Riitta Katila. 2023. Walking the Walk of AI Ethics: Organizational Challenges and the Individualization of Risk among Ethics Entrepreneurs. In *2023 ACM Conference on Fairness Accountability and Transparency (FAccT '23)*. ACM, 217–226. <https://doi.org/10.1145/3593013.3593990>
- [4] Yehia Ibrahim Alzoubi and Alok Mishra. 2024. Green artificial intelligence initiatives: Potentials and challenges. *Journal of Cleaner Production* 468 (Aug. 2024), 143090. <https://doi.org/10.1016/j.jclepro.2024.143090>
- [5] Lasse F. Wolff Anthony, Benjamin Kanding, and Raghavendra Selvan. 2020. Carbontracker: Tracking and Predicting the Carbon Footprint of Training Deep Learning Models. <https://doi.org/10.48550/arXiv.2007.03051>
- [6] Giovanni Baiocchi, Jan Minx, and Klaus Hubacek. 2010. The Impact of Social Factors and Consumer Behavior on Carbon Dioxide Emissions in the United Kingdom. *Journal of Industrial Ecology* 14, 1 (2010), 50–72. <https://doi.org/10.1111/j.1530-9290.2009.00216.x>
- [7] Mary Barreto, Evangelos Karapanos, and Nuno Nunes. 2013. Why don't families get along with eco-feedback technologies? a longitudinal inquiry. In *Proceedings of the Biannual Conference of the Italian Chapter of SIGCHI (CHI-taly '13)*. Association for Computing Machinery, New York, NY, USA, 1–4. <https://doi.org/10.1145/2499149.2499164>
- [8] Ben Bedwell, Enrico Costanza, and Michael O. Jewell. 2016. Understanding Energy Consumption at Work: Learning from Arrow Hill. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. Association for Computing Machinery, New York, NY, USA, 1337–1348. <https://doi.org/10.1145/2818048.2819993>
- [9] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Virtual Event Canada, 610–623. <https://doi.org/10.1145/3442188.3445922>
- [10] Mike Berners-Lee. 2020. *How Bad Are Bananas?: The carbon footprint of everything*. Profile Books.
- [11] Vincent Blok, Renate Wesselink, Oldrich Studynka, and Ron Kemp. 2015. Encouraging sustainability in the workplace: a survey on the pro-environmental behaviour of university employees. *Journal of Cleaner Production* 106 (2015), 55–67. <https://doi.org/10.1016/j.jclepro.2014.07.063>
- [12] Christine Boomsma, Julie Goodhew, Sabine Pahl, and Rory V. Jones. 2016. The feasibility of saving energy in challenging organisational contexts: Testing energy visualisation in a social services office in the United Kingdom. *Energy Research & Social Science* 15 (May 2016), 58–74. <https://doi.org/10.1016/j.erss.2016.02.004>
- [13] Aurélien Bourdon, Adel Noureddine, Romain Rouvoy, and Lionel Seinturier. 2013. Powerapi: A software library to monitor the energy consumed at the process-level. *ERCIM News* 92 (2013), 43–44.
- [14] Virginia Braun and Victoria Clarke. 2023. Thematic analysis. In *APA handbook of research methods in psychology: Research designs: Quantitative, qualitative, neuropsychological, and biological, Vol. 2, 2nd ed.* American Psychological Association, Washington, DC, US, 65–81. <https://doi.org/10.1037/0000319-004>
- [15] Christina Bremer, Oliver Bates, Christian Remy, Alexandra Gormally-Sutton, Bran Knowles, and Adrian Friday. 2023. COVID-19 as an Energy Intervention: Lockdown Insights for HCI. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*. Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3544549.3585896>
- [16] Dirk Brounen, Nils Kok, and John M. Quigley. 2013. Energy literacy, awareness, and conservation behavior of residential households. *Energy Economics* 38 (July 2013), 42–50. <https://doi.org/10.1016/j.eneco.2013.02.008>
- [17] Brian Calvert. 2024. AI already uses as much energy as a small country. It's only the beginning. <https://www.vox.com/climate/2024/3/28/24111721/climate-ai-tech-energy-demand-rising> (Retrieved January 21, 2025).
- [18] Amanda R. Carrico and Manuel Riemer. 2011. Motivating energy conservation in the workplace: An evaluation of the use of group-level feedback and peer education. *Journal of Environmental Psychology* 31, 1 (March 2011), 1–13. <https://doi.org/10.1016/j.jenvp.2010.11.004>
- [19] Cathy Cassell, Phil Johnson, and Ken Smith. 1997. Opening the Black box: Corporate codes of ethics in their organizational context. *Journal of Business Ethics* 16, 10 (1997), 1077–1093.
- [20] Nico Castelli, Niko Schönau, Gunnar Stevens, Tobias Schwartz, and Timo Jakobi. 2015. Role-based Eco-info Systems: An Organizational Theoretical View of Sustainable HCI at Work. *Info Systems* (2015).
- [21] Andrew R. Chow and Billy Perrigo. 2023. The AI Arms Race Is Changing Everything. <https://web.archive.org/web/20240727002702/https://time.com/6255952/ai-impact-chatgpt-microsoft-google/> (Retrieved January 21, 2025).
- [22] Enrico Costanza, Ben Bedwell, Michael O. Jewell, James Colley, and Tom Rodden. 2016. 'A bit like British Weather, I suppose': Design and Evaluation of the Temperature Calendar. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose California USA, 2016-05-07). ACM, 4061–4072. <https://doi.org/10.1145/2858036.2858367>
- [23] Benoît Courty, Victor Schmidt, Kamal Goyal, Coutarel Marion, Boris Feld, Jérémy Lecourt, LiamConnell, SabAmine, inimiz, supatomic, Mathilde Léval, Luis Blanche, Alexis Cruveiller, ouminasara, Franklin Zhao, Aditya Joshi, Alexis Bogroff, Amine Saboni, Hugues de Lavoreille, Niko Laskaris, Edoardo Abati, Douglas Blank, Ziyao Wang, Armin Catovic, alencon, Michał Stechly, Christian Bauer, Lucas-Otavio, JPW, and MinervaBooks. 2024. mlco2/codecarbon: v2.4.1. <https://doi.org/10.5281/zenodo.11171501>
- [24] Kate Crawford. 2021. *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale University Press, New Haven. <https://doi.org/10.2307/j.ctvt1ghv45t>
- [25] Kate Crawford. 2024. Generative AI's environmental costs are soaring — and mostly secret. *Nature* 626, 8000 (2024), 693–693. <https://doi.org/10.1038/d41586-024-00478-x>
- [26] Elizabeth Cripps. 2013. *Climate Change and the Moral Agent: Individual Duties in an Interdependent World*. OUP Oxford.
- [27] Sarah Darby et al. 2006. The effectiveness of feedback on energy consumption. *A Review for DEFRA of the Literature on Metering, Billing and direct Displays* 486, 2006 (2006), 26.
- [28] Alexander Dawson. 2024. Web Sustainability Guidelines (WSG). <https://w3c.github.io/sustainableweb-wsg/>. (Retrieved January 17, 2025).
- [29] Wesley Hanwen Deng, Manish Nagireddy, Michelle Seng Ah Lee, Jatinder Singh, Zhiwei Steven Wu, Kenneth Holstein, and Haiyi Zhu. 2022. Exploring How Machine Learning Practitioners (Try To) Use Fairness Toolkits. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Seoul Republic of Korea, 473–484. <https://doi.org/10.1145/3531146.3533113>
- [30] Catherine D'Ignazio. 2022. Creative data literacy: Bridging the gap between the data-haves and data-have nots. (2022), 6–18. <https://doi.org/10.1075/idj.23.1.03dig>
- [31] Catherine D'Ignazio and Lauren F. Klein. 2020. *Data Feminism*. The MIT Press. <https://doi.org/10.7551/mitpress/11805.001.0001>
- [32] Carl DiSalvo, Phoebe Sengers, and Hrönn Brynjarsdóttir. 2010. Mapping the landscape of sustainable HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (CHI '10). Association for Computing Machinery, New York, NY, USA, 1975–1984. <https://doi.org/10.1145/1753326.1753625>
- [33] Graham N. Dixon, Mary Beth Deline, Katherine McComas, Lauren Chambliss, and Michael Hoffmann. 2015. Using Comparative Feedback to Influence Workplace Energy Conservation: A Case Study of a University Campaign. *Environment and Behavior* 47, 6 (July 2015), 667–693. <https://doi.org/10.1177/0013916513520417>
- [34] M. Economidou, B. Atanasiu, C. Despret, J. Maio, I. Nolte, O. Rapf, J. Laustsen, P. Ruyssevelt, D. Staniaszek, D. Strong, and S. Zinetti. 2011. Europe's buildings under the microscope. A country-by-country review of the energy performance of buildings. *Buildings Performance Institute Europe* (Oct. 2011). <https://www.osti.gov/etdweb/biblio/21514343>
- [35] Ufuk Erdal. 2023. Quenching Big-data's Thirst: A Novel Water Recycling Strategy for Datacenter Cooling & Community Reuse. In *WEFTEC 2023*. Water Environment Federation. <https://doi.org/10.2175/193864718825159112>
- [36] EU Parliament. 2023. EU AI Act: first regulation on artificial intelligence. <https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence> (Retrieved January 21, 2025).
- [37] Casey Fiesler, Natalie Garrett, and Nathan Beard. 2020. What Do We Teach When We Teach Tech Ethics? A Syllabi Analysis. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education* (Portland, OR, USA) (SIGCSE '20). Association for Computing Machinery, New York, NY, USA, 289–295. <https://doi.org/10.1145/3328778.3366825>
- [38] Bj Fogg. 2009. A behavior model for persuasive design. In *Proceedings of the 4th International Conference on Persuasive Technology*. ACM, Claremont California USA, 1–7. <https://doi.org/10.1145/1541948.1541999>
- [39] Environmental Association for Universities and Colleges. 2024. Concordat for the Environmental Sustainability of Research and Innovation Practice. <https://wellcome.org/who-we-are/positions-and-statements/environmental-sustainability-concordat>. (Retrieved January 21, 2025).
- [40] Derek Foster, Shaun Lawson, Jamie Wardman, Mark Blythe, and Conor Linehan. 2012. "Watts in it for me?": design implications for implementing effective energy interventions in organisations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '12). Association for Computing Machinery, New York, NY, USA, 2357–2366. <https://doi.org/10.1145/2207676.2208396>

- [41] Augustin Fragnière. 2016. Climate change and individual duties. *WIREs Climate Change* 7, 6 (2016), 798–814. <https://doi.org/10.1002/wcc.422>
- [42] Charlotte Freitag, Mike Berners-Lee, Kelly Widdicks, Bran Knowles, Gordon S. Blair, and Adrian Friday. 2021. The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations. *Patterns* 2, 9 (Sept. 2021). <https://doi.org/10.1016/j.patter.2021.100340>
- [43] Immo Fritsche and Torsten Masson. 2021. Collective climate action: When do people turn into collective environmental agents? *Current Opinion in Psychology* 42 (Dec. 2021), 114–119. <https://doi.org/10.1016/j.copsyc.2021.05.001>
- [44] Jon Froehlich, Leah Findlater, and James Landay. 2010. The design of eco-feedback technology. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. Association for Computing Machinery, New York, NY, USA, 1999–2008. <https://doi.org/10.1145/1753326.1753629>
- [45] Óscar García, Ricardo S. Alonso, Javier Prieto, and Juan M. Corchado. 2017. Energy Efficiency in Public Buildings through Context-Aware Social Computing. *Sensors* 17, 4 (April 2017), 826. <https://doi.org/10.3390/s17040826>
- [46] Google. 2024. 2024 Environmental Report. <https://sustainability.google/reports/google-2024-environmental-report/> (Retrieved January 21, 2025).
- [47] Kirsten Gram-Hanssen. 2010. Standby Consumption in Households Analyzed With a Practice Theory Approach. *Journal of Industrial Ecology* 14, 1 (2010), 150–165. <https://doi.org/10.1111/j.1530-9290.2009.00194.x>
- [48] Alice Grønhoj and John Thøgersen. 2011. Feedback on household electricity consumption: learning and social influence processes. *International Journal of Consumer Studies* 35, 2 (2011), 138–145. <https://doi.org/10.1111/j.1470-6431.2010.00967.x>
- [49] Gianluca Guidi, Francesca Dominici, Jonathan Gilmour, Kevin Butler, Eric Bell, Scott Delaney, and Falco J. Bargagli-Stoffi. 2024. Environmental Burden of United States Data Centers in the Artificial Intelligence Era. <https://doi.org/10.48550/arXiv.2411.09786>
- [50] Sinem Görücü, Luiz A Morais, and Georgia Panagiotidou. 2025. A Critical Analysis of Machine Learning Eco-feedback Tools through the Lens of Sustainable HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, Yokohama, Japan, 28 pages. <https://doi.org/10.1145/3706598.3713198>
- [51] Walid A. Hanafy, Qianlin Liang, Noman Bashir, David Irwin, and Prashant Shenoy. 2023. CarbonScaler: Leveraging Cloud Workload Elasticity for Optimizing Carbon-Efficiency. *Proc. ACM Meas. Anal. Comput. Syst.* 7, 3 (Dec. 2023), 57:1–57:28. <https://doi.org/10.1145/3626788>
- [52] Mél Hogan. 2024. The Fumes of AI. *Critical AI* 2, 1 (April 2024). <https://doi.org/10.1215/2834703X-11205231>
- [53] Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé, Miro Dudik, and Hanna Wallach. 2019. Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need?. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3290605.3300830>
- [54] Sàde Hormio. 2023. Collective responsibility for climate change. *WIREs Climate Change* 14, 4 (2023), e830. <https://doi.org/10.1002/wcc.830>
- [55] Rachel A. Howell. 2018. Carbon management at the household level: a definition of carbon literacy and three mechanisms that increase it. *Carbon Management* 9, 1 (Jan. 2018), 25–35. <https://doi.org/10.1080/17583004.2017.1409045>
- [56] Robert Jackall. 1988. *Moral mazes : the world of corporate managers*. New York : Oxford University Press, 1988. <https://search.library.wisc.edu/catalog/99958842202121>
- [57] Mohit Jain, Ankit Agrawal, Sunil K. Ghai, Khai N. Truong, and Deva P. Seetharam. 2013. "We are not in the loop": resource wastage and conservation attitude of employees in indian workplace. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing (UbiComp '13)*. Association for Computing Machinery, New York, NY, USA, 687–696. <https://doi.org/10.1145/2493432.2493444>
- [58] Marc Jentsch, Marco Jahn, Ferry Pramudianto, Jonathan Simon, and Amro Al-Akkad. 2011. An Energy-Saving Support System for Office Environments. In *Human Behavior Understanding*. Albert Ali Salah and Bruno Lepri (Eds.). Springer, Berlin, Heidelberg, 83–92. https://doi.org/10.1007/978-3-642-25446-8_10
- [59] Lynn H. Kaack, Priya L. Donti, Emma Strubell, George Kamiya, Felix Creutzig, and David Rolnick. 2022. Aligning artificial intelligence with climate change mitigation. *Nature Climate Change* 12, 6 (June 2022), 518–527. <https://doi.org/10.1038/s41558-022-01377-7>
- [60] Peng Kuai, Xiaoxiao Zhang, Shuan Zhang, and Jing Li. 2022. Environmental awareness and household energy saving of Chinese residents: Unity of knowing and doing or easier said than done? *Journal of Asian Economics* 82 (Oct. 2022), 101534. <https://doi.org/10.1016/j.asieco.2022.101534>
- [61] Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. 2019. Quantifying the Carbon Emissions of Machine Learning. <https://doi.org/10.48550/arXiv.1910.09700>
- [62] Loïc Lannelongue, Juliette Propier, and Even Matencio. 2025. How to include environmental sustainability criteria in national AI funding schemes? Reflecting on the example of France and the Green Algorithms tool. (Jan. 2025). <https://doi.org/10.5281/ZENODO.14607021>
- [63] Loïc Lannelongue, Jason Grealey, and Michael Inouye. 2021. Green Algorithms: Quantifying the Carbon Footprint of Computation. *Advanced Science (Weinheim, Baden-Württemberg, Germany)* 8, 12 (June 2021), 2100707. <https://doi.org/10.1002/adv.202100707>
- [64] Loïc Lannelongue and Kirsty Pringle. 2024. Announcing Green DiSC, the new Digital Sustainability Certification! <https://www.software.ac.uk/GreenDiSC>. (Retrieved January 15, 2025).
- [65] Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren. 2023. Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. <https://doi.org/10.48550/arXiv.2304.03271>
- [66] Weixin Liang, Nazneen Rajani, Xinyu Yang, Ezinwanne Ozoani, Eric Wu, Yiqun Chen, Daniel Scott Smith, and James Zou. 2024. Systematic analysis of 32,111 AI model cards characterizes documentation practice in AI. 6, 7 (2024), 744–753. <https://doi.org/10.1038/s42256-024-00857-z>
- [67] David Liu, Priyanka Nanayakkara, Sarah Ariyan Sakha, Grace Abuhamad, Su Lin Blodgett, Nicholas Diakopoulos, Jessica R. Hullman, and Tina Eliassi-Rad. 2022. Examining Responsibility and Deliberation in AI Impact Statements and Ethics Reviews. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society* (Oxford, United Kingdom) (AI/ES '22). Association for Computing Machinery, New York, NY, USA, 424–435. <https://doi.org/10.1145/3514094.3534155>
- [68] Clara Hernanz Lizarraga and Olivia Solon. 2023. Thirsty data centres are making Europe's dry summers even drier. *The Straits Times* (July 2023). <https://www.straitstimes.com/world/europe/thirsty-data-centres-are-making-europe-s-dry-summers-even-drier>
- [69] Alexandra Sasha Luccioni. 2024. Generative AI and Climate Change Are on a Collision Course. <https://www.wired.com/story/true-cost-generative-ai-data-centers-energy/>. (Retrieved January 15, 2025).
- [70] Alexandra Sasha Luccioni, Yacine Jernite, and Emma Strubell. 2024. Power Hungry Processing: Watts Driving the Cost of AI Deployment?. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 85–99. <https://doi.org/10.1145/3630106.3658542>
- [71] Alexandra Sasha Luccioni, Sylvain Viguier, and Anne-Laure Ligozat. 2022. Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model. <https://doi.org/10.48550/arXiv.2211.02001>
- [72] Sasha Luccioni. 2024. Energy Scores for AI Models. <https://huggingface.co/blog/sasha/energy-star-ai-proposal>. (Retrieved January 21, 2025).
- [73] Michael Madaio, Shivani Kapania, Rida Qadri, Ding Wang, Andrew Zaldivar, Remi Denton, and Lauren Wilcox. 2024. Learning about Responsible AI On-The-Job: Learning Pathways, Orientations, and Aspirations. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAcCT '24)*. Association for Computing Machinery, New York, NY, USA, 1544–1558. <https://doi.org/10.1145/3630106.3658988>
- [74] Markey, Heinrich, Eshoo, Beyer. 2024. S.3732 - Artificial Intelligence Environmental Impacts Act of 2024. <https://www.congress.gov/bills/118th/congress/senate-bill/3732/text> (Retrieved January 21, 2025).
- [75] Patricia Mergen. 2024. Environmental Impact of Digital Technologies and Artificial Intelligence. *Biodiversity Information Science and Standards* 8 (2024), e141372.
- [76] Dan Milmo. 2024. Google's emissions climb nearly 50% in five years due to AI energy demand. *The Guardian* (July 2024). <https://www.theguardian.com/technology/article/2024/jul/02/google-ai-emissions> (Retrieved January 21, 2025).
- [77] Dennis K. Mumby, Robyn Thomas, Ignasi Marti, and David Seidl. 2017. Resistance Redux. *Organization Studies* 38, 9 (2017), 1157–1183. <https://doi.org/10.1177/0170840617717554>
- [78] FEBS Network. 2024. Heidelberg Agreement on Environmental Sustainability in Research Funding. <https://network.febs.org/posts/heidelberg-agreement-on-environmental-sustainability-in-research-funding>. (Retrieved January 21, 2025).
- [79] NeurIPS. 2023. Ethics Guidelines. <https://neurips.cc/Conferences/2023/EthicsGuidelines>. (Retrieved January 21, 2025).
- [80] Carman Neustaeter, Lyn Bartram, and Aaron Mah. 2013. Everyday activities and energy consumption: how families understand the relationship. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. Association for Computing Machinery, New York, NY, USA, 1183–1192. <https://doi.org/10.1145/2470654.2466153>
- [81] Michael Nye and Tom Hargreaves. 2010. Exploring the Social Dynamics of Proenvironmental Behavior Change: A Comparative Study of Intervention Processes at Home and Work. *Journal of Industrial Ecology* 14, 1 (Jan. 2010), 137–149. <https://doi.org/10.1111/j.1530-9290.2009.00193.x>
- [82] Anne Olhoff and John M. Christensen. 2020. *Emissions Gap Report 2020*. Report. UNEP DTU Partnership.

- [83] Brian Orland, Nilam Ram, Dean Lang, Kevin Houser, Nate Kling, and Michael Coccia. 2014. Saving energy in an office environment: A serious game intervention. *Energy and Buildings* 74 (May 2014), 43–52. <https://doi.org/10.1016/j.enbuild.2014.01.036>
- [84] Thanasis G. Papaioannou, Nikos Dimitriou, Kostas Vasilakis, Anthony Schoofs, Manolis Nikiforakis, Fabian Pursche, Nikolay Deliyiski, Amr Taha, Dimosthenis Kotsopoulos, Cleopatra Bardaki, Sarantis Kotsilitis, and Anastasia Garbi. 2018. An IoT-Based Gamified Approach for Reducing Occupants' Energy Wastage in Public Buildings. *Sensors* 18, 2 (Feb. 2018), 537. <https://doi.org/10.3390/s18020537>
- [85] David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Luis-Miquel Munguia, Daniel Rothchild, David R. So, Maud Texier, and Jeff Dean. 2022. The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink. *Computer* 55, 7 (July 2022), 18–28. <https://doi.org/10.1109/MC.2022.3148714>
- [86] James Pierce and Eric Paulos. 2012. Beyond energy monitors: interaction, energy, and emerging energy systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. Association for Computing Machinery, New York, NY, USA, 665–674. <https://doi.org/10.1145/2207676.2207771>
- [87] Zachary Pousman, Hafez Rouzati, and John Stasko. 2008. Imprint, a community visualization of printer data: designing for open-ended engagement on sustainability. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work (CSCW '08)*. Association for Computing Machinery, New York, NY, USA, 13–16. <https://doi.org/10.1145/1460563.1460568>
- [88] Hamed Nabizadeh Rafsanjani and Ali Ghahramani. 2020. Towards utilizing internet of things (IoT) devices for understanding individual occupants' energy usage of personal and shared appliances in office buildings. *Journal of Building Engineering* 27 (Jan. 2020), 100948. <https://doi.org/10.1016/j.jobe.2019.100948>
- [89] Bogdana Rakova, Jingying Yang, Henriette Cramer, and Rumman Chowdhury. 2021. Where Responsible AI meets Reality: Practitioner Perspectives on Enablers for Shifting Organizational Practices. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 7 (April 2021), 23 pages. <https://doi.org/10.1145/3449081>
- [90] Majken K. Rasmussen, Mia Kruse Rasmussen, Nervo Verdezoto, Robert Brewer, Laura L. Nielsen, and Niels Olof Bouvin. 2017. Exploring the flexibility of everyday practices for shifting energy consumption through clockcast. In *Proceedings of the 29th Australian Conference on Computer-Human Interaction (OzCHI '17)*. Association for Computing Machinery, New York, NY, USA, 296–306. <https://doi.org/10.1145/3152771.3152803>
- [91] Sarah Royston. 2014. Smart energy technologies in everyday life: smart Utopia? *Technology Analysis & Strategic Management* 26, 10 (Nov. 2014), 1242–1247. <https://doi.org/10.1080/09537325.2014.975789>
- [92] Jeffrey Saltz, Michael Skirpan, Casey Fiesler, Micha Gorelick, Tom Yeh, Robert Heckman, Neil Dewar, and Nathan Beard. 2019. Integrating Ethics within Machine Learning Courses. *ACM Trans. Comput. Educ.* 19, 4, Article 32 (Aug. 2019), 26 pages. <https://doi.org/10.1145/3341164>
- [93] Gabrielle Samuel, Hardcastle Faranak, and A. M. Lucassen. 2022. Environmental sustainability and biobanking: a pilot study of the field. *New Genetics and Society* 41, 2 (April 2022), 157–175. <https://doi.org/10.1080/14636778.2022.2093707>
- [94] Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni. 2020. Green AI. *Commun. ACM* 63, 12 (Nov. 2020), 54–63. <https://doi.org/10.1145/3381831>
- [95] Tobias Schwartz, Matthias Betz, Leonardo Ramirez, and Gunnar Stevens. 2010. Sustainable energy practices at work: understanding the role of workers in energy conservation. In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries (NordiCHI '10)*. Association for Computing Machinery, New York, NY, USA, 452–462. <https://doi.org/10.1145/1868914.1868966>
- [96] Tobias Schwartz, Gunnar Stevens, Timo Jakobi, Sebastian Deneff, Leonardo Ramirez, Volker Wulf, and Dave Randall. 2015. What People Do with Consumption Feedback: A Long-Term Living Lab Study of a Home Energy Management System. *Interacting with Computers* 27, 6 (Nov. 2015), 551–576. <https://doi.org/10.1093/iwc/iwu009>
- [97] James C. Scott. 1990. *Domination and the Arts of Resistance: Hidden Transcripts*. Yale University Press, New Haven, CT.
- [98] Omar Shaikh, Jon Saad-Falcon, Austin P Wright, Nilaksh Das, Scott Freitas, Omar Asensio, and Duen Horng Chau. 2021. EnergyVis: Interactively Tracking and Exploring Energy Consumption for ML Models. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems (CHI EA '21)*. Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3411763.3451780>
- [99] Anne Sharp and Meagan Wheeler. 2013. Reducing Householders' Grocery Carbon Emissions: Carbon Literacy and Carbon Label Preferences. *Australasian Marketing Journal* 21, 4 (Nov. 2013), 240–249. <https://doi.org/10.1016/j.ausmj.2013.08.004>
- [100] Jonathan Simon, Marco Jahn, and Amro Al-Akkad. 2012. Saving energy at work: the design of a pervasive game for office spaces. In *Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia (MUM '12)*. Association for Computing Machinery, New York, NY, USA, 1–4. <https://doi.org/10.1145/2406367.2406379>
- [101] Thomas, Leah. 2022. *The intersectional environmentalist: How to dismantle systems of oppression to protect people+ planet*. Souvenir Press, S.I.
- [102] Mikhail Tiutiunlikov, Vladimir D. Lazarev, A.N. Korovin, N. Zakharenko, Igor Doroshchenko, and Semen Budenny. 2023. eco4cast: Bridging Predictive Scheduling and Cloud Computing for Reduction of Carbon Emissions for ML Models Training. *Doklady Mathematics* 108, 2 (Dec. 2023), S443–S455. <https://doi.org/10.1134/S1064562423701223>
- [103] Evangelos Tolia, Enrico Costanza, Alex Rogers, Benjamin Bedwell, and Nick Banks. 2015. IdleWars: An Evaluation of a Pervasive Game to Promote Sustainable Behaviour in the Workplace. In *Entertainment Computing - ICEC 2015*, Konstantinos Chorianopoulos, Monica Divitini, Jannicke Baalsrud Hauge, Letizia Jaccheri, and Rainer Malaka (Eds.). Springer International Publishing, Cham, 224–237. https://doi.org/10.1007/978-3-319-24589-8_17
- [104] UNEP. 2024. AI has an environmental problem. Here's what the world can do about that. <https://www.unep.org/news-and-stories/story/ai-has-environmental-problem-heres-what-world-can-do-about>. (Retrieved December 17, 2024).
- [105] Roberto Verdecchia, June Sallou, and Luís Cruz. 2023. A systematic review of Green AI. *WIREs Data Mining and Knowledge Discovery* 13, 4 (July 2023), e1507. <https://doi.org/10.1002/widm.1507>
- [106] David Gray Widder and Dawn Nafus. 2023. Dislocated accountabilities in the "AI supply chain": Modularity and developers' notions of responsibility. *Big Data & Society* 10, 1 (Jan. 2023). <https://doi.org/10.1177/20539517231177620>
- [107] Sarah Williams. 2020. *Data Action: Using Data for Public Good*. The MIT Press. <https://doi.org/10.7551/mitpress/12261.001.0001>
- [108] Dustin Wright, Christian Igel, Gabrielle Samuel, and Raghavendra Selvan. 2023. Efficiency is Not Enough: A Critical Perspective of Environmentally Sustainable AI. <http://arxiv.org/abs/2309.02065>
- [109] Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga, Jinshi Huang, Charles Bai, et al. 2022. Sustainable ai: Environmental implications, challenges and opportunities. *Proceedings of Machine Learning and Systems* 4 (2022), 795–813.
- [110] Tien-Ju Yang, Yu-Hsin Chen, Joel Emer, and Vivienne Sze. 2017. A method to estimate the energy consumption of deep neural networks. In *2017 51st Asilomar Conference on Signals, Systems, and Computers*. 1916–1920. <https://doi.org/10.1109/ACSSC.2017.8335698> ISSN: 2576-2303.
- [111] Ray Yun, Azizan Aziz, Peter Scupelli, Bertrand Lasternas, Chenlu Zhang, and Vivian Loftness. 2015. Beyond Eco-Feedback: Adding Online Manual and Automated Controls to Promote Workplace Sustainability. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. Association for Computing Machinery, New York, NY, USA, 1989–1992. <https://doi.org/10.1145/2702123.2702268>
- [112] Ray Yun, Peter Scupelli, Azizan Aziz, and Vivian Loftness. 2013. Sustainability in the Workplace: Nine Intervention Techniques for Behavior Change. In *Persuasive Technology*, Shlomo Berkovsky and Jill Freyne (Eds.). Springer, Berlin, Heidelberg, 253–265. https://doi.org/10.1007/978-3-642-37157-8_30