What's Your Stake in Sustainability of AI?: An Informed Insider's Guide

Grace C. Kim, Annabel Rothschild, Carl DiSalvo, Betsy DiSalvo

School of Interactive Computing, Georgia Institute of Technology
85 5th St NW
Atlanta, Georgia 30332 USA
{gracekim3, arothschild, cdisalvo}@gatech.edu, bdisalvo@cc.gatech.edu

Abstract

It's no secret that AI systems come with a significant environmental cost. This raises the question: What are the roles and responsibilities of computing professionals regarding the sustainability of AI? Informed by a year-long informal literature review on the subject, we employ stakeholder identification, analysis, and mapping to highlight the complex and interconnected roles that five major stakeholder groups (industry, practitioners, regulatory, advocacy, and the general public) play in the sustainability of AI. Swapping the traditional final step of stakeholder methods (stakeholder *engagement*) for *entanglement*, we demonstrate the inherent entwinement of choices made with regard to the development and maintenance of AI systems *and* the people who impact (or are impacted by) these choices. This entanglement should be understood as a system of human and non-human agents, with the implications of each choice ricocheting into the use of natural resources and climate implications. We argue that computing professionals (AI-focused or not) may belong to multiple stakeholder groups, and that we all have multiple roles to play in the sustainability of AI. Further, we argue that the nature of regulation in this domain will look unlike others in environmental preservation (e.g., legislation around water contaminants). As a result, we call for ongoing, flexible bodies and policies to move towards the regulation of AI from a sustainability angle, as well as suggest ways in which individual computing professionals can contribute to fighting the environmental and climate effects of AI.

Introduction

The recent wave of generative AI is the latest in a continued trend of AI innovation. Among the many repercussions of AI systems are the resources used to make these systems possible – not least, those of the natural world, including electricity and water, used to maintain the extensive "cloud" infrastructures (Amoore 2020) on which these systems must run. For example, data centers alone are predicted to have an energy consumption roughly equal to that of the entirety of Japan by 2026 (Berreby 2024).

Given these substantial resource demands, addressing the sustainability of AI has become increasingly important. We note the distinction between two components that make up sustainable AI: AI for sustainability and sustainability of AI (Schwartz et al. 2019). The former refers to how AI is used to help achieve sustainability goals, while the latter refers to how AI is developed, deployed, and maintained in a manner that minimizes its environmental footprint (Sánchez-Pi and Martí 2021). While AI for sustainability could theoretically contribute to several UN Sustainable Development Goals (SDGs)[1], the use of AI can also hinder progress, as a result of the lack of sustainability in AI methods. Vinuesa et al. theorize that AI can facilitate 93% of the targets related to environmental SDGs, but it can also hinder 30% of them (Vinuesa et al. 2020). In this report, we are concerned with the sustainability of AI, rather than the potential of AI to promote sustainability, because sustainability of AI plays into all systems, including those used to promote sustainability.

Then, as computer science professionals, what sustainability issues are we entrusted to address when leveraging AI? Some among us who are AI system developers may see the immediate implications on sustainability, while those working in other subfields of computer science may not, which could lead them to perceive the sustainability of AI as less relevant to their work. In addition, there exists epistemological confusion surrounding the sustainability of AI systems; while some argue such systems are the only clear guide to counteracting continued climate change, others point to the massive energy consumption of such systems. In this work, we perform *stakeholder identification, analysis, and mapping* to illustrate both how far-reaching our duty is as computer science professionals and how we can act responsibly with regard to sustainability in our respective subfields.

Our work is predicated on a year-long stakeholder analysis and mapping, which encompassed an informal literature review and conversations with colleagues, including both self-titled AI "insiders" (Dwyer and Buckle 2009) and those in other subfields of computer science. Using our position as informed insiders, we share the results of our stakeholder analysis and mapping (Lelea et al. 2014), allowing us to locate the entanglement of various stakeholders in the wider domain of the sustainability of AI. In the Discussion, we draw on the concept of entanglement theory from

[1] Within the Environmental Group: Goals 13 (Climate Action), 14 (Life Below Water), 15 (Life on Land) (nat 2015).

the adjacent field of human-computer interaction (HCI) as a way to make sense of complex systems, where the number of variables and actors are both inseparable and numerous (Frauenberger 2019). Combining entanglement with stakeholder methods highlights the vital and multiple roles all computing professionals have regarding the environmental and, subsequently, climate impacts of AI. Through this process, we located five stakeholder groups (industry, practitioners, regulatory, advocacy, and the general public). In part, we identified these groups by analyzing a number of scholarly articles, popular news reports, and essays – and in doing so, we noticed that many of these works were written by a member of a given stakeholder group for other members of their same group. A goal of this work is to surface the overlap between groups, in part to demonstrate how collaboration and reflective action can be engendered.

Related Work

This work is predicated on two existing areas of study: first, the environmental and climate impacts of AI, and second, stakeholder analysis and mapping as a means of sustainable practice within both Environmental Studies (ES) and Software Engineering (SWE).

Environmental and Climate Impacts of AI

There are three primary environmental concerns of AI: energy consumption, water usage, and e-waste (electronic waste) in the form of exhausted hardware infrastructure. These concerns comprise significant segments in the larger lifecycle of AI systems, though we do not focus on hardware development, as infrastructure development (e.g., producing a GPU) is not necessarily specific to AI systems (OECD 2022). Despite much of the work happening "in the cloud," the physical nexus of both computation and impact is a network of data centers (Hogan 2015b). We additionally discuss the social and cultural implications of natural resource usage, as well as *red AI* (Schwartz et al. 2019), or the fixation on model accuracy that requires exponential energy consumption.

Energy Consumption The major contributing factor to AI's environmental footprint is its significant energy consumption for training and operation. Patternson et al. (Patterson et al. 2022) refined earlier estimations from the commonly referenced study by Strubell, Ganesh, and McCallum (2019) regarding the carbon emissions of large model training; they found that training GPT-3 consumed 1,287 MWh of energy and generated 552 tons of CO_2 equivalent emissions (Patterson et al. 2022). This is roughly the same amount of energy needed to power 120 U.S. households for a single year, with emissions roughly equivalent to that of 110 gas-powered cars in the U.S. in a year (Kojola 2023).

As models become larger and more complex, they also need even more energy to run. Consider GPT-3—the 2021 version of which had 175 billion parameters—compared to GPT-1's relatively mere 117 million in 2018 (Griffith 2023). The type of task being performed also plays a role in the emissions produced. Luccioni et al. found that classification tasks were the lowest energy consumers (0.002 to 0.007 kWh for 1,000 references), with generative tasks consuming more (0.05 kWh) and multimodal tasks taking even more (0.06-2.9 kWh) (Luccioni, Jernite, and Strubell 2023). Image tasks are unsurprisingly more energy intensive than text tasks due to their higher-dimensional data; to put it into perspective, if charging a smartphone requires about 0.012 kWh of energy, then for every 1,000 prompts, the text generation model charges 3.5 phones (0.042 kWh) while the image generation model charges 112.5 phones (1.35 kWh) (Luccioni, Jernite, and Strubell 2023).

Water Consumption Energy consumption is not the only environmental concern related to AI. Models are typically trained and run at data centers, often using vast amounts of water for cooling servers (Hogan 2015a). For example, Google's data centers used 4.3 billion gallons of water in 2021, with its average data center requiring about 450,000 gallons of water per day (Hölzle 2022). By 2022, their water consumption increased by 20% to 5.2 billion gallons (Langley 2023). Moving from companies to individual AI products, Li et al. found that training GPT-3 could require about 700,000 liters of water, which is roughly equivalent to the annual water consumption of 2,200 average Americans (Kojola 2023), and requires approximately 500 mL of water for every 10-50 user prompt responses (Li et al. 2023).

Can't we just...?

- *move data center locations?* A handful of nations have taken advantage of natural resources and geography to create low-emission energy sources. For example, Canada and Switzerland rely on hydroelectric power, a renewable energy source (Writer 2023). Nations with "free air cooling" (Jones 2018), such as Iceland, are also prime locations for data centers (Adalbjornsson 2019). Simply moving data centers to these locations, however, poses problems for national data sovereignty efforts, as well as compromising the benefits of transmission speed due to proximity.
- *compensate with carbon offsets?* Promises of carbon neutrality or negativity—of which two-thirds of the world's largest companies have agreed to (Pearson 2023)—are often predicated on the carbon offset system, in which purchased offsets are translated into environmental projects, such as planting forests or investing in renewable energy projects. However, offsets are controversial. For example, among the offsets purchased by the world's fifty largest companies, only 8% were used for projects that removed carbon from the atmosphere (mostly planting trees), while the majority focused on reducing carbon (Pearson 2023). Further, they can be used to suggest that fossil fuels are "carbon neutral" through the purchase of carbon offsets (Gardner, Adomaitis, and Nickel 2021). Additionally, in some cases, offsets (as they manifest land use) challenge traditional property ownership laws, and may not even be used for the intended purpose (Song and Moura 2019).

Implications of Natural Resource Use There are several factors involved in picking the site of a data center, including sensitivity to natural and human-made hazards, availability of utilities (including energy sources and water), and proximity to primary audience(s) (Covas, Silva, and Dias 2013). As a result, environmental sustainability has not historically been the primary rationale behind picking data center locations (Shehabi et al. 2011).

The impact of data center resource consumption falls disproportionately on surrounding communities. A popular argument by local officials and corporate executives has been the creation of jobs in the region chosen; however, this narrative of new job opportunities for residents has been disputed (Lenio 2015; Hardy 2016; Rayome 2016; Mayer 2023). Local communities appear to be more frequently negatively impacted by resource usage from data centers, such as in West Des Moines, Iowa, where training for GPT-4 occurred during a three-year drought (Adarlo 2023). In addition to running counter to environmental justice (Keller, Donoghoe, and Perry 2024) and potentially creating noise pollution (Gonzalez Monserrate 2022), the placement of data centers can violate local and ancestral approaches to caring for nature and natural resources (Lehuedé 2024) and displace Indigenous groups (Lehuedé 2022).

Beyond local natural resource consumption, data centers, necessary for the training and delivery of AI systems, produce vast amounts of e-waste. Gonzalez Monserrate terms this "immoral waste," given that much of the exhausted hardware is then exported to countries with weak labor protections, resulting in workers being exposed to the hazardous materials embedded in hardware (Gonzalez Monserrate 2022).

Exacerbating Implications: Red AI Given what we know about the environmental, social, and cultural implications of AI system development and deployment, it is a concern that these effects are only becoming more intensified. The so-called *Red AI* trend, in which accuracy (i.e., correct identification or generation) is prioritized over efficiency (i.e., resource usage), is becoming increasingly common (Dhar 2020). Schwartz et al. analyzed 60 papers from AI conferences, finding 90% of ACL 2018 papers to prioritize accuracy over efficiency, compared to 80% and 75% at the same year's NeurIPS and CCVPR conferences, respectively (Schwartz et al. 2019).

Red AI is concerning due to the logarithmic relationship between accuracy and efficiency – a linear improvement in accuracy requires that a model become exponentially larger (Mill, Garn, and Ryman-Tubb 2022). Brownlee et al. found that a 1% increase in accuracy increases energy costs by 30-50% (Brownlee et al. 2021). Mill et al. term this the *accuracy-energy paradox* (Mill, Garn, and Ryman-Tubb 2022), as studied by Spillo et al. in the case of recommendation algorithms (Spillo et al. 2023).

Stakeholder Mapping and Entanglement

Stakeholder identification, analysis, and mapping methods have been used in both Environmental Studies and Software Engineering to help identify how necessary change can occur. Following Rakova and Dobbe's argument that the fields of environmental studies and AI are inextricably linked (Rakova and Dobbe 2023), we employ stakeholder methods, given its precedent in both environmental studies and software engineering.

For *stakeholder identification*, Majumdar (2013) defines stakeholders in a given context as any individual, group, or organization that can affect (or is affected by) a topic. In software engineering, McGuire et al. demonstrate that the stakeholder model used in the software development process influences the final software product and that sustainability is a participatory process, requiring the involvement of many stakeholders; stakeholder concerns impact the software development process, adding necessary but otherwise overlooked dimensions (McGuire et al. 2023). Penzenstadler, Femmer, and Richardson (2013) argue that without identifying stakeholders, a sustainability objective in software systems will not receive support and attention to be effectively implemented. Further, when stakeholders are identified, Penzenstadler et al. argue that software engineers will be motivated to invest in sustainability since they will see that sustainability aligns with the interests and objectives of their stakeholders, therefore viewing sustainability as complementary rather than a sacrifice to their goals.

In *stakeholder analysis*, stakeholders are arranged based on their involvement, interest, and influence on a domain (Majumdar 2013). Then, *stakeholder mapping* can take place, in which a visual representation is made of relationships between stakeholders in order to help all parties navigate the social complexity. In Environmental Studies, complex social, political, and cultural dimensions can be put into conversation; for example, Raum uses stakeholder mapping to push for sustainable forestry practices that are more effective, equitable, and sustainable (Raum 2018). On the Software Engineering side, stakeholder mapping has been used for ethics and sustainability education for apprenticing software engineers. For example, Ayoola (2023) developed a curriculum for stakeholder mapping and personas as tools to guide students in socially sustainable software development.

Traditionally, in stakeholder analysis and mapping, the third step is *engagement*, or finding ways to convene representatives of all groups (Lelea et al. 2014; Leventon et al. 2016). However, given the scope and scale of the problems presented by the impacts of the lack of sustainability of AI, we argue that straightforward engagement may not be plausible. Instead, we argue that *entanglement* is a better substitute for engagement. Where *engagement* suggests a finite, scheduled site of interaction, *entanglement* is understood in computing as an irrevocable state of interconnectedness between actors, whether human or technical infrastructure (Frauenberger 2019). By suggesting a site-specific swap of engagement and entanglement, we do not intend to make an argument about the (in)appropriateness of engagement as a solution-oriented step in stakeholder methods more broadly. Our point in turning to entanglement is to reconcile with the unusually fast-evolving nature of the domain (AI development, practice, and infrastructure) and the large-scale ecology that it both produces and inhabits.

Methods

In performing the informal literature review, stakeholder analysis and mapping described in this work, we did so as relative "insiders" to the field of AI. The first author is an undergraduate student with expertise in state-of-the-art ML, and the second, third, and fourth authors are researchers in systems design and AI ethics.

The methods of this paper are two-part. First, we performed an informal literature review to gauge who the stakeholders might be. Then, we performed an analysis and mapping of the identified stakeholder groups.

Informal Literature Review

Our multi-method literature review began in August 2023 and continued through May 2024. Unlike a traditional, exhaustive literature review using pre-specified queries, we used a number of approaches to ground our stakeholder groups, namely 1) discussing with peers in AI and ML disciplines, gathering both suggested ideas and sources to look into, including specific article or paper recommendations and 2) reviewing articles and papers we encountered on social media (e.g., the HCI Social Mastodon) and a variety of newspapers and magazines, covering both tech-focused topics and general current events (e.g., The New York Times, Rest of World, Wired, MIT Tech Review). In addition, we searched several variations of "sustainability" in the ACM Digital Library and IEEE Xplore.

In total, we read more than one hundred research papers, newspaper and magazine articles, and project descriptions. The full list with categories is accessible at the following link: https://annabelrothschild.com/aies-163-appendix/.

Performing Stakeholder Analysis and Mapping

From this broad survey of concerns related to the sustainability of AI system development and execution, we began identifying groups of stakeholders. Our process was predicated on a guide by Lelea et al. (2014), which helped us employ a transdisciplinary approach. Following their stakeholder analysis approach, we completed the following steps:

- Step 1 (selecting human activity system focus): We focused on the sustainability of AI. We consider sustainability to be the impact of AI on both the environment and climate, which is why we focus on energy use and consequent pollution, as well as water usage.
- Step 2 (conducting actor identification and initial characterization): We conducted a broad literature survey with weekly discussions between the first and second authors over the course of roughly six months (August 2023 - January 2024). Each week, the two authors reported back on the articles they read and discussed what they were seeing, in terms of identifying the author, intended audience, and overarching point for each piece. Additionally, we discussed with colleagues in AI fields and gathered their impressions of specific tools or systems for measuring or improving environmental impact, as well as their broader perspectives on the sustainability of AI. From these weekly review meetings, the first author developed a preliminary list of stakeholders, which the first and second authors discussed and iterated through two more times. Changes to the list involved categorizing some stakeholders into groups (e.g., condensing governmental organizations, professional organizations, and nonprofit policymakers into a "policymakers" group). The final version is that used in this paper, with five stakeholder groups.
- Step 3 (formulating specific issue to be addressed): reflecting on our goals for the project, all authors, in multiple discussions, agreed that the priority for this project was highlighting the roles and responsibilities of various stakeholder groups, with the emphasis being on how negotiations of those respective stakeholder priorities and challenges could be a basis for cross-group collaboration.
- Step 4 (conducting stakeholder analysis): employing the stakeholder list, we spent the remaining months of the research project (February - May 2024) analyzing the concerns and priorities of the identified stakeholder groups.

In producing the resulting stakeholder list, we don't intend for it to be seen as a final or static compendium; rather, it should be seen as a list of major players at this point in time, which will no doubt fluctuate as does the ever-evolving field of AI does so. Additionally, though this method may not be "rigorous" in the traditional sense of a literature review (e.g., bounded by pre-defined queries, sourced from a set number of databases), we believe it is the most appropriate method, given that there has been relatively little written about the sustainability of AI, in comparison to the amount of literature on AI techniques and innovations. Further, by crowd-sourcing relevant papers and ideas, we give an overview of relevant stakeholders grounded not only in traditional academic venues, but also in popular media and governmental and NGO reports, which are more indicative of the concerns of non-AI practitioner stakeholders.

Identifying Stakeholders & Their Priorities and Challenges

There are five groups of stakeholders in the sustainability of AI, as yielded by our literature review: industry, practitioners, regulatory, advocacy, and the general public. In this section, we describe these groups, along with their priorities and challenges, and provide an overview in Table 1.

1 - Industry

Perhaps the most associated stakeholder group with AI is the industry group, which includes infrastructure companies (i.e., those that provide physical and virtual resources for AI operations, including data centers, cloud service providers, and hardware suppliers) and AI companies (i.e., those that produce commercial AI products, such as OpenAI). Given the size and complexity of this stakeholder group, we discuss each in turn.

Infrastructure Companies Infrastructure companies strive to be the preferred choice for AI companies by providing quality services at competitive rates. To appeal to customers, infrastructure companies are becoming increasingly concerned about their environmental footprint, aiming to align with the sustainability goals of their clients.

While infrastructure companies share in the concern for environmental issues, they face unique challenges as they

Stakeholder group	Stakeholder examples	Key concerns	Key challenges
Industry	1. AI companies, 2. Infrastructure companies: data centers, cloud service providers, hardware suppliers	Competitive advantage, innovation, regulatory compliance, reputation	Business implications of disclosing environmental footprint, need for alternative energy sources and scientific innovations for sustainability; AI companies: lack of transparency among infrastructure companies
Practitioners	1. Researchers (industry, academic), 2. System developers	Innovation, funding, ethical considerations	Red AI trend, lack of awareness and collective effort
Regulatory	1. Governments, 2. International organizations	Environmental standards and guidelines, environmental legislation, regulatory compliance, public welfare, economic development	Lack of transparency in environmental footprint, political opposition, resistance from industry, slow pace of government
Advocacy	1. Environmental organizations, 2. Environmental activists	Awareness, advocacy, environmental legislation, transparency in environmental footprint	Lack of transparency in environmental footprint
General public	1. Local communities (affected communities), 2. More broadly marginalized communities	Environmental impact, representation, equity, economic development	Lack of power and resources to influence decision-making processes

Table 1: Summary of stakeholder groups

contribute to the environmental footprint in different ways. Each infrastructure company, along with its challenges, is described below.

Data Centers. As sites where AI model training and running usually occur, data centers are significant contributors to the environmental footprint of AI. They face challenges due to their high energy and resource requirements, and while these issues can be mitigated through innovation, doing so often involves significant costs and complexities.

Data centers consume over 200 TWh each year, a larger amount than the annual energy consumption of some countries, such as Iran (Jones 2018). However, they are not used efficiently. Jones argues that many data centers have hardware that is outdated and hard to optimize (Jones 2018), pointing out Taylor and Koomey's research on the prevalence of "zombie" servers, or ones consuming power but providing no work[2].

Notably, the primary energy and resource consumption in data centers stems from cooling systems, such as pumps, chillers, and cooling towers (Zhang et al. 2021; Evans and Gao 2016). In 2014, cooling towers in the US, which operate by evaporating water to cool the facilities, used about 100 billion liters of water (Ristic, Madani, and Makuch 2015).

There are applications for sustainable AI at data centers; for example, Google's DeepMind has used machine learning to optimize the cooling systems in Google's data centers, achieving a 40% reduction in the energy used for cooling (Evans and Gao 2016). In addition, the waste heat from data centers can be reused, as demonstrated by IBM's data center in Switzerland, which uses excess heat to warm a nearby swimming pool, or the Condorcet data center in Paris, which directs heat to a research site studying the effects of high temperatures on vegetation (Jones 2018).

Cloud Service Providers. Cloud service providers operate data centers distributed across different regions to provide services over the Internet. Facing pressure to disclose their carbon emissions, cloud service providers seem concerned with the business implications of these disclosures (Issa, Chang, and Issa 2010). While they can employ a host of methods to reduce emissions, these efforts require investment (Balasooriya, Wibowo, and Wells 2016; Kumar and Buyya 2012).

Hardware Suppliers. Hardware also plays a significant role in the environmental footprint of AI. A challenge faced by hardware suppliers is the sustainable management of their products throughout their entire life cycle.

The production of hardware starts with the extraction of raw resources needed for electronics, a process that often involves burning fossil fuels and releasing carbon emissions due to mining activities. The manufacturing and transportation phases that follow are also energy-intensive; consider Williams' study of the environmental impacts of microchip manufacturing (Williams 2004). Manufacturing

[2] https://blog.anthesisgroup.com/zombie-servers-hunting-lost-capital

hardware requires high temperatures and controlled environments, which are maintained using high-energy machines and equipment. Distributing hardware, whether by ground, sea, or air, relies on fossil fuels, and once in use, the hardware requires electricity until it is eventually disposed of at "end of life" as e-waste. E-waste contains toxic chemicals, resulting in other environmental impacts like polluting air, water, land, and food around nearby communities (Miller 2022; Robinson 2009). Therefore, the entire lifecycle of hardware must be considered when addressing its environmental impact. The economic concerns and priorities of hardware suppliers, however, may be in opposition to building chips with a longer lifespan.

Shared Challenges. Several challenges faced by infrastructure companies are shared. For example, the rapid increase in energy needed to power AI may necessitate the innovation of new energy sources, according to OpenAI CEO Sam Altman, who advocates for the development of sustainable alternatives (Tangermann 2024).

AI Companies AI companies need to maintain competitive advantages while also balancing environmental responsibilities. By investing in sustainable practices, we argue, they can not only satisfy their regulatory compliance, and also sustain and improve their public image. Prioritizing sustainability is vital for enhancing their reputation among customers, investors, and the public.

It is important to recognize that AI companies aiming to reduce their emissions might face obstacles due to current limitations. Even if AI companies are required to be transparent, adhering to this regulation may be difficult because the infrastructure services they depend on lack transparency themselves (Adams 2023). All businesses, including those providing infrastructure, must operate with transparency.

2 - Practitioners

The practitioners group includes AI and ML researchers, in addition to professionals involved in the development, deployment, and management of AI systems. The distinction between practitioners and the industry group is that practitioners use and create AI systems, whereas industry members are organizations that provide the infrastructure of AI systems. We focus, in particular, on AI and ML researchers, particularly those in academic or industry research, as well as enterprise organizations. This group is concerned about advancing the state of the art in their field, while upholding their ethical responsibilities.

A key challenge for this group is the trend of *red AI*, or prioritizing accuracy over efficiency (Van Wynsberghe 2021). Reflecting on the relationship between accuracy and energy savings, referred to as the accuracy-energy paradox (Mill, Garn, and Ryman-Tubb 2022), Spillo et al. argue that algorithmic evaluation should also take into account the environmental impact of AI systems (Spillo et al. 2023).

There has been progress in efforts among practitioners to integrate environmental considerations into algorithmic evaluation with the creation of transparency tools that provide standardized ways for tracking carbon emissions of AI systems, including carbon trackers (Henderson et al. 2020; Lacoste et al. 2019), sustainability checklists (Tornede et al. 2023), and energy leaderboards (Henderson et al. 2020). However, the adoption of these tools has suffered from a lack of practitioner awareness (Hershcovich et al. 2022).

Other avenues for lessening the environmental impact include performing smarter experiments. Lacoste et al. argue that more efficient algorithms (e.g., random search over grid search) can save energy while providing better outcomes and that responsible scholarly practice (e.g., extensive literature review, careful system design, thorough debugging) can help prevent chances of failed experiments, in turn decreasing experimental cycles (Lacoste et al. 2019). Optimization techniques can make models more efficient, decreasing the computational power and, subsequently, the energy consumption of the hardware running them. Using pre-trained models and fine-tuning them can prevent researchers from having to train models from scratch. Methods like compression, quantization, pruning, distillation, and feature reduction can all decrease the training time needed for models (Kirkpatrick 2023). Finally, an optimized framework for federated learning called AdaFL has been shown to enhance FedNLP by reducing training speed and computational costs, thus reducing energy consumption (Cai et al. 2023).

Collaboration is another avenue practitioners can take to achieve better energy efficiency. One way is by sharing pre-trained models, which can help avoid duplicate efforts and conserve both energy and costs for training (Sánchez-Pi and Martí 2021). Another way is by increasing reproducibility; Mill et al. write about the reproducibility checklist, which some major AI conferences have adopted, to promote the sharing of code and dataset for full reproducibility (Mill, Garn, and Ryman-Tubb 2022). The ML community has been neglecting reproducibility, Gunderson & Kjensmo argue – in their survey of 400 research papers, none were fully reproducible, with only 6% sharing code and 30% sharing training data (Gundersen and Kjensmo 2018). This lack of sharing can cause other researchers to replicate existing models, increasing AI's carbon footprint. Online platforms like "Papers with Code" have been formed by volunteer authors to address this issue, offering access to over 70,000 papers under open license[3]. Finally, Tornede calls for the publication of unsuccessful experiment results to prevent others from repeating the same mistake, saving both energy and computational resources (Tornede et al. 2023).

An important final avenue is the strategic selection of hardware. The first approach is the selection of more energy-efficient hardware, with consideration to the e-waste generated by the disposal of old hardware. Patterson found that optimized ML hardware like tensor processing units (TPUs) and advanced graphics processing units (GPUs) are between 2 to 5 times more energy-efficient than general-purpose processors (Patterson et al. 2022). For example, TPU 4 has been recognized for its improved efficiency, being 2.4 times more efficient than TPU 2 (Sánchez-Pi and Martí 2021). Lacoste et al. also found that GPUs were 10 times more efficient than CPUs, showing the importance of carefully selecting the right hardware for specific tasks (Lacoste et al. 2019).

[3]https://paperswithcode.com/

Selecting the most optimal processors is crucial as they also influence the efficiency of data centers (Jones 2018).

Another approach is to use less electricity-consuming hardware alternatives. Although advanced hardware like GPUs and TPUs are necessary during the training phase, less electricity-consuming hardware alternatives, such as field programmable gate arrays (FPGAs) and application-specific instruction-set processors (ASIPs), can be chosen for usage (Schwartz et al. 2019). There is also an emerging area of low-precision computing, where using a lower bit-width to represent numbers can lead to faster and more space-efficient computations (Schwartz et al. 2019).

3 - Regulatory

Policymakers are responsible for establishing and enforcing regulations governing AI, and they are concerned about promoting public welfare and meeting international standards.

A current challenge for policymakers is the lack of regulations mandating transparency in carbon emissions throughout the entire life cycle of AI. Due to a lack of such regulations, businesses are not transparent about their emissions (Kirkpatrick 2023), and they are lobbying against regulations (Van Wynsberghe 2021). However, transparency in such information is needed to assess the true environmental impact of AI. For example, in an attempt to estimate the carbon footprint of ChatGPT, Ludvigsen estimated the unknown variables but noted this was difficult due to limited available information (Ludvigsen 2023). Missing was information about emissions during the data preparation and training process, the hardware used (which impacts power consumption), and the region where the model operated (namely, the data center location), along with the carbon intensity of the electricity used there. Additionally, transparency in the hardware supply chain is also important. Luccioni et al. faced similar difficulties in calculating the carbon footprint of BLOOM due to insufficient data regarding the hardware used and emissions from supply chains, including the large quantities of chemicals and minerals involved in emissions from the transportation process (Luccioni, Viguier, and Ligozat 2022).

A few efforts have been made to push regulation, specifically regarding the environmental impacts of AI. The European AI Act[4] considered mandating transparency around carbon emissions in 2021, even though this proposal did not make it to the final bill (Kirkpatrick 2023). Moreover, California Governor Gavin Newsom signed climate disclosure laws for the state, mandating transparency about carbon emissions and climate-related financial risks for companies, and these laws could impact over 10,000 companies, including numerous AI-intensive firms in Silicon Valley (Tangermann 2024). Given the scope and scale of the implications of unsustainable AI energy and water consumption, there is ample room for regulators to act.

However, these proposed regulations faced strong resistance from political and industry opposition. The public has been concerned with potential economic impacts as well as interest groups and corporations concerned with avoiding public scrutiny resist transparency requirements. This opposition presents a challenge to successful implementation of proposed regulations.

The slow pace of the government in responding to technological advancements presents yet another challenge in implementing regulations. To mitigate this challenge, Vinuesa et al. argue that the government should prioritize establishing policy and legislation frameworks for long-term guidance regarding new technologies, followed by regulatory oversight – the authors underscore that policymakers should have sufficient domain understanding in order to create effective policies (Vinuesa et al. 2020).

Beyond direct regulation, alternatives like tax incentives and certification systems can be considered. Tax incentives can be provided to cloud service providers for building data centers in regions that use cleaner energy, because carbon emissions vary by region, as suggested by Luccioni (Luccioni, Jernite, and Strubell 2023). In addition, certification systems similar to those used in sustainability can be adopted into the machine learning community; these certifications can verify responsible AI practices similar to the way sustainable certifications verify responsible environmental practices, encouraging best industry practices and increasing trust in AI systems (Matus and Veale 2022).

4 - Advocacy

Critics and advocates for impacted groups push the government to regulate the environmental footprint of AI. They range from formal policy organizations to artist collectives. For example, consider the Biden administration's 2023 Blueprint for an AI Bill of Rights (OSTP 2022), which mentioned sustainability, but only to encourage the use of AI to address climate change (David 2023). Seventeen advocacy groups signed a letter[5] in response, highlighting the increasing energy requirements of AI and pushing for the Bill to include transparency laws for businesses about their environmental impact (David 2023; Klar 2023).

Members of the advocacy group also help identify insufficient or misaligned practices regarding the carbon footprint of AI. For example, The Institute for Technology in the Public Interest created a report (ParisBurning 2024) against Frontier Climate (a consortium of Big Tech companies established to manage commitments related to carbon removal[6]), protesting Frontier Climate's emphasis on advanced market commitments for carbon removal. The report's authors advocate for collective action against and investment away from Frontier Climate, calling for a more inclusive and effective approach to climate change and transition to addressing the causes of climate change, rather than using carbon removal as a justification for continued fossil fuel use (ParisBurning 2024).

A significant challenge for members of this group is the lack of transparency in emissions that limits data on AI's

[4] https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

[5] Archived on the Climate Action Against Disinformation Website: https://caad.info/wp-content/uploads/2023/12/Climate-Response-to-WH-EO-on-AI.pdf

[6] https://frontierclimate.com/

carbon footprint. As a result, both critics and advocates face difficulties in quantifying and communicating the extent of AI's environmental impact, which hinders efforts to develop effective mitigation strategies.

5 - General Public

The impact of the sustainability of AI (or lack thereof) is felt most acutely by specific populations, particularly those living close to data centers, as discussed in the Related Work. Specifically, members of the general public living or working near data centers may be concerned about the health impacts of the environmental harm caused by AI. This includes air pollution, water quality deterioration, and climate change, which are all threats to public well-being. Additionally, those most impacted by climate change (e.g., climate refugees) are pushed into even more precarious positions due to the energy consumption of AI (Faber and Schlegel 2017).

Bender et al. examine the environmental impact on marginalized communities, highlighting *environmental racism* – that marginalized communities are the first to be impacted by climate change and are the ones to suffer the greatest, when most language models are built for those who are most privileged in society (Bender et al. 2021). At the same time, some companies have tried to address these ethical concerns by pledging to be carbon-neutral. After the Iowa incident, Microsoft claimed that it aimed to be "carbon negative, water positive, and zero waste by 2030" (Adarlo 2023). This pledge is predicated on the notion of carbon neutrality, which, as discussed in the Related Work, is a controversial solution. Another problem to consider regarding data centers is the domination of data centers by American companies. Dobbe & Whittaker note that because these data centers are located across the globe, there is a need for local agency over the infrastructure of AI (Dobbe and Whittaker 2019).

Due to a lack of power and resources, the public typically has limited ability to influence decision-making processes regarding AI's environmental footprint. First, powerful stakeholders like AI companies and the government can prevent ordinary citizens from influencing practices and policies. In addition, individuals usually have limited resources in terms of time and money to engage in advocacy efforts. These constraints make it challenging for the general public to significantly influence decisions about AI's environmental footprint.

Stakeholder Mapping

Based on the stakeholder analysis, the identified stakeholders were mapped using a four-quadrant interest-influence matrix (see Figure 1). This style of matrix is commonly used in stakeholder mapping as a way of identifying the respective levels of interest and influence of various stakeholders (Walker, Bourne, and Shelley 2008). Specifically, the horizontal axis measures stakeholder interest (the degree to which a topic affects the stakeholder), while the vertical axis measures stakeholder influence (the degree to which a stakeholder affects a topic).

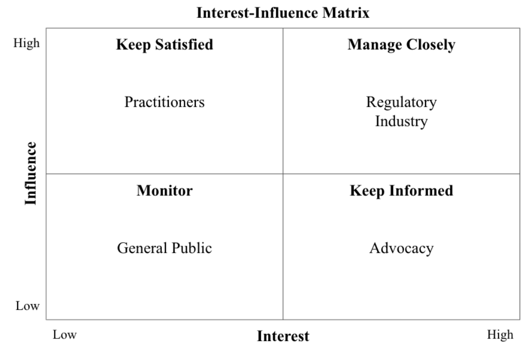

Figure 1: Stakeholder mapping via interest-influence matrix

From the matrix, we can see that the general public is the most overlooked stakeholder group. This is a result of their low influence (namely, they must go through other stakeholder groups, such as advocacy and regulatory channels, to enact change) and low interest. Here, low interest does not suggest negligence; instead, aside from communities that are acutely impacted by factors like physical proximity to data centers, there is not a great deal of public awareness, as demonstrated by a slew of popular news media pieces introducing the impacts to a broader audience (Erdenesanaa 2023; Kishan and Saul 2024; Langley 2023).

However, the general public is affected by the environmental footprint of AI, and we hope that this stakeholder mapping will ensure they are considered in discussions about the sustainability of AI. Additionally, the advocacy group currently does not have significant influence, but it is highly interested. This group has the potential to bridge the gap between the general public and other stakeholder groups (i.e., regulatory, industry, practitioners), enabling a more collaborative approach to sustainable AI development.

Finally, regulatory and industry groups, along with practitioners, have the most influence, as they make decisions that directly affect the environmental footprint of AI. The relationship between the regulatory and industry groups varies by regulatory jurisdiction; for example, in the United States, general AI regulation at the federal level has been heavily influenced by the industry, as seen by the composition of the new Department of Homeland Security's AI Safety and Security Board (DHS 2024; Shepardson 2024).

Further, regulatory and industry groups have high interest because the former is responsible for regulating the environmental impact, while the latter must adhere to the standards and guidelines set by the former. Practitioners, on the other hand, might have lower interest as individuals because they may feel their individual contributions are negligible compared to larger organizational efforts, or what Widder and Nafus (2023) describe as *dislocated accountabilities*. However, because practitioners have the power to have significant influence, due to the implications of even the smallest choices with regard to AI system development and implementation, it is important to raise awareness among them and encourage greater involvement and interest in sustainable practices.

Discussion: Making Sense of Our Stake(s)

This paper has two intended purposes. First, to be a resource for computing professionals interested in the sustainability of AI, but unsure where to begin. We hope our stakeholder analysis and mapping will support interested professionals in seeing opportunities for action toward more sustainable practices by identifying the stakeholder group(s) to which they belong. Second, to begin a conversation around the sustainability of AI at higher organizational levels, such as academic conferences and funding bodies, as called for by Thelisson (2018). Critically, we do not suggest a blanket policy of our own creation; the goal is for policy entities to recognize what is at stake and develop policy with regard to their local community, ideally by considering the needs of other stakeholders, such as marginalized communities.

As noted in Related Work, the third step of stakeholder methods is usually engagement – convening representatives of the stakeholder groups (Lelea et al. 2014; Leventon et al. 2016) – but we argue that entanglement is perhaps a better framing than engagement. Entanglement theory has shown up in the adjacent field of human-computer interaction (HCI) as a method for understanding complex HCI systems, where numerous variables and actors are intertwined and cannot be separated (Frauenberger 2019). Karen Barad's notion of *agential realism*, in which both the materiality of products and the processes by which those products come into being, traces inspiration to a grand example of complex systems (quantum physics) (Barad 2006). What entanglement offers us – and which engagement traditionally does not – is acknowledgment that this problem, or the sustainability of AI, is neither finite (no one convening of stakeholders will be able to generate a fitting ultimatum) nor fixed. Said differently, this number of stakeholders and materials involved will fluctuate (Lisle 2021), as will our understanding of what is a reasonable resource consumption by an AI system. This is not, however, an intractable problem; individual choices, such as the choice of hardware, do add up over time, given the enormous energy consumption posed by AI development and production at large.

This sense of fluctuation is unlike sustainability in the traditional policy sense; for example, consider contamination levels in drinking water. Once a contaminate is identified and an unacceptable level agreed upon (e.g., lead), policy can be enacted and revisited as necessary. However, what is considered environmentally sustainable will continue to fluctuate, based on both the complexity inherent to the number of agents (human and material) involved, as well as the unabated onward march of developments in AI technology and practice. While lead has a stable definition and popular conception, the sustainability of AI does not, nor will it likely arrive in the near future, despite the immediate and lasting impacts of its resource usage. As a comparison, it took more than fifty years for U.S. policy on asbestos to change (from flexible fire retardant "wonder material" to serious carcinogen), and major scientific method breakthroughs were needed to provide evidence (Barlow et al. 2017; Castleman 2006). AI, however, has an expedited development speed, with major infrastructural shifts occurring at infinitesimal intervals – for example, the transition from ChatGPT-3 to ChatGPT-4 happened in just a few years with a significant uptick in environmental footprint (Griffith 2023).

Consider the case of educators in this regard. At any given time, educators (namely at the university level) may be part of multiple stakeholder groups: they are policymakers within their courses and research programs, setting local regulations for student work like final projects and for broader research agendas. Each decision they make about the materials used and conditions set has politics (or, implications) (Winner 1980). If the professor has a selection of multiple data center offerings for their research group to use as core lab infrastructure, they might choose one based on pricing and geographical proximity, but should also consider the sustainability of their choice (i.e., whether the data center is powered by renewable energy). In both the research and course spaces, they can require proof of scopious literature review to ensure that experiments or final projects are not an unnecessary duplication of energy-consuming processes (e.g., training their own classifier vs. using an existing one).

Both of these scopes also pertain to setting standards for future computing professionals (research affiliates and students) and have a unique chance to imprint best practices regarding the sustainability of AI. For example, a professor teaching a course on social computing will reasonably be faced with students wanting to use AI or ML systems in their final course project. Here, the professor can suggest that they employ certain mechanisms (e.g., engaging with energy scorecards and carbon trackers to make sense of their actions) to best practice responsible AI. Likewise, a professor of an AI or ML course can impart the importance of sustainability and environmental resource conservation while also introducing state-of-the-art techniques.

At the same time, this professor may hold a partial appointment with an industry organization, perhaps for fundamental AI research, or hold grants or gifts from such entities. They can use their positions or resources from these industry partners to advocate, where possible, for sustainable development practices, or ensure that resources are responsibly used – for example, by ensuring proper stewardship of old hardware to more environmentally friendly disposal and recycling sites, or prioritizing energy efficiency in the purchase of new hardware.

In another complication, a university-level educator may also hold roles with funding bodies (e.g., U.S. National Science Foundation) or research venues, such as AI or ML conferences. Here, there is an opportunity to enact policy at a forceful level. This work spells out some of the ways in which professionals in such roles can show consideration for other stakeholder groups, such as the general public, and argue for more transparency and sustainability around AI research and practice.

And what about those of us who may not be AI or ML researchers? Why should we care about the sustainability of AI? Besides being members of the general public, whether local or global, even if we don't live or work near a data center, the implications of AI use in our field are contributing to climate change, which in turn will affect us. Said differently, even if we (non-AI/ML-focused computing professionals) do not consider ourselves to be in the practitioner stakeholder group, we are at least in the general public group and should take action based on what we know to be the experiences of our peers (e.g., the resource pressure placed on local communities by data centers). Furthermore, given the ever-pressing calls to use and integrate AI in other domains, including other fields of computing research and practice, it is likely we will be asked to review or otherwise critique work that makes use of AI and ML methods and systems. Thus, it is critical that we are equipped with some level of knowledge that allows us to be thoughtful interlocutors with such work, and, in particular, with regard to its implications for sustainability and climate and/or natural resource impact. Hopefully, this work also provides language and examples to help professionals whose primary research interest or affiliation is not in AI or ML critique these methods when used, with regard to sustainability and resource consumption.

What does acknowledging the sustainability of AI as an entanglement problem get us? Throughout this discussion, there has been not only an emphasis on the need to take action (congruent with most stakeholder engagement instantiations) but also on the multi-faceted human and material elements of this problem space. Specifically, the role of computing hardware and infrastructure, both of which have huge environmental impacts (i.e., graphics cards (Koski 2021; Griffin 2021) and data centers (Siddik, Shehabi, and Marston 2021)). Unlike traditional use of engagement in stakeholder methods, entanglement allows us to incorporate these non-human components as significant participants in the wider problem space, as an example of *intertwined integration*, or when the product of a human-computer interaction cannot be clearly traced to either party, given that the agency of the user and computer are deeply intertwined (Mueller et al. 2023). Further, by combining entanglement with stakeholder methods, we demonstrate that we all, as computing professionals, play a role (whether relative insider or outsider to AI practice) with regard to the environmental and, subsequently, climate impacts of AI.

Limitations

There are two limitations to this work. The first is that it represents the stakeholder groups at a snapshot in time. While we do not anticipate major variance in the immediate future, it is not unlikely that this list will fluctuate over time, particularly as AI systems continue to be the subject of development and their ramifications are increasingly felt by all cross-sections of the global population.

The second is that this work is predicated on a U.S. (and slightly more broadly North American and European) standpoint, as well as an English language one, particularly with regard to popular media reports and examples of local impact. This is not to say that these locales are the only sites of impact, but rather that they are the ones that we have encountered. In part, this standpoint is reflective of the over-representation of large AI industry players being headquartered in the U.S., but it should not be used to infer that the U.S. is the only relevant site for discussion.

Conclusion

In this work, we report on a year-long study of the roles and responsibilities that computing professionals hold with regard to the sustainability of AI systems. Given the immense environmental cost (through natural resource usage and pollution production) of developing and maintaining AI systems, whether or not we are direct designers, consumers, or users of those systems, we should be mindful of their environmental toll. We first employ stakeholder methods – namely identification, analysis, and mapping – to sort stakeholders into five key groups (industry, practitioners, advocacy, regulatory, and general public). We then eschew the traditional final step of stakeholder methods (*engagement*) for *entanglement*. Our rationale for doing so is to highlight the unusual speed with which AI developments occur, as well as the deeply intertwined nature of humans and machines. By turning to entanglement theory, we recognize that all computing professionals have the ability to affect the sustainability of AI and a responsibility to do so. We close with a call for policy and policy bodies that are ongoing, flexible sites, enabling us to keep pace with the ongoing evolution of AI systems and products.

Acknowledgments

We thank our anonymous reviewers for their feedback on this work. This work is supported by NSF grant #1951818 DataWorks: Building Smart Community Capacity. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Author Statements

Ethical Considerations

In conducting this work, our focus was on the impacts of AI on the environment, namely implications for climate change, as a result of natural resource usage and pollution production. This is not to say we are against AI in theory, nor is it a call to cease the use of AI systems and products entirely. Rather, we argue that as a community, we should better understand the implications of the systems we engage with and re-negotiate our interactions accordingly.

Researcher Positionality

This work was motivated by an interest in the sustainability of AI, and we approached it from the standpoint of trying to understand the environmental "costs" (or impacts) of the development and use of such systems. The first author is an undergraduate student with an emphasis on computational ML methods and technologies, the second author is a graduate student researching methods of dataset annotation and dataset production, and the last two authors are faculty members with backgrounds in design and learning sciences. All authors study and/or work in a computing department at a large research university in the U.S.

Adverse Impact

The key contribution of this work is not a computational system, but a way of knowing and understanding an ecology. Thus, our insights are, of course, shaped by our own perspectives. We hope that this work will provide a way to personalize the reader's relationship with the sustainability of AI (if not yet considered). The largest risk we imagine is overlooking an impacted community or stakeholder group; however, we are confident that the larger research community will address such errors with future work.

References

2015. Transforming our world: the 2030 Agenda for Sustainable Development. Technical report, United Nations.

Adalbjornsson, T. 2019. Iceland's data centers are booming—here's why that's a problem. *MIT Technology Review*.

Adams, C. 2023. Responsible Clouds and the Green Web Triangle.

Adarlo, S. 2023. Critics Furious Microsoft Is Training AI by Sucking Up Water During Drought. *Futurism*.

Amoore, L. 2020. *Cloud ethics: algorithms and the attributes of ourselves and others*. Durham: Duke University Press. ISBN 978-1-4780-0778-4 978-1-4780-0831-6.

Ayoola, B. 2023. Evaluation of Stakeholder Mapping and Personas for Sustainable Software Development. In *Proceedings of the 45th International Conference on Software Engineering: Companion Proceedings*, ICSE '23, 270–272. Melbourne, Victoria, Australia: IEEE Press. ISBN 9798350322637.

Balasooriya, P.; Wibowo, S.; and Wells, M. A. 2016. Green Cloud Computing and Economics of the Cloud: Moving towards Sustainable Future. *GSTF Journal on computing*.

Barad, K. 2006. *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning*. Duke University Press. ISBN 978-0-8223-8812-8.

Barlow, C. A.; Sahmel, J.; Paustenbach, D. J.; and Henshaw, J. L. 2017. History of knowledge and evolution of occupational health and regulatory aspects of asbestos exposure science: 1900–1975. *Critical Reviews in Toxicology*, 47(4): 286–316. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/10408444.2016.1258391.

Bender, E. M.; Gebru, T.; McMillan-Major, A.; and Shmitchell, S. 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*, FAccT '21, 610–623. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-8309-7.

Berreby, D. 2024. As Use of A.I. Soars, So Does the Energy and Water It Requires. *Yale E360*.

Brownlee, A. E.; Adair, J.; Haraldsson, S. O.; and Jabbo, J. 2021. Exploring the Accuracy – Energy Trade-off in Machine Learning. In *2021 IEEE/ACM International Workshop on Genetic Improvement (GI)*, 11–18. Madrid, Spain: IEEE. ISBN 978-1-66544-466-8.

Cai, D.; Wu, Y.; Wang, S.; Lin, F. X.; and Xu, M. 2023. Efficient Federated Learning for Modern NLP. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, 1–16. Madrid Spain: ACM. ISBN 978-1-4503-9990-6.

Castleman, B. 2006. Asbestos Products, Hazards, and Regulation. *International Journal of Health Services*, 36(2): 295–307. Publisher: SAGE Publications Inc.

Covas, M. T.; Silva, C. A.; and Dias, L. C. 2013. Multicriteria decision analysis for sustainable data centers location. *International Transactions in Operational Research*, 20(3): 269–299. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1475-3995.2012.00874.x.

David, E. 2023. Climate groups say Biden AI order doesn't address AI's climate impact.

Dhar, P. 2020. The carbon impact of artificial intelligence. *Nature Machine Intelligence*, 2(8): 423–425.

DHS. 2024. Over 20 Technology and Critical Infrastructure Executives, Civil Rights Leaders, Academics, and Policymakers Join New DHS Artificial Intelligence Safety and Security Board to Advance AI's Responsible Development and Deployment | Homeland Security.

Dobbe, R.; and Whittaker, M. 2019. AI and Climate Change: How they're connected, and what we can do about it.

Dwyer, S. C.; and Buckle, J. L. 2009. The Space Between: On Being an Insider-Outsider in Qualitative Research. *International Journal of Qualitative Methods*, 8(1): 54–63. Publisher: SAGE Publications Inc.

Erdenesanaa, D. 2023. A.I. Could Soon Need as Much Electricity as an Entire Country. *The New York Times*.

Evans, R.; and Gao, J. 2016. DeepMind AI Reduces Google Data Centre Cooling Bill by 40%.

Faber, D.; and Schlegel, C. 2017. Give Me Shelter from the Storm: Framing the Climate Refugee Crisis in the Context of Neoliberal Capitalism. *Capitalism Nature Socialism*. Publisher: Routledge.

Frauenberger, C. 2019. Entanglement HCI The Next Wave? *ACM Transactions on Computer-Human Interaction*, 27(1): 2:1–2:27.

Gardner, T.; Adomaitis, N.; and Nickel, R. 2021. Clean crude? Oil firms use offsets to claim green barrels. *Reuters*.

Gonzalez Monserrate, S. 2022. The staggering ecological impacts of computation and the cloud.

Griffin, A. 2021. Waste and Emissions from the Nvidia GPU. University of California - Davis.

Griffith, E. 2023. GPT-4 vs. ChatGPT-3.5: What's the Difference? *PCMAG*.

Gundersen, O. E.; and Kjensmo, S. 2018. State of the Art: Reproducibility in Artificial Intelligence. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1). Number: 1.

Hardy, Q. 2016. Cloud Computing Brings Sprawling Centers, but Few Jobs, to Small Towns. *The New York Times*.

Henderson, P.; Hu, J.; Romoff, J.; Brunskill, E.; Jurafsky, D.; and Pineau, J. 2020. Towards the systematic reporting of the energy and carbon footprints of machine learning. *The Journal of Machine Learning Research*, 21(1): 248:10039–248:10081.

Hershcovich, D.; Webersinke, N.; Kraus, M.; Bingler, J. A.; and Leippold, M. 2022. Towards Climate Awareness in NLP Research. ArXiv:2205.05071 [cs].

Hogan, M. 2015a. Data flows and water woes: The Utah Data Center. *Big Data & Society*, 2(2): 2053951715592429. Publisher: SAGE Publications Ltd.

Hogan, M. 2015b. Facebook Data Storage Centers as the Archive's Underbelly. *Television & New Media*, 16(1): 3–18. Publisher: SAGE Publications.

Hölzle, U. 2022. Our commitment to climate-conscious data center cooling.

Issa, T.; Chang, V.; and Issa, T. 2010. The impact of Cloud Computing and organizational sustainability. In *Annual international conference on Cloud Computing and visrtualization*, 163–169. CCV& GSTF. ISBN 978-981-08-5864-3. Accepted: 2017-01-30T11:24:40Z.

Jones, N. 2018. How to stop data centres from gobbling up the world's electricity. *Nature*, 561(7722): 163–166.

Keller, J.; Donoghoe, M.; and Perry, A. 2024. The US must balance climate justice challenges in the era of artificial intelligence.

Kirkpatrick, K. 2023. The Carbon Footprint of Artificial Intelligence. *Communications of the ACM*, 66(8): 17–19.

Kishan, S.; and Saul, J. 2024. AI Needs So Much Power That Old Coal Plants Are Sticking Around. *Bloomberg.com*.

Klar, R. 2023. Biden pressed to address AI impact on climate change.

Kojola, E. 2023. ChatCO2 – Safeguards Needed For AI's Climate Risks -. Section: News.

Koski, R. 2021. Nvidia Founders' Edition GPU Raw Materials Acquisition and Manufacture. University of California - Davis.

Kumar, S.; and Buyya, R. 2012. Green Cloud Computing and Environmental Sustainability. In *Harnessing Green It*, 315–339. John Wiley & Sons, Ltd. ISBN 978-1-118-30539-3.

Lacoste, A.; Luccioni, A.; Schmidt, V.; and Dandres, T. 2019. Quantifying the Carbon Emissions of Machine Learning.

Langley, H. 2023. Google's water use is soaring. AI is only going to make it worse. *Business Insider*.

Lehuedé, S. 2022. Territories of data: ontological divergences in the growth of data infrastructure. *Tapuya: Latin American Science, Technology and Society*, 2035936.

Lehuedé, S. 2024. An Elemental Ethics for Artificial Intelligence: Water as Resistance Within AI's Value Chain.

Lelea, M. A.; Roba, G. M.; Christinck, A.; and Kaufmann, B. 2014. *Methodologies for stakeholder analysis: for application in transdisciplinary research projects focusing on actors in food supply chains: reload reducing losses adding value*. Witzenhausen: DITSL. ISBN 978-3-945266-00-7.

Lenio, J. 2015. The Mystery Impact of Data Centers on Local Economies Revealed. Section: Property Markets.

Leventon, J.; Fleskens, L.; Claringbould, H.; Schwilch, G.; and Hessel, R. 2016. An applied methodology for stakeholder identification in transdisciplinary research. *Sustainability Science*, 11(5): 763–775.

Li, P.; Yang, J.; Islam, M. A.; and Ren, S. 2023. Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. ArXiv:2304.03271 [cs].

Lisle, D. 2021. A Speculative Lexicon of Entanglement. *Millennium*, 49(3): 435–461. Publisher: SAGE Publications Ltd.

Luccioni, A. S.; Jernite, Y.; and Strubell, E. 2023. Power Hungry Processing: Watts Driving the Cost of AI Deployment?

Luccioni, A. S.; Viguier, S.; and Ligozat, A.-L. 2022. Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model.

Ludvigsen, K. G. A. 2023. The Carbon Footprint of ChatGPT.

Majumdar, S. I. 2013. Thorny Issues of Stakeholder Identification and Prioritization in Requirement Engineering Process. *IOSR Journal of Computer Engineering*, 15(5): 73–78.

Matus, K. J.; and Veale, M. 2022. Certification systems for machine learning: Lessons from sustainability. *Regulation & Governance*, 16(1): 177–196.

Mayer, V. 2023. When do we go from here? Data center infrastructure labor, jobs, and work in economic development time and temporalities. *New Media & Society*, 25(2): 307–323. Publisher: SAGE Publications.

McGuire, S.; Schultz, E.; Ayoola, B.; and Ralph, P. 2023. Sustainability is Stratified: Toward a Better Theory of Sustainable Software Engineering. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*, 1996–2008. ISSN: 1558-1225.

Mill, E.; Garn, W.; and Ryman-Tubb, N. 2022. Managing Sustainability Tensions in Artificial Intelligence: Insights from Paradox Theory. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, 491–498. Oxford United Kingdom: ACM. ISBN 978-1-4503-9247-1.

Miller, A. 2022. The Future of E-Waste Depends on AI and Machine Learning.

Mueller, F. Semertzidis, N.; Andres, J.; Marshall, J.; Benford, S.; Li, X.; Matjeka, L.; and Mehta, Y. 2023. Toward Understanding the Design of Intertwined Human–Computer Integrations. *ACM Transactions on Computer-Human Interaction*, 30(5): 73:1–73:45.

OECD. 2022. Measuring the environmental impacts of artificial intelligence compute and applications. Technical Report 341, OECD Publishing, Paris. Type: doi:https://doi.org/10.1787/7babf571-en.

OSTP. 2022. Blueprint for an AI Bill of Rights.

ParisBurning. 2024. Frontier Climate: a bug report · Issue #147 · stripe/openapi.

Patterson, D.; Gonzalez, J.; Holzle, U.; Le, Q.; Liang, C.; Munguia, L.-M.; Rothchild, D.; So, D. R.; Texier, M.; and Dean, J. 2022. The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink. *Computer*, 55(7): 18–28.

Pearson, J. G., Tom. 2023. Analysis: How some of the world's largest companies rely on carbon offsets to 'reach net-zero'. *Carbon Brief*.

Penzenstadler, B.; Femmer, H.; and Richardson, D. 2013. Who is the advocate? Stakeholders for sustainability. In *2013 2nd International Workshop on Green and Sustainable Software (GREENS)*, 70–77.

Rakova, B.; and Dobbe, R. 2023. Algorithms as Social-Ecological-Technological Systems: an Environmental Justice Lens on Algorithmic Audits. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23, 491. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701924.

Raum, S. 2018. A framework for integrating systematic stakeholder analysis in ecosystem services research: Stakeholder mapping for forest ecosystem services in the UK. *Ecosystem Services*, 29: 170–184.

Rayome, A. D. 2016. Why data centers fail to bring new jobs to small towns.

Ristic, B.; Madani, K.; and Makuch, Z. 2015. The Water Footprint of Data Centers. *Sustainability*, 7(8): 11260–11284. Number: 8 Publisher: Multidisciplinary Digital Publishing Institute.

Robinson, B. H. 2009. E-waste: An assessment of global production and environmental impacts. *Science of The Total Environment*, 408(2): 183–191.

Schwartz, R.; Dodge, J.; Smith, N. A.; and Etzioni, O. 2019. Green AI. ArXiv:1907.10597 [cs, stat].

Shehabi, A.; Masanet, E.; Price, H.; Horvath, A.; and Nazaroff, W. W. 2011. Data center design and location: Consequences for electricity use and greenhouse-gas emissions. *Building and Environment*, 46(5): 990–998.

Shepardson, D. 2024. US Homeland Security names AI safety, security advisory board. *Reuters*.

Siddik, M. A. B.; Shehabi, A.; and Marston, L. 2021. The environmental footprint of data centers in the United States. *Environmental Research Letters*, 16(6): 064017. Publisher: IOP Publishing.

Song, L.; and Moura, P. 2019. An (Even More) Inconvenient Truth.

Spillo, G.; De Filippo, A.; Musto, C.; Milano, M.; and Semeraro, G. 2023. Towards Sustainability-aware Recommender Systems: Analyzing the Trade-off Between Algorithms Performance and Carbon Footprint. In *Proceedings of the 17th ACM Conference on Recommender Systems*, 856–862. Singapore Singapore: ACM. ISBN 9798400702419.

Strubell, E.; Ganesh, A.; and McCallum, A. 2019. Energy and Policy Considerations for Deep Learning in NLP. ArXiv:1906.02243 [cs].

Sánchez-Pi, N.; and Martí, L. 2021. Towards a green AI: evolutionary solutions for an ecologically viable artificial intelligence. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, GECCO '21, 1135–1140. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-8351-6.

Tangermann, V. 2024. Sam Altman Says AI Using Too Much Energy, Will Require Breakthrough Energy Source. *Futurism*.

Thelisson, E. 2018. Towards a Computational Sustainability for AI/ML to Foster Responsibility. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '18, 386–387. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-6012-8.

Tornede, T.; Tornede, A.; Hanselle, J.; Mohr, F.; Wever, M.; and Hüllermeier, E. 2023. Towards Green Automated Machine Learning: Status Quo and Future Directions. *Journal of Artificial Intelligence Research*, 77: 427–457.

Van Wynsberghe, A. 2021. Sustainable AI: AI for sustainability and the sustainability of AI. *AI and Ethics*, 1(3): 213–218.

Vinuesa, R.; Azizpour, H.; Leite, I.; Balaam, M.; Dignum, V.; Domisch, S.; Felländer, A.; Langhans, S. D.; Tegmark, M.; and Fuso Nerini, F. 2020. The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature Communications*, 11(1): 233.

Walker, D. H. T.; Bourne, L. M.; and Shelley, A. 2008. Influence, stakeholder mapping and visualization. *Construction Management and Economics*, 26(6): 645–658. Publisher: Routledge _eprint: https://doi.org/10.1080/01446190701882390.

Widder, D. G.; and Nafus, D. 2023. Dislocated accountabilities in the "AI supply chain": Modularity and developers' notions of responsibility. *Big Data & Society*, 10(1): 20539517231177620. Publisher: SAGE Publications Ltd.

Williams, E. D. 2004. Environmental impacts of microchip manufacture. *Thin Solid Films*, 461(1): 2–6.

Winner, L. 1980. Do Artifacts Have Politics? *Daedalus*, 109(1): 121–136.

Writer, S. 2023. Carbon Footprint Of Training GPT-3 And Large Language Models.

Zhang, Q.; Meng, Z.; Hong, X.; Zhan, Y.; Liu, J.; Dong, J.; Bai, T.; Niu, J.; and Deen, M. J. 2021. A survey on data center cooling systems: Technology, power consumption modeling and control strategy optimization. *Journal of Systems Architecture*, 119: 102253.