Taking (policy) action to enhance the sustainability of AI systems

The SustAIn Project: Synthesis, Critical Reflection and Policy Considerations

Friederike Rohde, Anne Mollen, Andreas Meyer, Josephin Wagner, Gesa Marken, Vivian Frick, Frieder Schmelzle, Lina Engel

sustain

AI and the Challenge of Sustainability

Editor:

|i|ö|w
INSTITUTE FOR ECOLOGICAL
ECONOMY RESEARCH

Institut für ökologische
Wirtschaftsforschung (IÖW) GmbH, gemeinnützig
Institute for Ecological Economy Research
Potsdamer Straße 105
D-10785 Berlin
Tel. +49 – 30 – 884 594-0
Fax +49 – 30 – 882 54 39
E-Mail: mailbox@ioew.de
www.ioew.de

Berlin, January 2024

Lead author contact:
friederike.rohde@ioew.de

In cooperation with:

ALGORITHM
WATCH

AW AlgorithmWatch gGmbH
Linienstr. 13
D-10178 Berlin
E-Mail: info@algorithmwatch.org

DAI-Labor
TU Berlin

Technische Universität Berlin / DAI-Labor
Fakultät IV für Elektrotechnik und Informatik
Sekretariat TEL 14
Ernst-Reuter-Platz 7
D-10587 Berlin
E-Mail: sekretariat@dai-labor.de

Funding:

Federal Ministry
for the Environment, Nature Conservation,
Nuclear Safety and Consumer Protection

Contents

1 Summary

This synthesis report gives insights into the findings, learnings, and policy implications from the SustAIn project regarding a more sustainable development of artificial intelligence (AI) systems. Within the recent debate on AI and sustainability, it is becoming increasingly clear that AI systems create significant risks for society and our planet. On the other hand, it continues to be seen as a strategic technology (Durant et al., 1998) without which we will presumably be unable to understand the complexities with which the present societal and environmental challenges are associated. These two sides of this dichotomy are inseparable and neither the potential of AI nor the dangers and damaging consequences associated with it can be ignored. As part of the SustAIn project, pioneering work has been undertaken in compiling comprehensive indicators that can be used to assess the social, environmental, and economic sustainability of AI systems. The project has shown that, besides organisations that develop and implement AI, a policy approach is needed to foster more sustainable AI systems. With this synthesis report, we want to share the most important insights and results of our work and facilitate and structure the discourse on the sustainability of AI along the AI lifecycle.

2 On the relationship between AI and sustainability

The rise in the number of applications that are referred to as Artificial Intelligence (AI) has led to debates about the impacts those systems might have on society and the planet. To be more precise, we are specifically referring to machine learning (ML) and its subfield deep learning as subfields of AI. We speak of AI *systems* and conceive of AI as complex socio-technical-ecological systems (Rohde et al., 2024) that are associated with multiple interrelated social, environmental, and economic challenges.

The potential associated with AI systems is often considered to be far-reaching and extensive. The fields of application for the methods and technologies that fall under the term "Artificial Intelligence" are essentially unlimited, in line with their role as a general purpose technology (Cockburn et al., 2019). From the finance sector, to health, education, online marketing, the energy sector, or public administration many application possibilities fuel the further implementation of the related systems. At the same time, it is becoming increasingly clear that those systems will probably not only contribute to dealing with some complex issues but also create a whole range of new problems that have to be dealt with. These issues include discrimination through bias, stereotypes, and representational harms (e.g., Bender et al., 2021; Solaiman et al., 2023); the environmental implications of training the systems (Luccioni et al., 2022; Strubell et al., 2019); other ecological impacts, such as CO_2 emissions (Luccioni et al., 2022; Patterson et al., 2022) and the water footprint (Li et al., 2023) of the digital infrastructure needed to maintain the systems (Robbins and van Wynsberghe, 2022); and issues related to market power and infrastructural monopolies (Png, 2022; Widder et al., 2023). We argue that the discussion on the sustainability of AI deserves more accuracy, more nuance, more scrutiny, and more evidence (SustAIn Magazine 1, 2022).

Current debates contain different but interrelated perspectives and approaches to the relationship between AI systems and sustainability. The rise of applications based on ML models has fuelled ambitions to address global challenges related to the Sustainable Development Goals (SDGs), such as health, education, climate change, or water and biodiversity issues. There is an ongoing debate on whether AI is more likely to contribute to or inhibit accomplishing the SDGs targets (Galaz et al., 2021; Sætra, 2021a; Vinuesa et al., 2020). A basic distinction between the debates' perspectives can be made between AI for sustainability and the sustainability of AI (Rohde et al., 2021; van Wynsberghe, 2021). Whereas AI for sustainability asks in which area AI can be applied to support the SDGs (Vinuesa et al., 2020), we seek to explore the sustainability impacts along the lifecycle of AI systems more generally. Ultimately, a discussion on AI for sustainability cannot be based on evidence as long as the sustainability impacts of such systems themselves are not accounted for. Further, only through understanding sustainability impacts can they be addressed and reduced.

AI is being rolled out increasingly widely, and recognition has grown that its impact on society and the planet is becoming increasingly problematic. Relatedly, there is growing structural evidence that a major problem is emerging and that society is becoming dependent on an unsustainable digital infrastructure (Robbins and van Wynsberghe, 2022). The ML-community itself is also raising debates in which it discusses the risks, for example, of foundation models (Bommasani et al., 2022), generative AI (Solaiman et al., 2023), energy consumption (Strubell et al., 2019), or the carbon footprint (Luccioni et al., 2022) and water footprint (Li et al., 2023). The findings of that research make it clear that the precautionary principle must be applied to AI and problematic effects need to be minimised. Currently, there do not appear to be AI systems that can claim to be sustainable since the growing implementation and the hardware, data streams, and digital infrastructure needed to keep AI applications running are reputed to create spillover and induction effects that may impede sustainability outcomes (Robbins and van Wynsberghe, 2022; Sætra, 2021a). We argue that, in the context of the AI Act being negotiated at the European level, there is a need to measure the environmental footprint and take sustainability considerations into account across all levels: from organisational governance, through AI development and application. But beyond such specific legislative attempts to regulate AI systems, further policy approaches are needed. With the end of the SustAIn project, the work on more sustainable AI systems has to start.

3 Sustainability criteria along the AI life cycle

The sustainability of an AI system depends on many decisions taken during its life cycle. Within the SustAIn project, we have done pioneering work in compiling comprehensive indicators that can be used to assess the social, environmental, and economic sustainability of AI systems. For detailed descriptions of each indicator and how we arrived at it see Rohde et al., 2024, 2021 and SustAIn Magazine 2, 2023.

Our criteria and indicators address different phases of an AI system life cycle, which can be divided into conceptualisation, data management, model development, model implementation, and model use and decision making. In addition, we have added 'organisational embedding' to the life cycle phases. With that, we want to clarify that many aspects we consider with our sustainability criteria should be embedded by management into the corporate structure and organisational culture, whereas other indicators have to be addressed on the regulatory level. Out of the criteria set, we developed a digital app-based self-assessment tool that contains the indicators allowing organisations to test the sustainability of any AI systems they have (available only in German).

4 Selected sustainability effects from socio-technical-ecological AI systems

The sustainability of AI systems can be addressed on various levels. In our understanding, AI systems can be described as socio-technical-ecological systems (Ahlborg et al., 2019; Rohde et al., 2024), and we follow an understanding of society, technology, and environment as co-constituted and co-emergent entities. It is the responsibility of individual developers, of organisations, and of regulators to avoid negative effects on sustainability from emerging AI developments.

While, on individual and organisational levels, the SustAIn project has invested in developing awareness-raising tools and guidelines for creating more sustainable AI systems, some crucial sustainability-relevant aspects surpass individual and organisational responsibility and must be addressed as a regulatory concern. For example, one way to mitigate sustainability risks stemming from severe market concentration in the AI industry might be to invest in open data initiatives and data-sharing practices. However, since sharing data might be counterintuitive for organisations from a competitive advantage perspective, a regulatory approach addressing data governance should be seen as pertinent for sustainable AI development. All sustainability indicators identified and developed in the SustAIn project can accordingly be assigned to the individual, organisational or regulatory level, with many relating to more than just one of these levels. Table 1 gives an overview over the sustainability indicators that predominantly require regulatory initiatives.

Resource consumption	Working conditions	Safeguarding competition	Sustainable application
— Monitoring of environmental impacts — Data centre certification — Efficiency metrics for data centres — Monitoring of efficiency metrics — Hardware disposal	— De-skilling and monotony — Fair wages along entire value chain — Analysis of effects on working conditions — Analysis of effects on labour market	— Open data pools — Open AI development — Accessibility for ML models and data — Support for SMEs — Adaptation requirements — Multihoming and compatibility — Inclusivity in application	— Promotion of sustainable consumption patterns — Promotion of sustainable objectives — Promotion of resource-efficient AI development and application

Table 1: Sustainability Indicators that predominantly require regulatory initiatives

CRITERIA

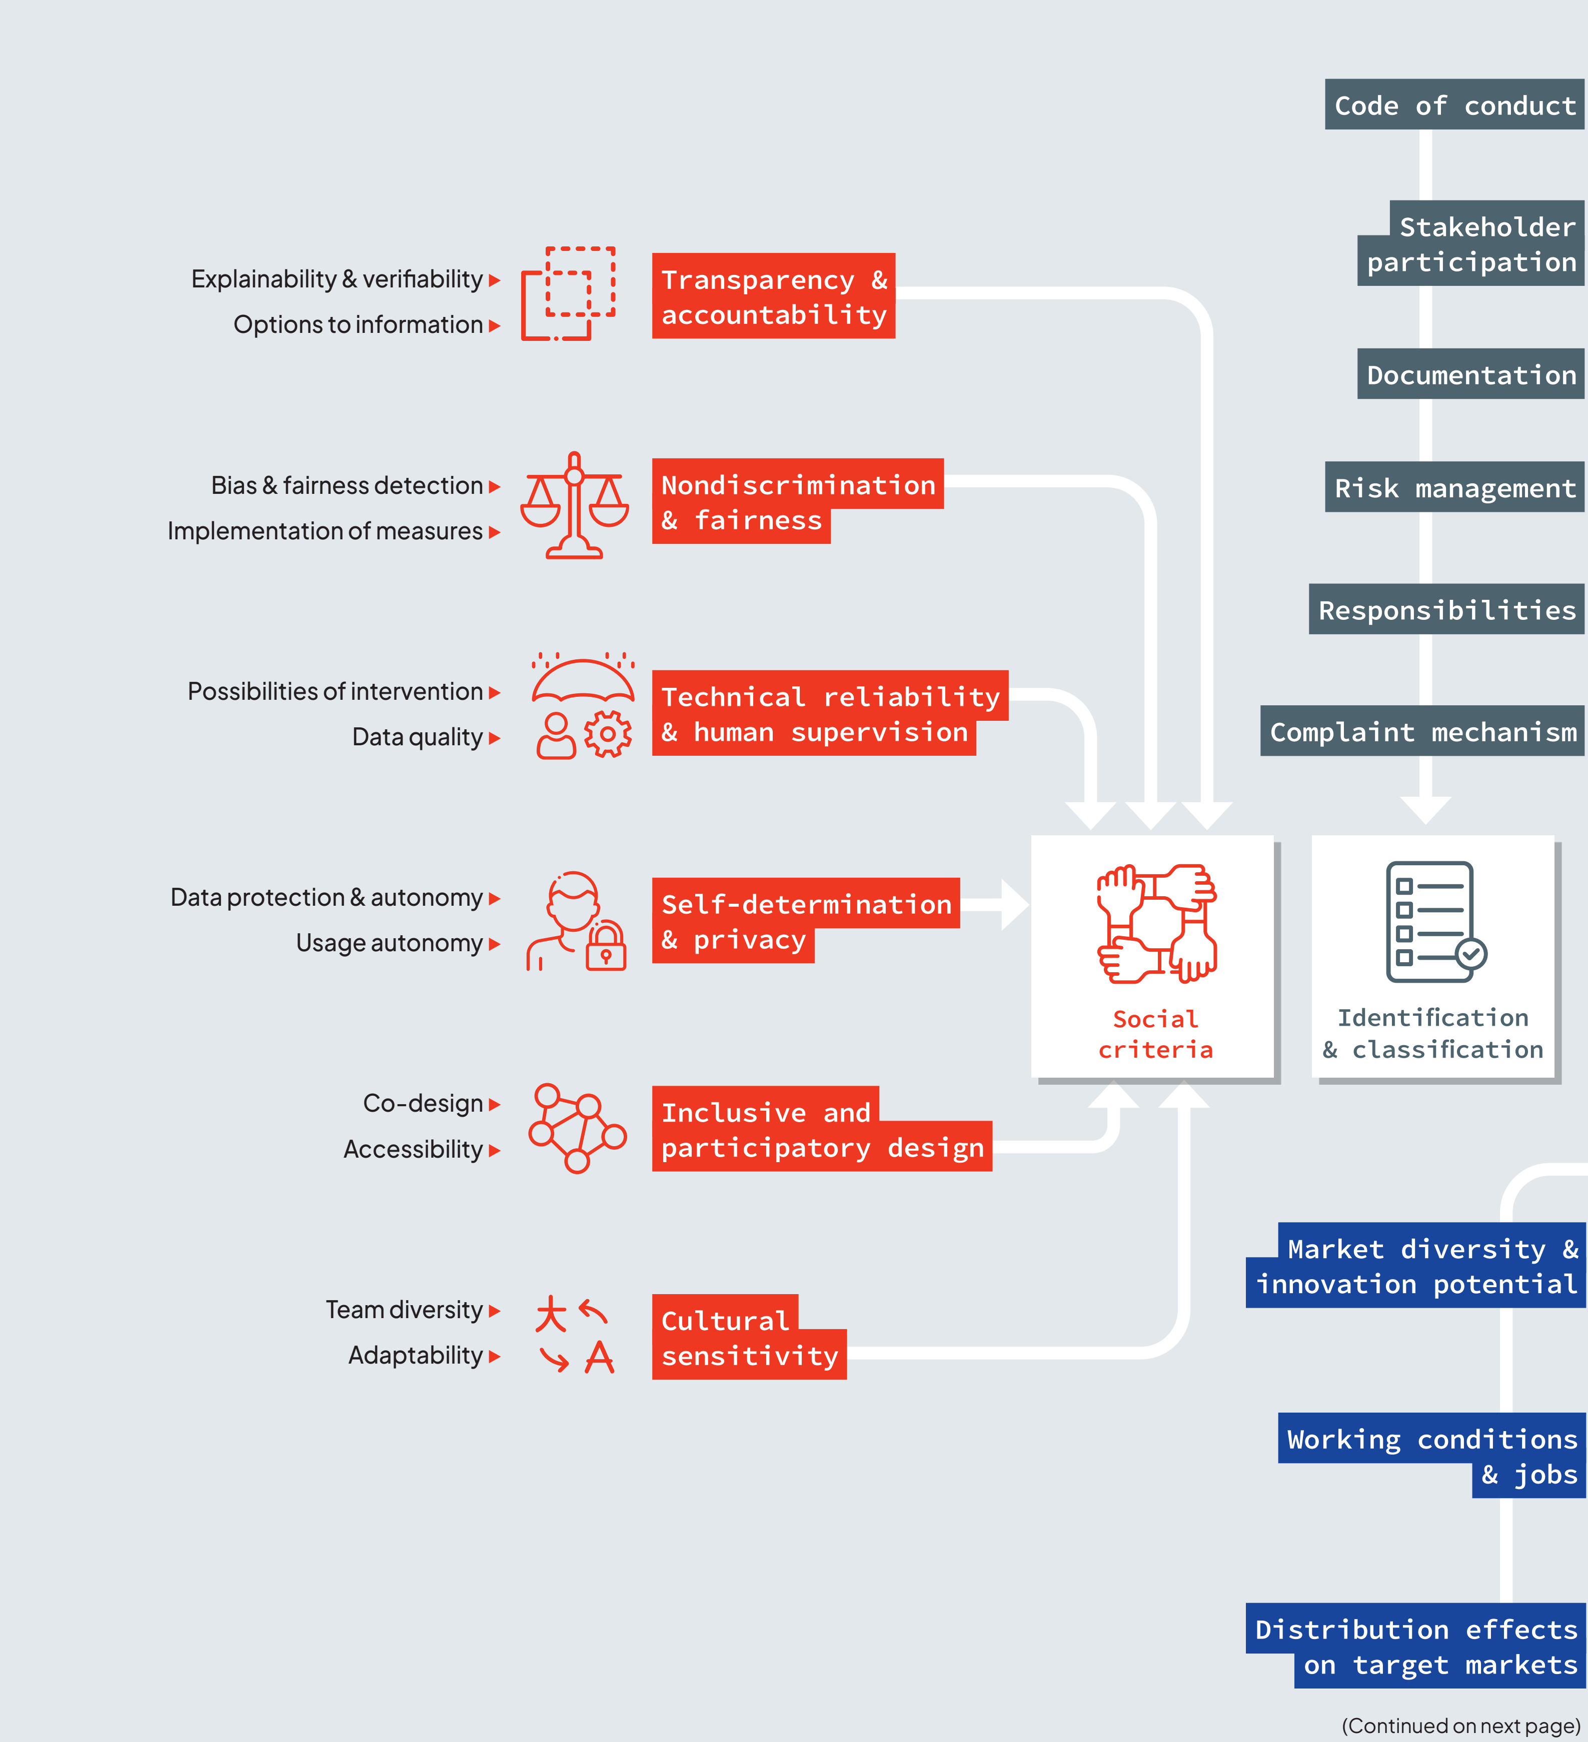

(Continued on next page)

Figure 1: Sustainability criteria and indicators identified and developed in the SustAIn Project (for

a detailed description of indicators see Rohde et al., 2024).

In the following, we present selected sustainability criteria and indicators for AI systems and possible regulatory steps to ensure more sustainable AI development. The selected criteria described in more detail below cover energy consumption and carbon emissions (4.1), embodied and shared resource consumption of hardware (4.2), cultural sensivity and global distributional injustices (4.3), working conditions and jobs (4.4), and market concentration (4.5).

4.1 Energy consumption & carbon emissions

Technological advances in AI allow systems to become more complex, and the size of ML models and the amount of data used for training them has increased significantly (especially in the domain of natural language processing) over the last few years (see Figure 2). This increase, moreover, results in an increasing demand for computing power needed for developing and training AI models. Even though the efficiency of computational hardware has risen substantially over the same time, with an exponential decrease in energy consumption per computation, that development has been superimposed by the trend towards more sophisticated AI systems with an improved prediction quality. While, initially, the increasing size of ML models was accompanied by a strong increase in performance in various areas, diminishing returns are now becoming increasingly apparent so that any improvements are now accompanied by disproportionate energy consumption (Thompson et al., 2021). Along with the advancing proliferation of AI use, the consumption of resources by AI systems in general and the associated ecological consequences are gaining increasing importance.

Moving towards logging impacts along the entire AI life cycle

As depicted in Figure 3, AI systems consume power along their whole life cycle, which includes data collection and storage, model development, training, and deployment (model inference). The development of new models, in which new architectures are designed in elaborate experiments, is particularly energy intensive. This intensity might be one of the reasons why research regarding green AI and the energy consumption of AI systems often focuses on the development and training phase (Verdecchia et al., 2023). The energy consumption of one individual training, and especially in an individual inference phase, are significantly lower. However, inference is repeated frequently: deployed AI systems (e.g., for virtual assistants or chatbots) are now being inferred millions of times a day. The inference phase thus might be responsible for up to 90 % of

POPULAR AI-MODELS SINCE 2010

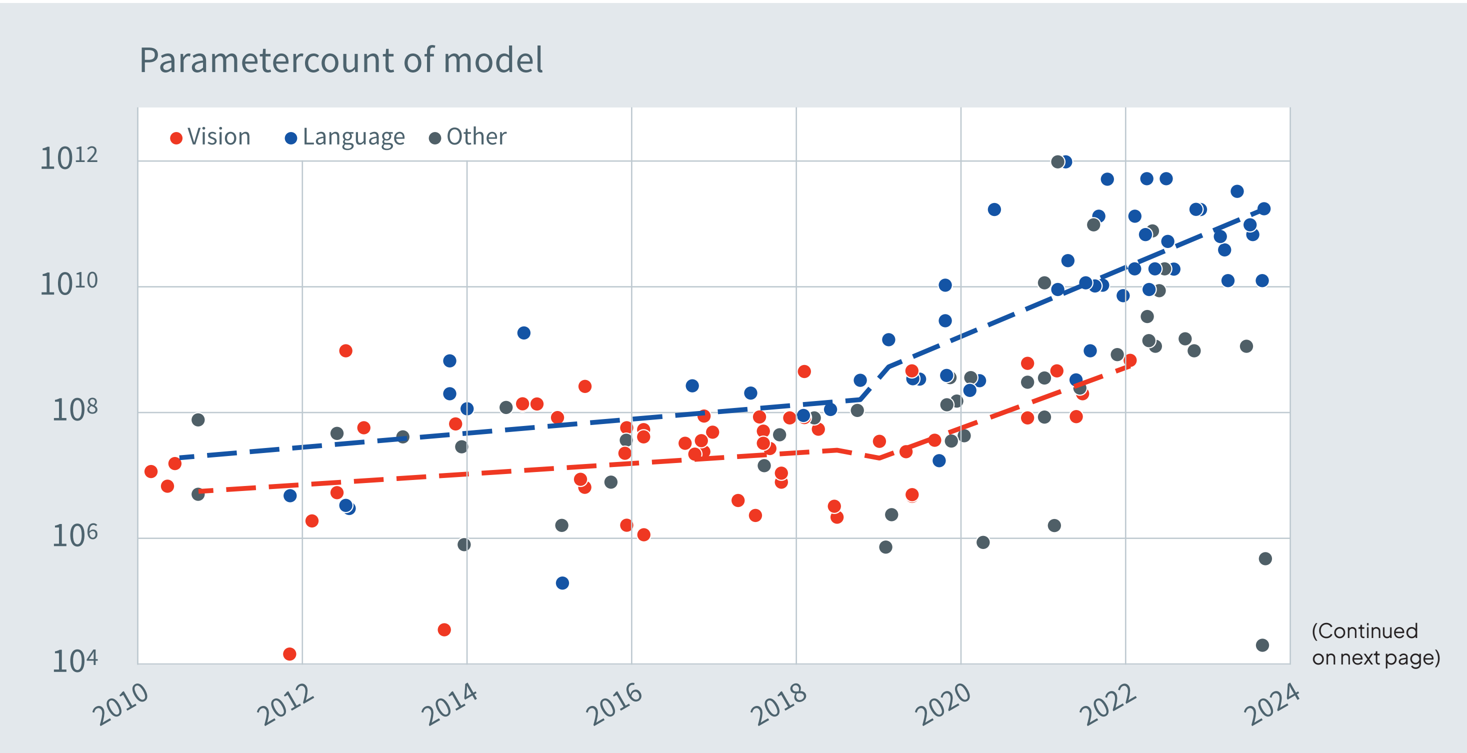

(Continued on next page)

Figure 2: Popular ML-Models since 2010 (source: own figure based on Dodge et al., 2022 and Sevilla et al., 2022)

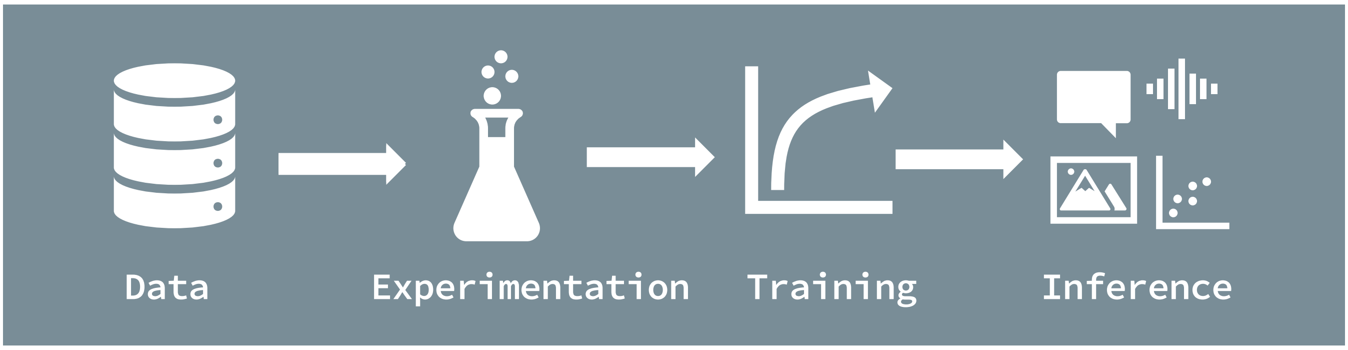

Figure 3: Power consumption along the AI lifecycle

the AI costs (Barr, 2019; Desislavov et al., 2023; Freund, 2019; Hernandez and Brown, 2020; Leopold, 2019; McDonald et al., 2022) and 60 % of its energy consumption (Patterson et al., 2022). With the broadening everyday use of AI systems, more attention should be paid to the inference phase when investigating their sustainability impact.

CO_2 emissions arise proportionally to the energy consumption in all life cycles of the AI system. If systems become more efficient and energy consumption is reduced, the carbon emissions can be reduced as well. In addition to energy consumption itself, CO_2 emissions depend on the specific time and location of consumption (see Figure 4). These two factors play a role in the CO_2 intensity of a data centre's energy mix. CO_2 emissions can be reduced by 1) shifting workloads to less carbon-intense locations and 2) training and deploying systems at times with high regenerative production.

Figure 5 shows that the CO_2 intensity of major AI system operators currently varies greatly. The particularly large players Amazon, Google, and Microsoft all define targets of CO_2 neutrality, CO_2 freedom, or even CO_2 negativity by 2030.

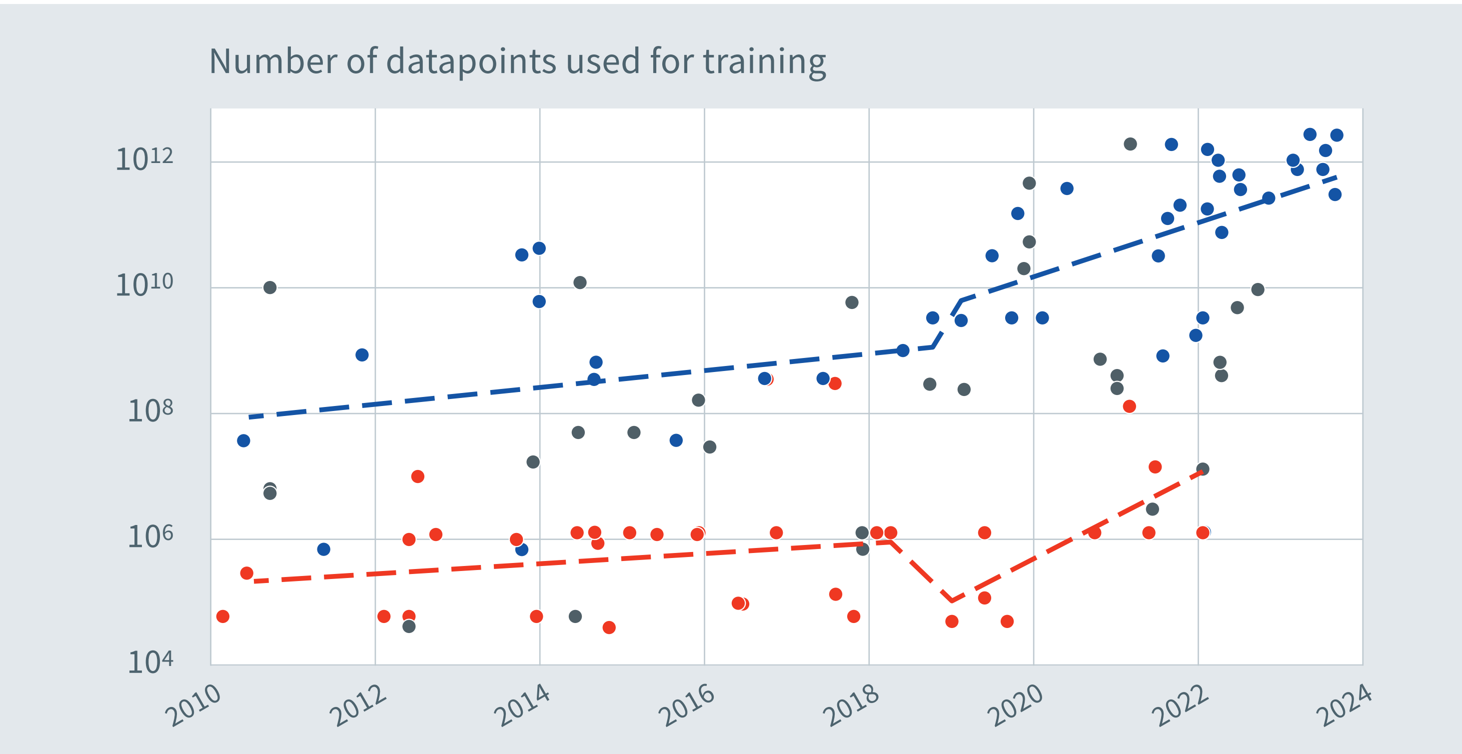

CO_2 GRAMS EMITTED, BERT LANGUAGE MODELLING

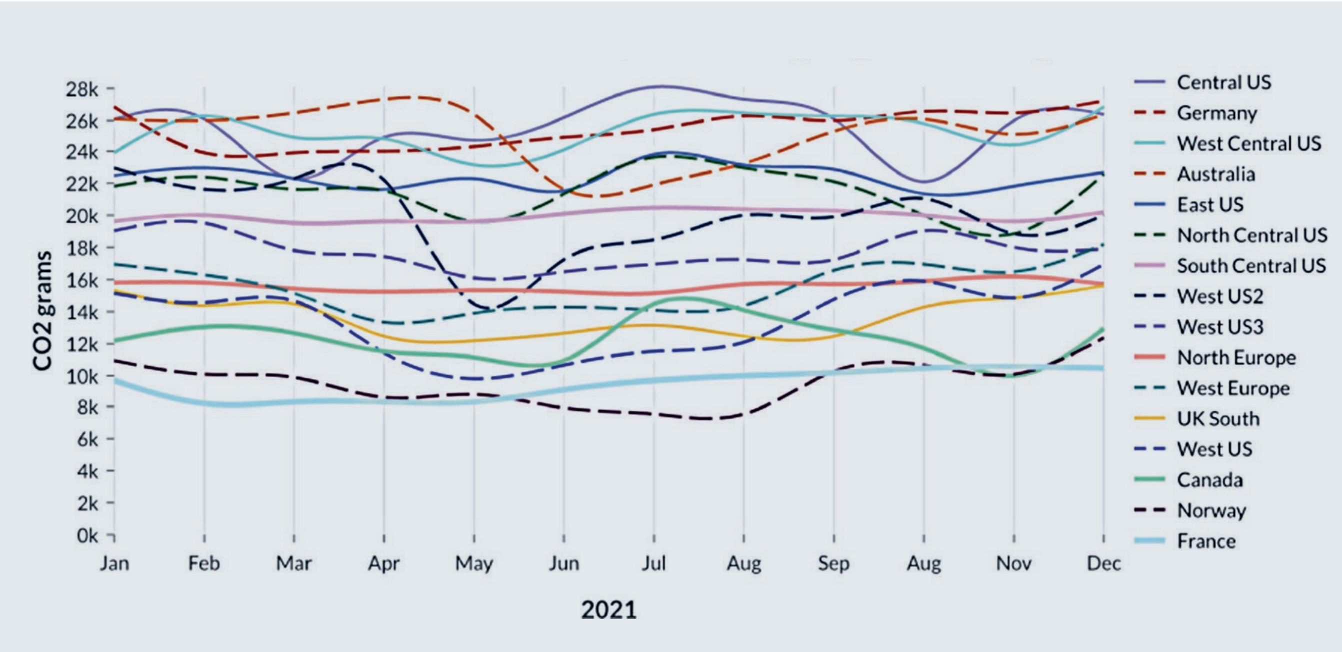

Figure 4: CO_2 in grams emitted in different locations, BERT Language Modeling (source: own figure based on data from https://www.climatiq.io/)

In principle, even if the envisioned balance-sheet neutrality is achieved by increasing the available renewable energy, that increase will not suffice to cover the accompanying rise in overall demand for energy. But, the corresponding additional generation capacities for renewable energies must still be created and the energy demand competes with other electricity consumers such as electrified mobility systems. Only when both the current and any future demand for energy can be covered

CO_2EQ EMISSIONS PER CPU-HOUR AT DIFFERENT LOCATIONS OF THE THREE BIGGEST CLOUD PROVIDERS

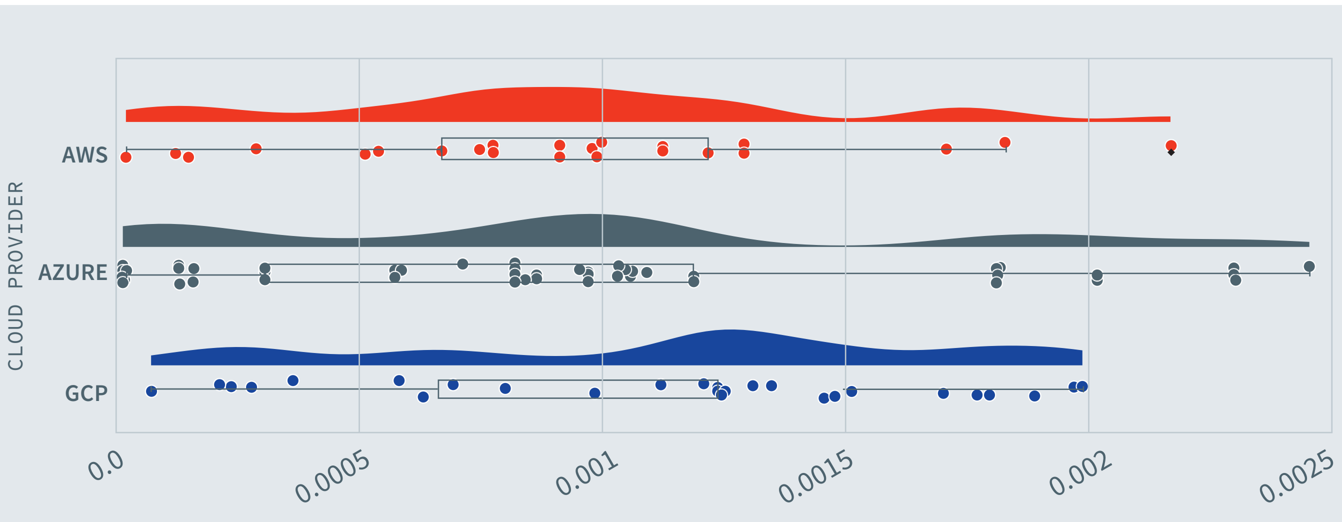

Figure 5: CO_2eq emissions per CPU-hour at different locations of the three biggest cloud providers (source: own figure based on data from https://huggingface.co/

AI-MODELS UPLOADED TO HUGGING FACE OVER TIME

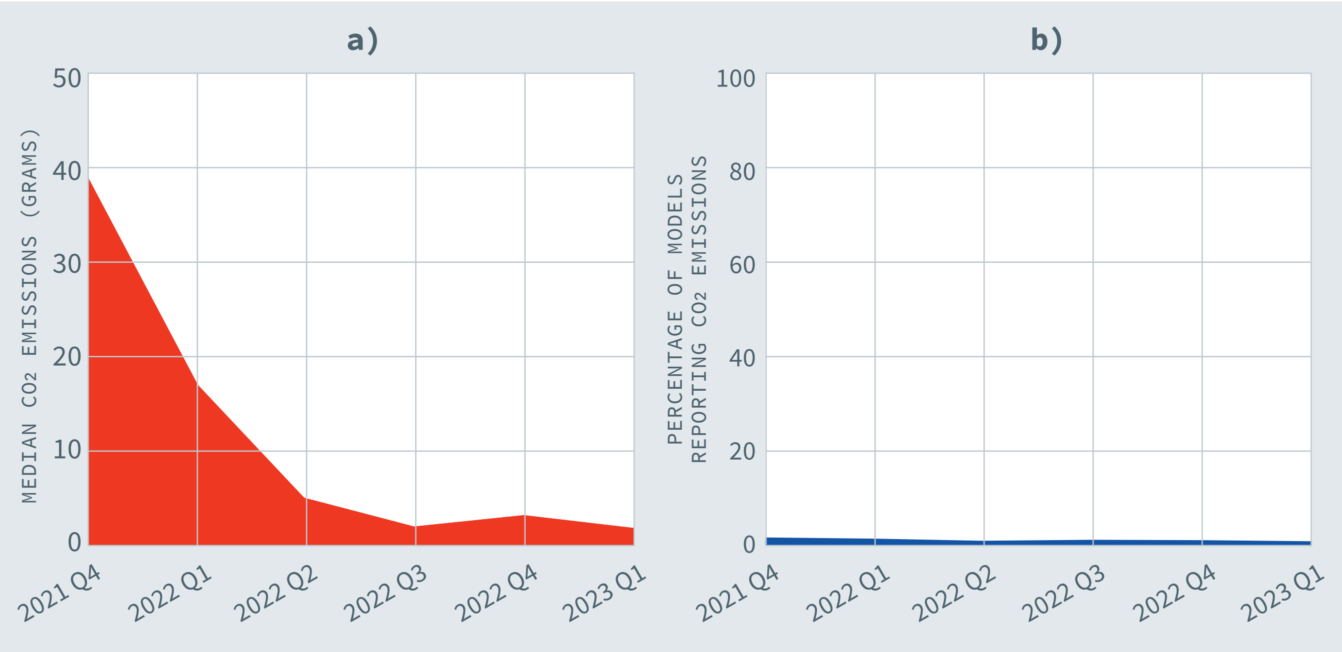

Figure 6: ML-Models uploaded to Hugging Face (source: own figure based on data from Luccioni et al., 2022)

by renewables can a positive contribution be made to climate protection. In this context, sustainability is closely related to political objectives such as the energy transition. An *absolute* reduction in energy consumption – also and especially regarding the use of digital infrastructure – is therefore of particular importance with a view to climate protection.

Establishing efficiency as a guiding goal of AI research and development

The development of large language models (LLM) can cause large CO_2 emissions (see Figure 4). However, such computationally intensive development processes rarely take place and are performed by only a few organisations (Kaack et al., 2021). Data from Hugging Face, a platform providing pre-trained models for developers, suggests that the emissions generated by model training are relatively low – for example, lower than the emissions of one hour of video streaming in 4k resolution on a TV set (see Figure 6a) However only around 1 % of the models added to the platform are annotated with information on CO_2 emissions, showing the lack of awareness for this topic among developers (see Figure 6b).

Preventing increasing environmental burdens from AI development requires focusing not only on model performance and prediction quality but also on establishing efficiency as a new (additional) guiding goal of AI research and development. Energy efficiency here refers to the ratio of achieved model performance to energy consumption. For this reason, it is important that companies and researchers developing and deploying AI systems provide the necessary information about the development process and the AI model. This information is essential when comparing results and creating incentives to produce more efficient and climate-friendly results (Strubell et al., 2019). There are more and more tools that allow energy consumption to be measured with little effort, such as CodeCarbon, Eco2AI, Cloud Carbon Footprint and Green Algorithms. In addition to more transparency in model development and deployment, developers also need to be skilled to ensure transparency is applied, to enhance the necessary measuring tools and methods, and to raise awareness.

Recent work by Luccioni et al. (2022) compared the power consumption and ensuing carbon emissions of several LLMs to investigate the scale of emissions of different sizes of LLMs (see also Luccioni, 2023). They determined that, depending on the energy source used for training and its carbon intensity, the training of an LLM emits between 30 und 552 metric tons of CO_2eq (see Figure 7), when taking into account the power usage effectiveness (PUE) of the data centre and the carbon intensity of the grid used (Luccioni et al., 2022). However, that calculation considers neither the manufacturing of the hardware used for training the models nor the emissions incurred when those models are deployed. We, thus, focus on the embodied and shared resource consumption in the next section.

ESTIMATED TRAINING CARBON FOOTPRINTS OF DIFFERENT LARGE LANGUAGE MODELS

Number of parameters (in billion)	Power Usage Effectiveness (PUE) data center	Carbon intensity of grid used (g CO_2eq/kWh)	Power consumption (MWh)	Emissions (tonnes CO_2eq)
175	1,1	428	1.287	502
280	1,08	330	1.066	352
175	1,09	231	324	70
176	1,2	57	433	25

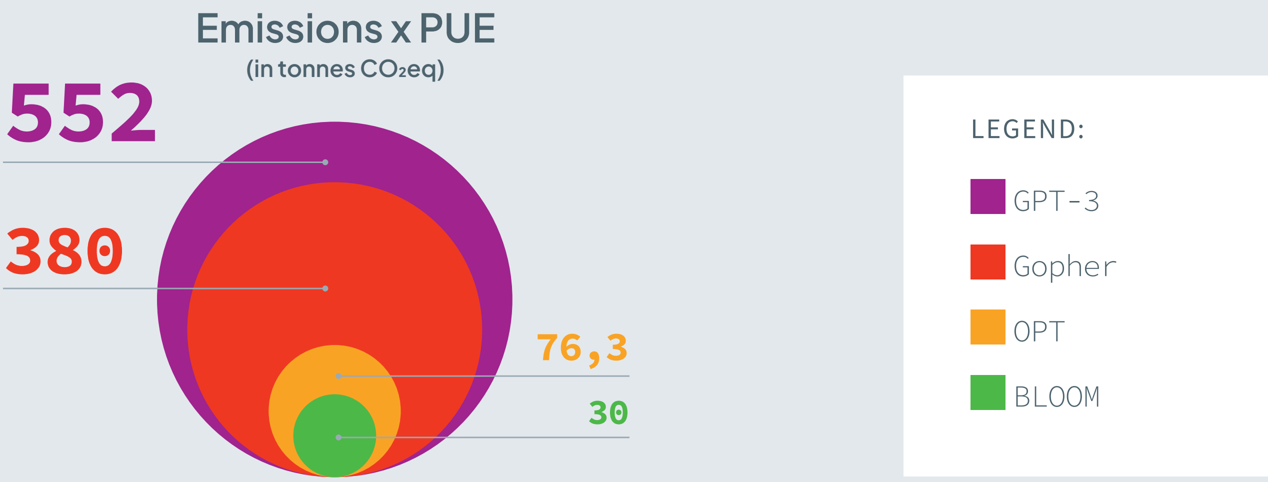

Figure 7: Estimated Training Carbon Footprints of different Large Language Models (LLMs) (source: own figure based on Luccioni et al., 2022)

4.2 Embodied and shared resource consumption of hardware infrastructure

The hardware infrastructure required for developing and deploying AI systems causes further environmental impacts in addition to the energy consumption and CO_2 emissions caused directly by the systems - especially by data centres. These impacts arise along the value chain of the hardware used (see Figure 9).

Acknowledging broad environmental impacts of AI infrastructures

In addition to the energy requirements of the hardware itself, additional energy is consumed in operating an AI system, for example, through cooling, lighting, and other consumption in a data centre (Whitehead et al., 2014). Currently, on average, 55 % of the energy required to operate the AI hardware is additionally required to maintain the infrastructure (Uptime Institute, 2021).

Depending on the energy mix, this energy consumption results in additional CO_2 emissions. Furthermore, to increase the energy efficiency of data centres, water is increasingly used for cooling. Particularly in areas where water is scarce, this use poses a problem. Furthermore, the water discharged from data centres is contaminated with various impurities that can have a negative impact on the environment (Andrews et al., 2021). While the energy efficiency of IT use in data centres is mostly recorded, CO_2 emissions and water use are often unknown (Uptime Institute, 2021). Analogous to carbon intensity, the water needed depends on external factors such as the weather. Figure 8 shows the difference in water needed to train the LLM LaMDA in four American data centres in different months. The third graph in Figure 8 also shows that, at times, there may be a conflict between reducing carbon emissions and reducing water usage since times when carbon emissions are low might be related to times when water use is intensive, stressing the need for holistically assessing the ecological impact of AI systems.

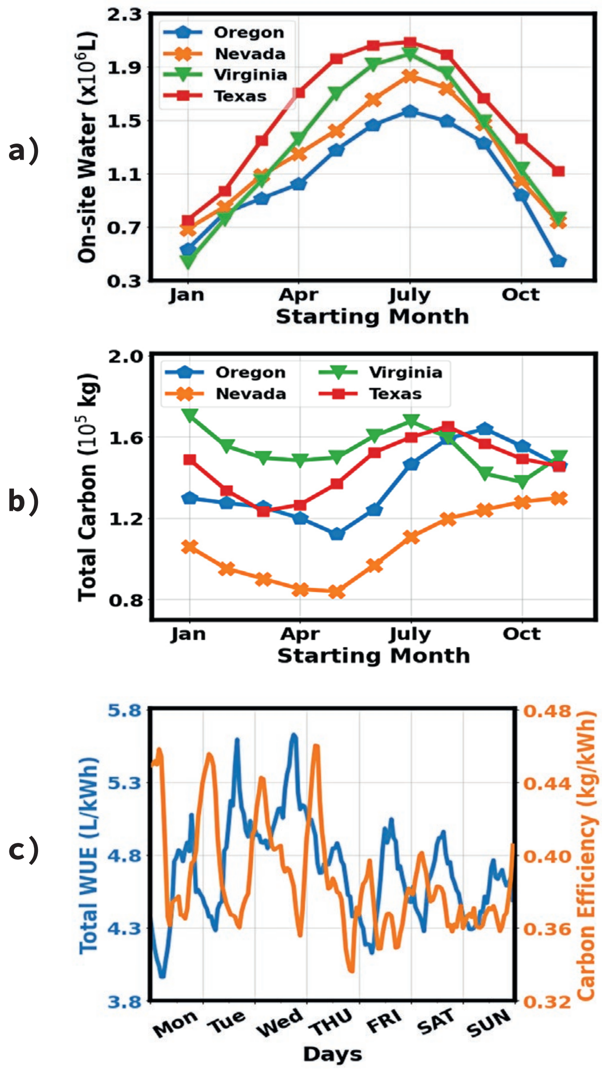

Figure 8: Estimated water (a) and carbon footprints (b) of training the LLM LaMDA with different starting months in 2022 in data centers at different US locations, and hourly carbon efficiency and on-site WUE for the first week of August 2022 in Oregon (c) (source: Li et al., 2023).

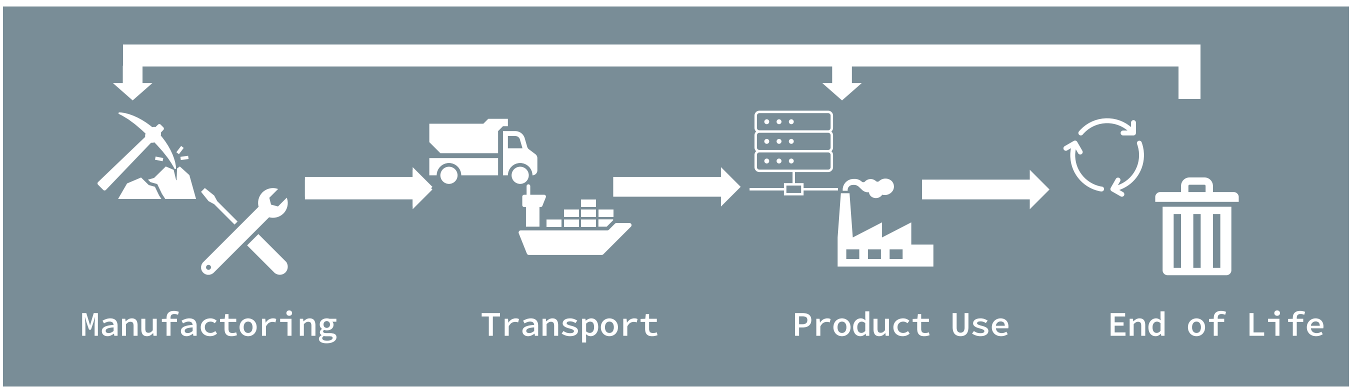

Figure 9: Value Chain of the Hardware used for ML Training and Inference

The increased use of and demands on AI have also led to an increase in the size of the infrastructure required to support it. For example, Facebook's hardware used for AI training and inference have increased, respectively, by a factor of 4 and 3.5 in less than two years (Gupta et al., 2020; Naumov et al., 2020; Park et al., 2018). To support these new applications, mobile devices such as smartphones are also incorporating more transistors and specialised circuitry than previous ones (Gupta et al., 2020). Furthermore, AI contributes to the accelerated and increased consumption of technology products, for example through technical obsolescence (Khakurel et al., 2018; Sætra, 2021b). The production of electronic and non-electronic components consumes electricity, raw materials (including precious metals and critical raw materials), chemicals, and water and generates hazardous waste. All of these factors can contribute to environmental problems (Uddin and Rahman, 2012). Impacts are amplified as equipment is regularly replaced (for example, data centre servers must be replaced every 1-5 years, batteries to ensure the uninterrupted power supply every 10 years) (Andrews and Whitehead, 2019).

With the increased use of AI and the hardware required to support it, which must be replaced regularly due to rapidly evolving technologies, the amount of electronic waste generated by AI is expected to increase. While about 70 to 90 % of the weight of each electronic product could now be recycled, few of these many different elements are actually recycled. For example, most, if not all, rare earth elements end up in landfills. Reasons for low recycling rates include complex material composition, lack of infra-structure in the form of high-tech recycling facilities, lack of incentives for manufacturing companies to optimise the longevity and recyclability of their products, and a lack of systems for recovering and recycling e-waste, as well as a lack of incentive for consumers to use existing systems (Kreps and Fors, 2020). While studies have been conducted on the generation of global electronic waste (Forti et al., 2020), AI's impact on it is largely unexplored.

4.3 Cultural sensitivity and global distributional injustices

AI developments are unevenly distributed across the globe, with a heavy dominance of US and China in the AI industry (Png, 2022; SustAIn Magazine 1, 2022). This is not only the case, when it comes to AI more specifically but holds true for digital technologies more generally (Coding Rights, n.d.). While Europe, parts of Asia, and North America are homes to big technology companies and people benefit from digital infrastructures, the rest of the world does not profit in the same way even though it provides the resources and labour and deals with the e-waste. Generative AI systems, which use available data to generate new data, have reinforced the entrenching of only a few cultural values into these systems' outputs (Birhane et al., 2022) and market concentration in the Global North is leading to forms of cultural hegemony in model outputs (SustAIn Magazine 3, 2023). These global distributional injustices (Crawford, 2021) and the heavy market concentration in the AI industry that comes with them (Widder et al., 2023), further manifest themselves in AI systems.

Issues of global distributional injustices and cultural hegemony exemplify how social, economic, and environmental sustainability impacts overlap. In their complexity, they are not easy to address on an organisational or regulatory level. Diversity in teams and cultural sensitivity in AI development could be one approach. Also, transparency measures, e.g., by disclosing information about training data through data sheets (Gebru et al., 2018) or model cards (Mitchell et al., 2019) could help ensure that models are only applied in appropriate contexts. Transparency requirements imposed on AI developers could, at least, assist in generating insights into the cultural values inscribed into AI models. Equally, open-sourcing ML models could help as those models would make it possible to adapt or re-train other models to fit with local environments. There is no one-size-fits-all solution to addressing this complex issue of intercultural appropriateness. What is certain, however, is that market concentration in the AI industry needs to be addressed.

4.4 Working conditions and jobs

From an economic perspective, securing fair working conditions along the entire AI value chain has to be considered when discussing the sustainability of AI systems. Ample reports refer to exploitative, precarious, inhumane, and dangerous working conditions, from extracting the minerals necessary for producing computing hardware, to the annotating work done by click workers, to the health-threatening work of people disposing of e-waste. These working conditions often remain unseen and unacknowledged, part of the hidden labour behind AI technologies (Gray and Suri, 2019; Miceli and Posada, 2022). While profits from AI development are usually made in the Global North, the precarious working conditions are often outsourced, through sub-contractors hired by big tech companies, to the Global South (Williams et al., 2022).

Precarious working conditions in producing and disposing of computer hardware have been apparent for a long time and are not AI-specific. However, in the interest of sustainability and with the ever-increasing amount of hardware needed to develop and run AI systems, the organisations developing and using those systems need to acknowledge their global responsibilities in ensuring safe and fair working conditions along the entire AI life cycle. What are AI specific are working conditions of click workers who are annotating or curating data sets and AI model outputs. They are often underpaid and do not receive adequate support in light of the often toxic and extremely disturbing content they have to deal with (Perrigo, 2023).

While it is certainly difficult for smaller and medium-sized companies developing or using AI systems to ensure that fair working conditions exist along the entire AI life cycle, big tech companies must live up to their responsibilities. With only a few major tech companies dominating the online platform industry and to some extent also the AI industry, those companies have a powerful position in negotiating contracts with subcontractors in Global South countries, where click workers are often being employed. Investigative research has demonstrated that those companies only ensure fair working conditions when put under some sort of pressure (Perrigo, 2023). It is necessary for regulators to step in here – with adequate supply chain regulation. The European Union's planned Corporate Sustainability Due Diligence Directive (CSDD) (European Commission, 2022) could assist in achieving this aim – also for smaller and medium-sized companies.

4.5 Market concentration

Concentration in AI markets is a significant risk in terms of economic sustainability. This risk stems from entry into the AI development market being essentially determined by three factors: i) access to a large and diverse amount of data, (ii) the availability of high computing power, (iii) access to expertise for the development of algorithms (Hall and Pesenti, 2017; Vollhardt et al., 2021). Big tech companies such as GAFAM and BATX,[1] whose business models are based on engaging as many users as possible and collecting and processing their data, have already emerged as AI market leaders, particularly thanks to large and closed data pools (Simon, 2019). Concentration in AI markets means that the application purposes (i.e. target functions of the algorithms) for which AI-based technologies are developed are determined by only a small number of players, which implies a concentration of power. Given that these few private players operate with a clear profit orientation rather than with a focus on the common good, power concentration in the development of technologies with potentially widespread application fields poses problems from a sustainability perspective, with the global distributional injustices outlined above being some of them.

Market concentration implications have reached new levels when it comes to foundation models. Foundation models are trained on broad data and can be adapted to a wide range of downstream tasks (Bommasani et al., 2022). The emergent phenomenon of in-context learning exhibited by foundation models gives rise to homogenization of models because in-context learning "enables users to provide a query and a few examples from which a model derives an answer without being trained on such queries" (Schneider 2022: 1). This means that a few foundation models might be able to replace a myriad of task-specific models, concentrating power over AI development in the hands of a few private players controlling these foundation models, or put more succinctly "any tendency towards monopolization may mean that incumbents do not compete just in a single market or even for a single market: in the limit, they may compete for the entire economy"(Vipra and Korinek, 2023: 5).

Foundation models generally tend towards natural monopoly due to their technological properties that generate significant economies of scale and certain economies of scope. Resultingly, competition policy is in high demand. Particular attention should be paid to strategic behaviour, such as vertical integration, pricing strategies, or strategic lobbying, that incumbent players might engage in to increase their market power and set up additional entry barriers. Regulatory control is needed for pricing strategies and when it comes to mergers, acquisition, and other forms of financial ties to limit vertical integration (ibid). If needed, principles from banking regulations could be adapted for separation in AI markets (Khan, 2017) and the functionally-based divisions in telecommunication groups could model future break ups of foundation model companies (Vipra and Korinek, 2023). In a similar vein, the governance of computational power (as a major limited resource and market entry barrier) could aim at preventing cloud providers from being invested in foundation model developers (Vipra and Korinek, 2023).

5 Transformative potential of AI for the sustainability transition: Voiced expectations and narratives of AI futures in the energy and mobility sector

In recent years, AI has increasingly been referred to as a transformative technology, albeit that the meaning of the term 'transformative' remains ambiguous (Gruetzemacher and Whittlestone, 2022).The label transformative AI involves the assessment that AI systems qualify as so-called general purpose technology due to their autonomous learning ability and their ability to make accurate predictions based on large amounts of data. As general purpose technology, AI systems are expected to be widely used, have many applications, and cause many spillover effects. It is assumed that these properties will lead to a practically irreversible change in human life, although assessments differ on how broad and extreme this change will be. Furthermore, by unlocking their transformative potential, AI systems are increasingly being attributed a key role in overcoming previously insurmountable challenges related to the necessary sustainability transitions.

1 GAFAM is an acronym for the "American giants" Google, Apple, Facebook, Amazon, and Microsoft, while BATX stands for the "Chinese giants" Baidu, Alibaba, Tencent, and Xiaomi.

In two case studies within the SustAIn project (see Wagner et al. 2023), we examined narratives and voiced expectations of AI futures in the energy and mobility sector in Germany – two sectors associated with high expectations related to achieving climate protection goals (Yigitcanlar and Cugurullo, 2020). Both sectors face challenges that are widely argued to be overcomable by implementing AI systems (e.g., Federal Government of Germany, 2022). A particular challenge for energy transitions is ensuring electricity grid stability at all times and thus to avoid power outages. However, electricity from renewable sources fluctuates depending on the weather (e.g., generation by wind or photovoltaic plants) and renewable energy plants are small and spatially decentralised compared to conventional power plants. Controlling grid flexibly is an essential condition for the success of the energy transition and high shares of renewable energies in the electricity mix. Due to advances in the field of Big Data analysis through ML, many AI-based use cases in the energy sector have emerged or been extended, e.g., for forecasting, demand-side management, or grid maintenance.

For the mobility sector, with its essential mobility transition and the related task of shifting the modal split from private to public transportation, providing public transport services in rural areas poses a major challenge, especially financially, due to sparse population, low use, and wide service areas (Klinge, 2021). While mobility offers within the concept of 'Mobility as a Service' represent a lucrative business field for the mobility industry in densely populated urban areas, rural regions remain economically unattractive. AI-supported mobility is often seen as a way out, with autonomously driven and networked minibuses expected to save labour costs and be used flexibly to supplement mobility services in rural areas (von Mörner and Boltze, 2018, Sinner et al., 2017).

As AI futures in both the energy and the mobility sectors are still uncertain, examining the narratives and voiced expectations around AI allowed us to uncover which developments concerning AI are considered relevant, urgent, possible, or inevitable (Konrad and Böhle, 2019). We, therefore, centred our analysis around two research questions: First, which expectations towards the future are voiced concerning the use of AI in the energy and mobility sectors? And second, how do the narratives of those AI futures envision solutions for sector-specific sustainability challenges?

For the energy sector, we investigated which promises are associated with using AI in the smart grid to integrate renewable energies. We analysed reports, position papers, and strategy papers as well as studies by various actor groups that have made statements on AI's role in the energy transition. Additionally, we conducted five interviews with actors connected to the specific case of energy optimisation in an energy neighbourhood project. For the mobility sector, we investigated the role attributed to AI-based autonomous and connected driving in the context of shifting the modal split from private to public transport, using the example of autonomous minibuses in rural areas. Expectations and promises were extracted from six strategy papers from different federal states and supplementary desktop research. All analysed documents address digital technologies and/or transitions in the mobility sector. In addition to political strategies, we evaluated self-descriptions and publicly available information about projects dealing with autonomous buses in rural areas. Moreover, we interviewed staff members of four of those projects.

We found that both sectors differ with regard to how AI future solutions are envisioned for sector-specific sustainability challenges. In the energy sector, AI is expected to enable renewable energy integration by dealing with complexity, improving the security of supply and system stability, and enhancing acceptance and participation in the energy transition. Therefore, AI futures envisioned for the energy sector have a clear orientation towards sustainability. However, they reveal a strong focus on opportunities while potential (sustainability) risks are underrepresented. Additionally, ambivalent developments are being muffled for the sake of strong narratives: e.g., the vision of a democratic and decentralised energy transition versus Big Data as a basis for AI-enabled renewable energy integration that brings advantage only for players with access to Big Data. With their narratives, actors promote AI as a solution to urgent societal challenges, e.g., climate change, and these voiced expectations promote a convergence of AI and sustainability visions. As such, climate protection also functions as a legitimisation for AI implementation in the energy sector, which has also been found in other areas of 'smart energy' developments (Rohde and Santarius, 2023).

In the mobility sector, by contrast, AI's expected contribution to climate protection remains vague. Rural areas are addressed as areas of action, but at the same time, AI-enabled autonomous driving is rarely expected to help shift the modal split. Therefore, AI futures envisioned for the mobility sector lack a clear orientation towards sustainability. Quite the opposite, the envisioned AI futures for the mobility sector reveal that implementing AI-enabled autonomous driving technologies currently only aims at incremental change instead of a mobility transition. More tangible than the expected contributions of AI to sustainable rural mobility is the (federal and) state intention to increase the attractiveness of industrial locations in German regions by implementing autonomous driving test fields or required infrastructures. Thus, the focus is generally on strengthening automotive and digital industries.

Finally, our findings invoke the question of whether the lack of vision for enabling a modal shift with the use of AI stems from there being no technical fix for the mobility transitions. Consequently, the following must be asked: If AI does not contribute to this sustainability challenge, is the use of AI in the mobility sector appropriate at all?

6 AI as a barrier to the sustainability transition: the case of online advertising

As shown above, AI applications do not hold sustainability potentials themselves. In fact, there are systems that tend to impede sustainable developments by causing additional ecological burdens and putting pressure on individuals or social groups. The case of personalised online marketing shows most prominently and clearly such direct and indirect negative effects on sustainability (Kish, 2020). Relying heavily on ML models that extract sales-related information from personal data online, the widespread targeting of potential customers on the web uses AI for profit-oriented purposes with questionable practices.

6.1 Ecological risks

With the internet and smartphones having become omnipresent, technology companies now have more detailed customer information at their fingertips than ever before. Detailed user information enables advertisers to tailor their ads to extremely precise target groups. This ability has revolutionised online advertising. Instead of developing broad advertising strategies aimed at reaching as many people as possible, companies have begun personalising their ads and placing them only where they will have the greatest effect – meaning where they convince potential consumers most effectively to visit a company's website or purchase a certain product, thus, increasing sales and profits. To achieve the most accurate targeting possible, AI systems are used to analyse huge amounts of user data and produce detailed user profiles. From those profiles, customers are divided up into target groups, and for each of those groups, a demand forecast is generated, which determines the advertising content shown to the users.

This practice has generated a significant debate, with most of the discussion focusing on data protection and ethical concerns. The potential ecological risks, however, are less frequently addressed. According to estimates, internet use is consuming over 400 TWh of electricity each year (González-Cabañas et al., 2023). Personalised advertising's contribution to this consumption through AI use and data analysis processes is not yet well researched. However, it can be expected that it could intensify these developments.

The user data used for personalisation comes from a number of different sources, including websites, social media platforms, and mobile applications. Network infrastructure and data centres are required for the data transfer, all of which consume energy. The data collected must be stored and managed for extended periods in data centres and on hard drives – and here too, energy is consumed. Furthermore 'real-time bidding' and rendering of content such as images and video, so that the advertisement can be presented on end devices, also requires energy (for details see Marken et al., 2024).

In addition to these direct ecological impacts stemming from the energy and material consumption of the technological infrastructure related to ML and personalised advertising, there are further indirect ecological impacts that are seldom discussed. Personalised online advertising and its AI use are often legitimised with the value proposition that every consumer is addressed with the most suitable, interesting and attractive product and service offer. However, this proposition is implicitly based on the assumption that every person who is online actually wants to buy something. From an environmental perspective, every additionally purchased product means additional environmental impacts in terms of consuming resources, such as raw materials, energy, and water, and pollution, environmental degradation, and CO_2 emissions (e.g., from manufacturing or transportation). Thus, every 'successful' placement of AI-powered advertising can increasingly burden the environment. Companies' high expenditure on online marketing indicates that advertising is not just a matter of shifting preferences and simplifying the product selection of an already existing purchase intention but that the purchase intention is being created in the first place and overall sales are being increased. Additionally, online advertising reinforces 'psychological obsolescence', where consumers are encouraged to replace a fully functioning product with a newer, more modern and more fashionable one, thus, creating needs that did not exist before.

6.2 Social risks

To personalise advertising, the online marketing industry depends on personal data. That data is continuously collected across websites and platforms via cookies, device, or browser fingerprints, advertising IDs and other digital identifiers. Even if single pieces of data are very small, the overall picture of a user can be detailed and disclose highly personal, sometimes intimate, insights such as personality traits, beliefs, or information on a person's income or family situation. However, the technical procedures and mechanisms used to collect this information are opaque and barely comprehensible to average online users (Armitage et al., 2023; Christl, 2017). Even though large companies make a profit from the data, online users usually do not know what personal data is captured at which point. The use of AI fuels effectively a mode of exploitation while regulators struggle to keep pace with technological developments and practices in the advertising industry.

Personalised advertising can be effective, in other words, persuasive. Given the largely 'invisible' mechanisms behind the endeavour of nudging individuals to pre-set directions, the targeting practices can even be labelled as manipulative (González-Cabañas et al., 2021; Petropoulos, 2022). In the advertising industry, it is a common commercial service to provide digital access to minors, people with financial problems or missing health insurance, politicians, homosexuals, depressive or pregnant people (Dachwitz, 2023). In light of these possibilities and practices, the freedom of choice of many individuals is clearly at risk.

6.3 Economic risks

Lastly, AI-based personalised online advertising entails significant economic risks – both on the level of the companies in the online advertising value chain and on a macroeconomic level. On the specific company level, using AI poses a risk, especially in content creation. AI is used to combine elements of texts and images in advertisements in order to personalise the advertisement towards the target group. This procedure is prone to errors, resulting in unintended messages being sent to potential consumers. Only a fraction of the vast number of conveyed personalised messages can be reviewed by an advertising company or advertising agency. This limited reviewing ability poses a risk to a company's image and reputation and can, from a sustainability perspective, lead to an uncontrolled misuse of green claims and the greenwashing of products and services. Further, while some companies greatly influence the data-driven marketing world due to their access to and power over data, technologies and infrastructures, smaller companies have a limited scope of action. By deliberately driving and expanding AI- and data-based advertising, the big players are forcing all other actors to follow their lead. Smaller companies are urged to participate in the modern advertising sector if they want to have a chance in the competition. As difficult as it is for internet users to decide against providing data and receiving advertising as difficult it is for companies of a certain size to decide against personalised advertising and the use of data. Here, those companies have limited means to exert pressure on bigger companies, and they only have a small scope of action for influencing data collection and analysis, design of campaigns, and data protection and environmental monitoring and measures.

From the macroeconomic perspective, personalised online advertising is characterised by market concentration and monopoly formation. The market is – despite the many actors involved in the 'advertising life cycle' – dominated by two big players: Alphabet and Meta. They are involved in collecting and analysing data, they employ some of the biggest intermediaries for brokering advertising spaces, and they publish the personalised advertisings on their platforms. They, thus, have immense influence, as all three main activities of advertising, mediating, and publishing lie within the same company. This 'duopoly' (Christl, 2017; Kingaby, 2020) is problematic for several reasons. The companies concerned concentrate large parts of data volumes and know-how within their company, making it impossible for other companies to compete and enabling a further power concentration. They have control over prices, practices, and technical standards, allowing them to influence the market in their favour. The power concentration of the big players, the speed at which the advertising industry develops, and the large information asymmetries that exist between Alphabet and Meta on one side and governments and users on the other have made it difficult to hold these actors accountable and put effective regulations in place (Amnesty International, 2019), as can be seen in the difficulties with enforcing current regulations such as the European GDPR.

Resolute measures are needed to limit the risks and effects of using AI for personalised online advertising. Some proposed measures show promising ways forward: bans on (AI- and data-based) advertising, the development of public digital infrastructure (e. g., a public search index), the development of business models that do not depend on data, and the education and training of individuals on effective data protection measures. However, they are still to be implemented.

7 Policy actions to promote the development of more sustainable AI systems

There seems to be consensus that more information about the sustainability impact of AI technologies is needed. While research in ML has exemplified, especially in relation to environmental impacts, that AI models can have immense impacts, a systematic approach to assessing such impacts is missing – especially on a regulatory level. While research points to structural causes for considerable sustainability impacts of AI technologies, the absence of publicly available information hampers the development and enactment of effective policies.

On the topic of environmental sustainability impacts, the European Union's AI Act, which is currently being drafted, could for the first time require companies to measure and disclose information on the environmental impact of certain AI systems. The European Parliament has proposed requiring companies to measure the energy and resource consumption of foundation models and high-risk systems. This requirement would mean that ways of collecting the relevant data are integrated into these systems.

Critics often claim that the obligation of measuring the environmental impacts of AI systems is too complicated and overburdens small and medium-sized enterprises in particular – and ultimately hinders innovation. But easy-to-use measurement methods already exist, e.g. CodeCarbon, Eco2AI, Cloud Carbon Footprint and Green Algorithms. They could be used to easily monitor energy consumption, CO_2 equivalent emissions, water consumption, the use of minerals for hardware, and the generation of electronic waste to assess the sustainability of AI systems.

7.1 Getting the whole picture

Without a comprehensive life cycle analysis, the environmental footprint of ML models cannot be adequately captured. Providers of LLMs, in particular, like to disclose only the direct energy consumption and emissions for one training cycle (Chowdhery et al., 2022). The result is an incomplete picture. Consider, for example, the training of the BLOOM model. The energy consumption during the training phase emits around 24.7 tons of CO_2 equivalents. But if you factor in hardware production and operational energy, the emissions value doubles (Luccioni et al., 2022). Moreover, this model does not yet encompass the continuous emissions produced while the model is being applied. Reliable figures from this inference phase are lacking, although first indicators suggest that emissions could be immense – during both the production of the necessary hardware for the application and its operation (SustAIn Magazine 2, 2023). It is now time to measure how AI systems impact the climate throughout their entire life cycle – from raw material extraction to disposal – so that informed decisions and targeted policies can be made based on solid knowledge.

While data centre operations and hardware production are known to contribute significantly to global carbon emissions (Rozite et al., 2023), there is a lack of specific and meaningful

Energy Consumption During System Development and Training				
Impact	**To Report**	**Process**	**Source**	**Purpose**
Energy	Hardware used (e.g., number GPU models)	Documentation	Provider/data centre	— Calculation energy use & emissions — Resource consumption manufacturing
	Number of FLOPs	Documentation	Provider/data centre	— Calculation energy use & emissions
	Computing time	Documentation	Provider/Data centre	— Calculation energy use & emissions
	GPU hours (equivalent depending on hardware)	Documentation	Provider/data centre	— Calculation energy use & emissions
	Energy used	Documentation	Provider/data centre	— Calculation emissions
	Power Usage Effectiveness of data centre (PuE)	Documentation	Provider/data centre	— Calculation energy use & emissions
	Hardware energy consumption — Infrastructure consumption (consumption without computing) — Idle consumption (consumption with computing standby) — Dynamic consumption (consumption with computing running)	Documentation/ information request	Provider/data centre	— Calculation energy use & emissions

Table 2: Data that should be recorded with regard to Energy Consumption during System Development and Training

data on AI systems' contribution in the related processes. This lack concerns the production and disposal of their hardware as well as their energy consumption along with all the resulting environmental damage, such as CO_2 emissions, pollution, resource extraction, and water consumption.

7.2 Logging relevant data

Companies can already automatically log and report much of the data needed to assess the sustainability of AI systems, such as operational data from computer systems – e.g. how often calculations are performed and how long these processes take. This metadata, for example stored in a spreadsheet, can be used to generate efficiency metrics. Metrics such as "Power Usage Effectiveness" (PuE), for example, show how much energy a data centre uses for computing in relation to its overall energy consumption. This parameter makes it possible to compare the energy efficiency of data centres. Tracking the power consumption allows the energy mix of the data centre, the carbon intensity of the energy grid, and the percentage of CO_2 the provider is potentially compensating to be calculated.

During system development and training alone, the following data (see Table 2) should be recorded to allow the energy consumption of AI systems to be comprehensively assessed and compared. Similar requirements can be formulated for all other environmental impacts, such as emissions, water consumption, mineral extraction, and hardware disposal.

7.3 Detailed and standardised reporting needed

The life cycle approach demonstrates that various stakeholders need to provide accurate measurements. For instance, hardware manufacturers, such as Nvidia, should disclose environmental data on products that are widely used in developing and applying ML models.

Numerous measurement methods already exist for assessing the environmental impact during system development and training, material extraction, hardware manufacturing and disposal, as do different ones for tracking carbon. Some hardware manufacturers already report emissions levels for some of their products. Other approaches for assessing environmental impacts during system deployment need to be developed – reliable metrics and comparable units of measurement for assessing emissions during the application phase, for example.

To be as accurate as possible about the environmental impact during the deployment phase, we propose that AI providers define various standard usage scenarios prior to market launch.

7.4 Measuring environmental impacts during system deployment

Procedures and methods for assessing environmental impact during system deployment have not yet been widely established in practice. Developers of AI systems can record their energy consumption during training. However, under the requirements formulated within the AI Act to document energy consumption, recording will most likely not be feasible during inference. Thus, energy consumption during the application phase and the extent of emissions generated during this phase must be estimated. To that end, we are proposing two basic options, possibly in combination:

- Before an AI product is released on the market, different standard use scenarios (low-, middle-, high-use) should be evaluated based on test runs or, preferably, simulations.
- After market introduction, the de-facto average energy consumption over a certain period of time should be calculated. This calculation would allow the estimated standard use scenarios to be evaluated and adjusted if they deviate significantly from the actual value.

7.5 Greater transparency is feasible – and overdue

There is no lack of technical means for measuring the environmental impact of AI systems. There is, however, still a lack of political will to make AI more sustainable. This lack is all the more irresponsible considering that AI is a resource- and energy-intensive technology that is becoming increasingly pervasive. The European Parliament has taken some important steps in the right direction to ensure that AI does not further harm the environment, the climate, people, or the planet. Nevertheless, more data on environmental impacts of AI systems is indispensable. Clear and comprehensive requirements must be introduced for publicly available reports that include such data. These requirements could make AI systems more environmentally sustainable while simultaneously distributing their risks, harmful consequences, and benefits more equitably around the world.

7.6 Sustainability impacts beyond the environment

With sustainability aspects of AI technologies mostly being associated with environmental impacts, policy discussion often centres around energy or water consumption, resource extraction, and end of life/disposal issues. But from a broad sustainability perspective, also considering economic and social sustainability aspects, further policy approaches could contribute towards making AI technologies more sustainable.

From an economic and social perspective, problematic tendencies in AI development lie rooted in a heavy market concentration in the AI industry leading to global distributional injustices, high entry barriers for new market actors, unequal access to data, exploitation of labour etc. (see above). Especially in the interest of a strong European AI industry, it should thus be in the prime interest of European and national policy makers to address market concentration. Precedence has been created in the tech industry, with the Digital-Markets-Act (DMA) regulation of large online platforms. But further policy initiatives could contribute towards reducing market concentration and ensuring sustainable development and use of AI systems:

- Data Governance (e. g. the European Union's Data act, Data Governance Act, national data governance initiatives), especially relevant for data sharing, data access, and data value creation
- Supply Chain Legislation (ensuring fundamental rights and fair as well as safe working conditions along the entire value chain)
- Eco Design Regulation (e.g., the European Union's Eco Design Directive) aiming for more circularity and more environmentally sustainable product design
- Ethical Frameworks for AI development.

If the EU is serious about aligning the use of AI with the common good, then it should put all people, and not just Europeans, at the centre of its focus. Whatever form the AI legislation in Europe eventually takes, people will only be truly protected from the negative consequences of AI systems if the impact of those systems on sustainability is effectively monitored.

8 References

Ahlborg, H., Ruiz-Mercado, I., Molander, S., Masera, O., 2019. Bringing Technology into Social-Ecological Systems Research—Motivations for a Socio-Technical-Ecological Systems Approach. Sustainability 11, 2009. https://doi.org/10.3390/su11072009

Amnesty International, 2019. Surveillance Giants: How the Business Model of Google and Facebook threatens Human Rights. Amnesty International Ltd, London.

Andrews, D., Newton, E.J., Adibi, N., Chenadec, J., Bienge, K., 2021. A Circular Economy for the Data Centre Industry: Using Design Methods to Address the Challenge of Whole System Sustainability in a Unique Industrial Sector. Sustainability 13, 6319. https://doi.org/10.3390/su13116319

Andrews, D., Whitehead, B., 2019. Data Centres in 2030: Comparative Case Studies that Illustrate the Potential of the Design for the Circular Economy as an Enabler of Sustainability. Presented at the Sustainable Innovation 2019: 22nd International Conference Road to 2030: Sustainability, Business Models, Innovation and Design, Epsom, Surrey.

Armitage, C., Botton, N., Dejeu-Castang, L., Lemoine, L., 2023. Study on the impact of recent developments in digital advertising on privacy, publishers and advertisers: final report. Publications Office, LU.

Barr, J., 2019. Amazon EC2 Update – Inf1 Instances with AWS Inferentia Chips for High Performance Cost-Effective Inferencing. AWS News Blog. https://aws.amazon.com/de/blogs/aws/amazon-ec2-update-inf1-instances-with-aws-inferentia-chips-for-high-performance-cost-effective-inferencing/ (accessed 11.30.23).

Bender, E.M., Gebru, T., McMillan-Major, A., 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. Association for Computing Machinery, New York, NY, USA, pp. 610–623. https://doi.org/10.1145/3442188.3445922

Birhane, A., Kalluri, P., Card, D., Agnew, W., Dotan, R., Bao, M., 2022. The Values Encoded in Machine Learning Research, in: 2022 ACM Conference on Fairness, Accountability, and Transparency. Presented at the FAccT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency, ACM, Seoul Republic of Korea, pp. 173–184. https://doi.org/10.1145/3531146.3533083

Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J.Q., Demszky, D., Donahue, C., Doumbouya, M., Durmus, E., Ermon, S., Etchemendy, J., Ethayarajh, K., Fei-Fei, L., Finn, C., Gale, T., Gillespie, L., Goel, K., Goodman, N., Grossman, S., Guha, N., Hashimoto, T., Henderson, P., Hewitt, J., Ho, D.E., Hong, J., Hsu, K., Huang, J., Icard, T., Jain, S., Jurafsky, D., Kalluri, P., Karamcheti, S., Keeling, G., Khani, F., Khattab, O., Koh, P.W., Krass, M., Krishna, R., Kuditipudi, R., Kumar, A., Ladhak, F., Lee, M., Lee, T., Leskovec, J., Levent, I., Li, X.L., Li, X., Ma, T., Malik, A., Manning, C.D., Mirchandani, S., Mitchell, E., Munyikwa, Z., Nair, S., Narayan, A., Narayanan, D., Newman, B., Nie, A., Niebles, J.C., Nilforoshan, H., Nyarko, J., Ogut, G., Orr, L., Papadimitriou, I., Park, J.S., Piech, C., Portelance, E., Potts, C., Raghunathan, A., Reich, R., Ren, H., Rong, F., Roohani, Y., Ruiz, C., Ryan, J., Ré, C., Sadigh, D., Sagawa, S., Santhanam, K., Shih, A., Srinivasan, K., Tamkin, A., Taori, R., Thomas, A.W., Tramèr, F., Wang, R.E., Wang, W., Wu, B., Wu, J., Wu, Y., Xie, S.M., Yasunaga, M., You, J., Zaharia, M., Zhang, M., Zhang, T., Zhang, X., Zhang, Y., Zheng, L., Zhou, K., Liang, P., 2022. On the Opportunities and Risks of Foundation Models.

Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H.W., Sutton, C., Gehrmann, S., Schuh, P., Shi, K., Tsvyashchenko, S., Maynez, J., Rao, A., Barnes, P., Tay, Y., Shazeer, N., Prabhakaran, V., Reif, E., Du, N., Hutchinson, B., Pope, R., Bradbury, J., Austin, J., Isard, M., Gur-Ari, G., Yin, P., Duke, T., Levskaya, A., Ghemawat, S., Dev, S., Michalewski, H., Garcia, X., Misra, V., Robinson, K., Fedus, L., Zhou, D., Ippolito, D., Luan, D., Lim, H., Zoph, B., Spiridonov, A., Sepassi, R., Dohan, D., Agrawal, S., Omernick, M., Dai, A.M., Pillai, T.S., Pellat, M., Lewkowycz, A., Moreira, E., Child, R., Polozov, O., Lee, K., Zhou, Z., Wang, X., Saeta, B., Diaz, M., Firat, O., Catasta, M., Wei, J., Meier-Hellstern, K., Eck, D., Dean, J., Petrov, S., Fiedel, N., 2022. PaLM: Scaling Language Modeling with Pathways. https://doi.org/10.48550/ARXIV.2204.02311

Christl, W., 2017. Corporate Surveillance In Everyday Life. How Companies Collect, Combine, Analyze, Trade, and Use Personal Data on Billions.

Cockburn, I.M., Henderson, R., Stern, S., 2019. The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis, in: The Economics of Artificial Intelligence. University of Chicago Press, pp. 115–148. https://doi.org/10.7208/9780226613475-006

Coding Rights, n.d. Tech Cartographies - Your Cloud is in Territories. Tech Cartographies. https://www.cartografiasdainternet.org/en (accessed 12.6.23).

Crawford, K., 2021. Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence. Yale University Press. https://doi.org/10.12987/9780300252392

Dachwitz, I., 2023. Microsofts Datenmarktplatz Xandr: Das sind 650.000 Kategorien, in die uns die Online-Werbeindustrie einsortiert. netzpolitik.org.

Desislavov, R., Martínez-Plumed, F., Hernández-Orallo, J., 2023. Trends in AI inference energy consumption: Beyond the performance-vs-parameter laws of deep learning. Sustainable Computing: Informatics and Systems 38, 100857. https://doi.org/10.1016/j.suscom.2023.100857

Dodge, J., Prewitt, T., Tachet des Combes, R., Odmark, E., Schwartz, R., Strubell, E., Luccioni, A.S., Smith, N.A., DeCario, N., Buchanan, W., 2022. Measuring the Carbon Intensity of AI in Cloud Instances, in: 2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22. Association for Computing Machinery, New York, NY, USA, pp. 1877–1894. https://doi.org/10.1145/3531146.3533234

Durant, J., Bauer, M.W., Gaskell, G., 1998. Biotechnology in the public sphere: a European sourcebook. Science Museum, London.

European Commission, 2022. Corporate sustainability due diligence - Fostering sustainability in corporate governance and management systems. https://commission.europa.eu/business-economy-euro/doing-business-eu/corporate-sustainability-due-diligence_en (accessed 12.6.23).

Federal Government of Germany (2022). „Digital Strategy - Creating Digital Values Together

Forti, V., Balde, C.P., Kuehr, R., Bel, G., 2020. The Global E-waste Monitor 2020: Quantities, flows and the circular economy potential. United Nations University/ United Nations Institute for Training and Research, International Telecommunication Union, and International Solid Waste Association.

Freund, K., 2019. Google Cloud Doubles Down On NVIDIA GPUs For Inference. Forbes. https://www.forbes.com/sites/moorinsights/2019/05/09/google-cloud-doubles-down-on-nvidia-gpus-for-inference/ (accessed 11.29.23).

Galaz, V., Centeno, M.A., Callahan, P.W., Causevic, A., Patterson, T., Brass, I., Baum, S., Farber, D., Fischer, J., Garcia, D., McPhearson, T., Jimenez, D., King, B., Larcey, P., Levy, K., 2021. Artificial intelligence, systemic risks, and sustainability. Technology in Society 67, 101741. https://doi.org/10.1016/j.techsoc.2021.101741

Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J.W., Wallach, H., Daumé, H., Crawford, K., 2018. Datasheets for Datasets. https://doi.org/10.48550/ARXIV.1803.09010

González-Cabañas, J., Callejo, P., Cuevas, R., Svatberg, S., Torjesen, T., Cuevas, Á., Pastor, A., Kotila, M., 2023. CarbonTag: A Browser-Based Method for Approximating Energy Consumption of Online Ads. IEEE Trans. Sustain. Comput. 1≣12. https://doi.org/10.1109/TSUSC.2023.3286916

González-Cabañas, J., Cuevas, Á., Cuevas, R., López-Fernández, J., García, D., 2021. Unique on Facebook: Formulation and Evidence of (Nano)targeting Individual Users with non-PII Data, in: Proceedings of the 21st ACM Internet Measurement Conference. pp. 464≣479. https://doi.org/10.1145/3487552.3487861

Gray, M.L., Suri, S., 2019. Ghost work: how to stop Silicon Valley from building a new global underclass. Houghton Mifflin Harcourt, Boston New York NY.

Gruetzemacher, R., Whittlestone, J., 2022. The transformative potential of artificial intelligence. Futures 135, 102884. https://doi.org/10.1016/j.futures.2021.102884

Gupta, V., Choudhary, D., Tang, P.T.P., Wei, X., Wang, X., Huang, Y., Kejariwal, A., Ramchandran, K., Mahoney, M.W., 2020. Fast Distributed Training of Deep Neural Networks: Dynamic Communication Thresholding for Model and Data Parallelism. Preprint.

Hall, D.W., Pesenti, J., 2017. Growing the artificial intelligence industry in the UK. (Part of the Industrial Strategy UK and the Commonwealth). Department for Digital, Culture, Media & Sport and Department for Business, Energy & Industrial Strategy.

Hernandez, D., Brown, T., 2020. AI and efficiency. OpenAI. https://openai.com/research/ai-and-efficiency (accessed 11.30.23).

Kaack, L.H., Donti, P.L., Strubell, E., Kamiya, G., Creutzig, F., Rolnick, D., 2021. Aligning artificial intelligence with climate change mitigation.

Khakurel, J., Penzenstadler, B., Porras, J., Knutas, A., Zhang, W., 2018. The Rise of Artificial Intelligence under the Lens of Sustainability. Technologies 6, 100. https://doi.org/10.3390/technologies6040100

Khan, L.M., 2017. Amazon's Antitrust Paradox. Yale Law Journal 3, 564–907.

Kingaby, H., 2020. AI and Advertising - A consumer perspective.

Kish, K., 2020. Paying Attention: Big Data and Social Advertising as Barriers to Ecological Change. Sustainability 12, 10589. https://doi.org/10.3390/su122410589

Klinge, A. (2021, Nov. 18). *Ländliche Mobilität*. Bundeszentrale für politische Bildung. https://www.bpb.de/themen/stadt-land/laendliche-raeume/335912/laendliche-mobilitaet/.

Konrad, K., Böhle, K., 2019. Socio-technical futures and the governance of innovation processes—An introduction to the special issue. Futures 109, 101–107. https://doi.org/10.1016/j.futures.2019.03.003

Kreps, D., Fors, P., 2020. A Resource Perspective on E-Waste: A Global Problem with Local Solutions?, in: Strous, L., Johnson, R., Grier, D.A., Swade, D. (Eds.), Unimagined Futures – ICT Opportunities and Challenges, IFIP Advances in Information and Communication Technology. Springer International Publishing, Cham, pp. 129–141. https://doi.org/10.1007/978-3-030-64246-4_11

Leopold, G., 2019. AWS to Offer Nvidia's T4 GPUs for AI Inferencing. HPCwire. https://www.hpcwire.com/2019/03/19/aws-upgrades-its-gpu-backed-ai-inference-platform/ (accessed 11.30.23).

Li, P., Yang, J., Islam, M.A., Ren, S., 2023. Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. https://doi.org/10.48550/ARXIV.2304.03271

Luccioni, A.S., Viguier, S., Ligozat, A.-L., 2022. Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model. https://doi.org/10.48550/ARXIV.2211.02001

Luccioni, S., 2023. The mounting human and environmental costs of generative AI. Ars Technica. https://arstechnica.com/gadgets/2023/04/generative-ai-is-cool-but-lets-not-forget-its-human-and-environmental-costs/ (accessed 11.30.23).

Marken, G., Frick, V., Schmelzle, F., Meyer, A., 2024. The (Un-)Sustainability of Artificial Intelligence in Online Marketing - A Case Study on the Environmental, Social and Economic Impacts of Personalized Advertising. IÖW Text Series 228/24.

McDonald, J., Li, B., Frey, N., Tiwari, D., Gadepally, V., Samsi, S., 2022. Great Power, Great Responsibility: Recommendations for Reducing Energy for Training Language Models. https://doi.org/10.48550/ARXIV.2205.09646

Miceli, M., Posada, J., 2022. The Data-Production Dispositif. Proc. ACM Hum.-Comput. Interact. 6, 1–37. https://doi.org/10.1145/3555561

Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I.D., Gebru, T., 2019. Model Cards for Model Reporting. Proceedings of the Conference on Fairness, Accountability, and Transparency 220–229. https://doi.org/10.1145/3287560.3287596

Naumov, M., Kim, J., Mudigere, D., Sridharan, S., Wang, X., Zhao, W., Yilmaz, S., Kim, C., Yuen, H., Ozdal, M., Nair, K., Gao, I., Su, B.-Y., Yang, J., Smelyanskiy, M., 2020. Deep Learning Training in Facebook Data Centers: Design of Scale-up and Scale-out Systems. https://doi.org/10.48550/ARXIV.2003.09518

Park, J., Naumov, M., Basu, P., Deng, S., Kalaiah, A., Khudia, D., Law, J., Malani, P., Malevich, A., Nadathur, S., Pino, J., Schatz, M., Sidorov, A., Sivakumar, V., Tulloch, A., Wang, X., Wu, Y., Yuen, H., Diril, U., Dzhulgakov, D., Hazelwood, K., Jia, B., Jia, Y., Qiao, L., Rao, V., Rotem, N., Yoo, S., Smelyanskiy, M., 2018. Deep Learning Inference in Facebook Data Centers: Characterization, Performance Optimizations and Hardware Implications. https://doi.org/10.48550/ARXIV.1811.09886

Patterson, D., Gonzalez, J., Holzle, U., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D.R., Texier, M., Dean, J., 2022. The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink. Computer 55, 18–28. https://doi.org/10.1109/MC.2022.3148714

Perrigo, B., 2023. Exclusive: OpenAI Used Kenyan Workers on Less Than $2 Per Hour to Make ChatGPT Less Toxic. Time Magazine.

Petropoulos, G., 2022. The dark side of artificial intelligence: manipulation of human behaviour [WWW Document]. Bruegel | The Brussels-based economic think tank. https://www.bruegel.org/blog-post/dark-side-artificial-intelligence-manipulation-human-behaviour (accessed 3.5.23).

Png, M.-T., 2022. At the Tensions of South and North: Critical Roles of Global South Stakeholders in AI Governance, in: 2022 ACM Conference on Fairness, Accountability, and Transparency. Presented at the FAccT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency, ACM, Seoul Republic of Korea, pp. 1434–1445. https://doi.org/10.1145/3531146.3533200

Robbins, S., van Wynsberghe, A., 2022. Our New Artificial Intelligence Infrastructure: Becoming Locked into an Unsustainable Future. Sustainability 14, 4829. https://doi.org/10.3390/su14084829

Rohde, F., Gossen, M., Wagner, J., Santarius, T., 2021. Sustainability challenges of Artificial Intelligence and Policy Implications. Ökologisches Wirtschaften - Fachzeitschrift 36, 36≣40. https://doi.org/10.14512/OEWO360136

Rohde, F., Santarius, T., 2023. Emerging sociotechnical imaginaries – How the smart home is legitimized in visions from industry, users in homes and policymakers in Germany. Futures 151, 103194. https://doi.org/10.1016/j.futures.2023.103194

Rohde, F., Wagner, J., Meyer, A., Reinhard, P., Voss, M., Petschow, U., & Mollen, A. (2024). Broadening the perspective for sustainable artificial intelligence: sustainability criteria and indicators for Artificial Intelligence systems. Current Opinion in Environmental Sustainability, 66, 101411. https://doi.org/10.1016/j.cosust.2023.101411

Rozite, V., Bertoli, E., Reidenbach, B., 2023. Energy system Buildings Data Centres and Data Transmission Networks Data Centres and Data Transmission Networks. International Energy Agency. https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks (accessed 12.6.23).

Sætra, H.S., 2021a. AI in Context and the Sustainable Development Goals: Factoring in the Unsustainability of the Sociotechnical System. Sustainability 13, 1738. https://doi.org/10.3390/su13041738

Sætra, H.S., 2021b. A Framework for Evaluating and Disclosing the ESG Related Impacts of AI with the SDGs. Sustainability 13, 8503. https://doi.org/10.3390/su13158503

Schneider, J., 2022. Foundation models in brief: A historical, socio-technical focus. https://doi.org/10.48550/ARXIV.2212.08967

Sevilla, J., Heim, L., Ho, A., Besiroglu, T., Hobbhahn, M., Villalobos, P., 2022. Compute Trends Across Three Eras of Machine Learning, in: 2022 International Joint Conference on Neural Networks (IJCNN). Presented at the 2022 International Joint Conference on Neural Networks (IJCNN), IEEE, Padua, Italy, pp. 1–8. https://doi.org/10.1109/IJCNN55064.2022.9891914

Solaiman, I., Talat, Z., Agnew, W., Ahmad, L., Baker, D., Blodgett, S.L., Daumé, H., Dodge, J., Evans, E., Hooker, S., Jernite, Y., Luccioni, A.S., Lusoli, A., Mitchell, M., Newman, J., Png, M.-T., Strait, A., Vassilev, A., 2023. Evaluating the Social Impact of Generative AI Systems in Systems and Society. https://doi.org/10.48550/ARXIV.2306.05949

Strubell, E., Ganesh, A., McCallum, A., 2019. Energy and Policy Considerations for Deep Learning in NLP, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. Presented at the ACL 2019, Association for Computational Linguistics, Florence, Italy, pp. 3645–3650. https://doi.org/10.18653/v1/P19-1355

SustAIn Magazine 1, 2022. Sustainable AI in Practice. https://algorithmwatch.org/en/sustain-magazine-2022

SustAIn Magazine 2, 2023. AI and the Challenge of Sustainability. https://algorithmwatch.org/en/sustain-magazine-march-2023/

SustAIn Magazine 3, 2023. A Different Take on AI: We Decide What AI Has To Do for Us. https://algorithmwatch.org/en/sustain-magazine-november-2023/

Thompson, N.C., Greenewald, K., Lee, K., Manso, G.F., 2021. Deep Learning's Diminishing Returns: The Cost of Improvement is Becoming Unsustainable. IEEE Spectr. 58, 50–55. https://doi.org/10.1109/MSPEC.2021.9563954

Uddin, M., Rahman, A.A., 2012. Energy efficiency and low carbon enabler green IT framework for data centers considering green metrics. Renewable and Sustainable Energy Reviews 16, 4078–4094. https://doi.org/10.1016/j.rser.2012.03.014

Uptime Institute, 2021. 2021 Data Center Industry Survey Results.

van Wynsberghe, A., 2021. Sustainable AI: AI for sustainability and the sustainability of AI. AI Ethics 1, 213–218. https://doi.org/10.1007/s43681-021-00043-6

Verdecchia, R., Sallou, J., Cruz, L., 2023. A systematic review of Green AI. WIREs Data Min & Knowl 13, e1507. https://doi.org/10.1002/widm.1507

Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S.D., Tegmark, M., Fuso Nerini, F., 2020. The role of artificial intelligence in achieving the Sustainable Development Goals. Nat Commun 11, 233. https://doi.org/10.1038/s41467-019-14108-y

Vipra, J., Korinek, A., 2023. Market concentration implications of foundation models: The Invisible Hand of ChatGPT (Working Paper No. #9). Center on Regulation and Markets.

Vollhardt, S., Schmidt, K., Kask, S., Noga, M., 2021. Das intelligente Unternehmen: Effiziente Prozesse mit Künstlicher Intelligenz von SAP ≣ Wie Unternehmen die hohen Erwartungen an die KI erfüllen können, in: Buxmann, P., Schmidt, H. (Eds.), Künstliche Intelligenz. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 119≣137. https://doi.org/10.1007/978-3-662-61794-6_7

von Mörner, M., & Boltze, M. (2018). Sammelverkehr mit autonomen Fahrzeugen im ländlichen Raum: Zur Zukunft des ÖPNV in dünn besiedelten Gebieten. *NAHVERKEHR*, *36*(11).

Wagner, J., Rohde, F., Schmelzle, F., 2023: Sustainability expectations towards Artificial Intelligence in the energy and mobility sector, in: Proceedings of the STS Conference Graz 2023. Critical Issues in Science, Technology and Society Studies, 8 – 10 May 2023, Graz, Austria, pp. 347-362. https://openlib.tugraz.at/download.php?id=657c310e30295&location=browse

Whitehead, B., Andrews, D., Shah, A., Maidment, G., 2014. Assessing the environmental impact of data centres part 1: Background, energy use and metrics. Building and Environment 82, 151–159. https://doi.org/10.1016/j.buildenv.2014.08.021

Widder, D.G., West, S., Whittaker, M., 2023. Open (For Business): Big Tech, Concentrated Power, and the Political Economy of Open AI. SSRN Journal. https://doi.org/10.2139/ssrn.4543807

Williams, A., Miceli, M., Gebru, T., 2022. The Exploited Labor Behind Artificial Intelligence. Noema Magazine. https://www.noemamag.com/the-exploited-labor-behind-artificial-intelligence/ (accessed 12.6.23).

Yigitcanlar, T., Cugurullo, F., 2020. The Sustainability of Artificial Intelligence: An Urbanistic Viewpoint from the Lens of Smart and Sustainable Cities. Sustainability 12, 8548. https://doi.org/10.3390/su12208548

About the SustAIn Project

The project SustAIn developed criteria for the sustainability assessment of AI and used case studies to investigate the transformative potential of applications with a high maturity level in sectors of high relevance for sustainability goals such as Energy, Mobility and Online-Shopping. In order to strengthen the societal discourse and the development in terms of sustainability, dialogue processes for science, industry, civil society and politics ("Sustainable AI Labs") were organised and guidelines for sustainable AI development were set up. In our three SustAIn Magazines, we are promoting the debate on the sustainability impacts of AI.

www.algorithmwatch.org/en/sustain/

The Institute for Ecological Economy Research (IÖW)

The Institute for Ecological Economy Research (IÖW) is a leading scientific institute in the field of practice-orientated sustainability research. Around 70 employees develop strategies and approaches for a sustainable economy - for an economy that enables a good life and preserves the natural foundations. The institute works on a non-profit basis and without basic public funding. The IÖW is a member of the Ecological Research Network (Ecornet), the network of non-university, non-profit environmental and sustainability research institutes in Germany.

www.ioew.de/en/ | Twitter (X) | Mastodon | Newsletter

AlgorithmWatch

AlgorithmWatch is a non-profit organisation with the aim of observing and classifying processes of algorithmic decision-making that have social relevance - i.e. that either predict or predetermine human decisions or make decisions automatically.

We strive for a world in which technology in general and algorithmic systems in particular benefit people. The systems should make societies fairer, more democratic, more inclusive and more sustainable - be it in terms of ascribed origin and gender, racialisation, sexual orientation, age, class and wealth or resource consumption.

www.algorithmwatch.org/en/ | Twitter (X) | LinkedIn | Mastodon | Facebook | Instagram

Distribituted Artificial Intelligence Laboratory (DAI-Laboratory)

The DAI Laboratory at the TU Berlin sees itself as an intermediary between university research and industrial utilisation. With our interdisciplinary team, we generate innovations and transfer university research into everyday applications. This is done in close co-operation with other scientific and industrial institutions.

www.dai-labor.de/en/home/ | Twitter (X) | LinkedIn | Facebook | Instagram

Further Readings from the SustAIn Project

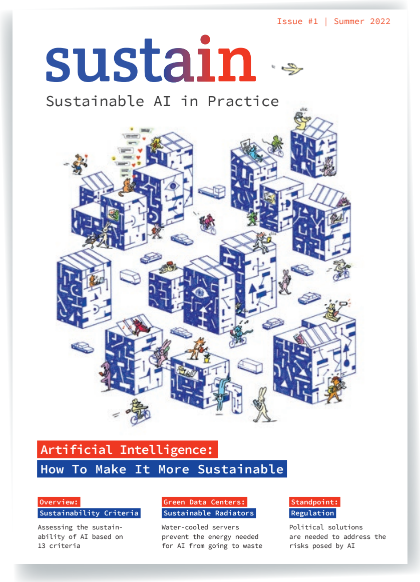

sustain
Sustainability Index
for Artificial Intelligence

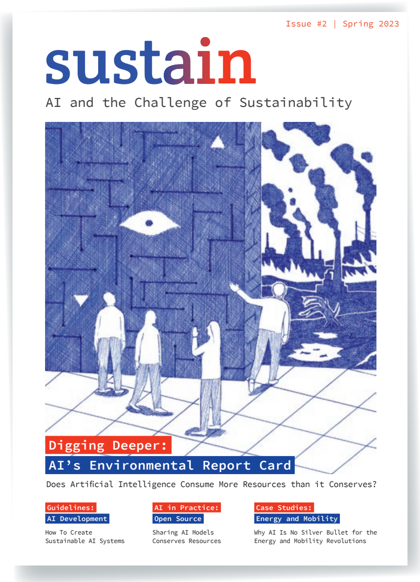

Project Partners:

ALGORITHM
WATCH

i ö w
INSTITUTE FOR ECOLOGICAL
ECONOMY RESEARCH

DAI-Labor
TU Berlin

Funding:

Federal Ministry
for the Environment, Nature Conservation,
Nuclear Safety and Consumer Protection