Tackling Climate Change with AI: Research Gaps and Needs to Align Machine Learning with Greenhouse Gas Reductions

Jan Bieser
Institute Public Sector Transformation
Business School
Bern University of Applied Sciences
Bern, Switzerland
jan.bieser@bfh.ch

Abstract— **Machine learning (ML) promises to revolutionize our socio-economic landscape, yet its impacts on greenhouse gas (GHG) emissions and strategies to harness ML for climate protection are not well understood. This discussion paper reviews key research on ML's GHG effects, highlighting significant research gaps and needs for a climate-oriented ML transformation. The results show that research on GHG emissions caused during model development, training, and operation is progressing. However, there is no comprehensive overview of effective measures to reduce them along the entire ML software and hardware life cycle. (Industrial) research on the GHG effects of ML applications focuses mainly on GHG reduction potentials while neglecting the possibility that ML applications also increase emissions. Thus, research in at least three key areas is needed to align ML with GHG reductions. First, robust methods to assess and report the GHG impacts of ML models and applications are required to systematically compare them and identify best practices. Second, comprehensive GHG assessments at every effect level are essential to identify measures to increase the GHG efficiency of ML models and exploit their climate protection potential. Third, analysing ML business models is crucial to propose measures that incentivize ML providers and users to reduce GHG emissions. Addressing these issues is essential for mindfully steering ML toward GHG reductions. Otherwise, there is a risk that the GHG footprint of ML will skyrocket, that ML applications will primarily accelerate GHG-intensive activities, and that an opportunity for decoupling (economic) growth and GHG emissions will be missed.**

Keywords—machine learning, artificial intelligence, greenhouse gas, emissions, climate protection, climate change

I. INTRODUCTION

Many digital technologies, such as mobile phones or social media, have had the steepest technology adoption curves in history, paving the way for the age of the so-called digital acceleration [1], [2]. Machine learning, the ability of computers to automatically learn and adapt without explicit programming [3], is expected to accelerate this trend even more. ML shortens digital applications' development cycles, expanding their capabilities and application domains [4]. The new opportunities provided by ML are expected to fundamentally change our patterns of production and consumption, potentially initiating a new socio-economic development path [5]. Given these expectations about ML's transformative impact, the pivotal question lingers: Will the new possibilities present additional challenges or viable solutions to the most pressing challenge of the 21st century: climate protection?

Despite the urgent need to reduce global greenhouse gas (GHG) emissions by 43% compared to 2019 to meet the Paris Agreement, emissions continue to rise [6]. Many companies involved in the development of ML spread the hope that it will pave the way for significant GHG emission reductions by facilitating the dematerialization of physical processes, enhancing resource and energy efficiency, or closing material cycles [7], [8]. The President of the European Commission, Ursula von der Leyen, captured this sentiment in her 2023 State of the Union Address: "The same should be true for artificial intelligence. It will improve healthcare, boost productivity, address climate change" [9, p. 1].

Such hopes, however, are not based on scientific studies on the positive and negative climate impacts of ML. On the one hand, research on the energy consumption and GHG emissions caused by the development, training, and operation of ML models is progressing. There is a growing body of GHG assessments for different types of ML models, which cover an increasing scope of GHG effects along the entire life cycle of ML models and the required hardware [10], [11], [12], [13], [14]. Some academic studies and media articles also extrapolate the results from individual case studies to an aggregate level based on the sales figures of graphical processing units (GPUs) or the number of user requests for popular models such as Open AI's ChatGPT [15] to emphasize the relevance at a societal level. Such studies led to several popular media articles warning about the growing energy use and GHG footprint of ML [15], [16]. However, there are also studies predicting a softer growth in ML's future energy use, still recognizing that a potential future pervasive use of ML models poses a risk [17].

On the other hand, there are only a few studies on the impacts of ML applications on other sectors (e.g., transport, buildings, agriculture) and the consequences for GHG emissions [18]. Although the GHG impacts of some applications, such as automated driving, have already been intensively studied [19], [20], this remains the exception. Many of the available (industrial) studies are characterized by narrow system boundaries and tend to overlook crucial climate risks associated with ML. For example, Rolnick et al. [8], Microsoft and PwC [21], and the Capgemini Research Institute [22] review the GHG effects of ML applications, focusing mainly on the potential for GHG reductions.

However, there are empirical studies on the climate impacts of digital technologies in general (without a special focus on ML), and their results challenge the optimism of the ML industry and politics. In fact, the most comprehensive and rigorous studies of digitalization's climate impacts indicate that digital technologies led, at best, to a marginal reduction of GHG emissions or that GHG-increasing and -reducing effects cancel each other out [23], [24]. Reasons for that are that digital technologies are resource- and energy-intensive in

production and operation and often intensify GHG-intensive activities due to rebound effects [25], [26].

While digital technologies possess significant potential for climate protection, this potential remains largely untapped because it would need purposeful action. A similar scenario may unfold for the special case of ML unless there is enough effort to actively guide its implementation in a manner consistent with climate protection goals. A prerequisite, therefore, is to systematically develop a nuanced understanding of the interplay between ML and GHG emissions.

Berkhout and Hertin [27] first conceptualized the environmental impacts of digital technologies in a framework that was later adapted several times [28], [29], [30] and most recently applied to the specific case of ML and climate protection [18], [31]. The framework distinguishes three types of ML effects on GHG emissions:

- *Technology effects* describe the GHG emissions that arise during the provision, operation, and disposal of ML software and the required hardware.

- *Application effects* describe the impacts of individual ML applications and how they increase or decrease GHG emissions in other sectors.

- *Systemic effects* describe fundamental changes in economic structures and lifestyles arising from the wider penetration of ML applications in society and the consequences for GHG emissions.

Technology effects are also called direct effects, compute-related effects, or GHG footprints, whereas application and systemic effects are also called indirect effects or GHG handprints. The aim of this discussion paper is to summarize relevant research along this framework in order to identify critical research gaps and needs that have to be addressed to generate the knowledge required for a climate-oriented ML transformation.

Following the best practices for text recycling [32], I acknowledge that this paper is based on an unpublished research grant proposal previously submitted to the Swiss National Science Foundation. It has been extended with additional analysis and discussions.

II. METHOD

To identify relevant research, I searched Google Scholar and Google for various combinations of search terms for machine learning and GHG emissions and related terms (e.g., artificial intelligence, AI, climate, climate protection, GHG reduction, energy consumption). I used Google Scholar because it is the most comprehensive database of academic publications [33], and Google, because it allows to identify relevant grey literature, such as company reports. I only included studies that explicitly address the technology, application, or systemic impacts of machine learning (or AI in general) on GHG emissions, excluding those focusing on ML's role in adapting to climate change or on other environmental impacts beyond GHG emissions. Partly, studies on the energy use of ML are also taken into account because ML's energy use is a main source of GHG emissions in the development and operation of ML models. The references considered include scientific studies, reports from international organizations like the OECD, and industry analyses.

I structured the identified literature into research on technology, application, and systemic effects, summarized, conceptualized, and discussed their findings. Based on this analysis, I outlined research gaps, methodological limitations, and derived necessary fields for further research to support the development and use of machine learning in ways that support GHG reductions. Partly, I also build on the literature on the GHG impacts of digital technology in general, but only if it is relevant to the specific case of ML.

It is important to note that this is not a systematic literature review whose ambition is to identify all available sources. There are additional potentially relevant sources not covered in this article. For example, ML researchers are intensively concerned with improving algorithms from a technical perspective and thereby explicitly or implicitly address energy use and GHG emissions of ML models. There are also many studies that deal with the optimization of specific technologies and applications (e.g., solar systems) using ML, which also influence GHG emissions. These studies are not considered. Instead, the focus is on research that primarily addresses machine learning through the lens of GHG emissions and climate protection.

III. TECHNOLOGY EFFECTS

Technology effects describe the GHG impacts caused throughout the life cycle of ML models. These are not caused by the ML model itself but throughout the life cycle of the hardware required to develop, train, fine-tune, deploy (or infer), and phase out ML models. Fig. 1 shows the key life cycle stages of ML models and the environmental impacts caused by the required hardware.

A systematic literature review by Veredecchia et al. [34] showed that most studies on ML's energy use and GHG emissions focus on the *(final) training phase*. The computational effort required to train large ML models has increased rapidly in recent years. Mehonic and Kenyon [16] and Amodei and Hernandez [35] showed that computational cost for training large ML models doubled every 24 months until 2012, every 3.4 months after the introduction of GPUs in 2012, and every 2 months since the increasing adoption of transformer models during 2019 which are the basis of the prominent large language models (LLMs) such as the GPT-models of OpenAI. The foundation for this development were rapid increases in ML hardware performance and energy efficiency [36]. For example, the performance of NVIDIA GPUs increased by a factor of over 300 between 2012 and 2021, greatly exceeding Moore's Law [16]. Further significant increases in demand and performance of hardware for training ML models are expected in the near future. However, physical constraints limit the possible increases in digital computing systems' performance and energy efficiency that are based on the conventional von Neumann architecture. Therefore, various innovative approaches are currently being explored, such as quantum computing or brain-inspired computing by "co-locating memory and processing, encoding information in a wholly different way or operating directly on signals, and employing massive parallelism" [16, p. 255]. Their future impact on energy use remains uncertain.

Some factors driving training GHG emissions are already known. Table I shows the results of four studies on energy use and GHG emissions in the *final training phase* of LLMs. The results suggest that the model size in terms of the number of
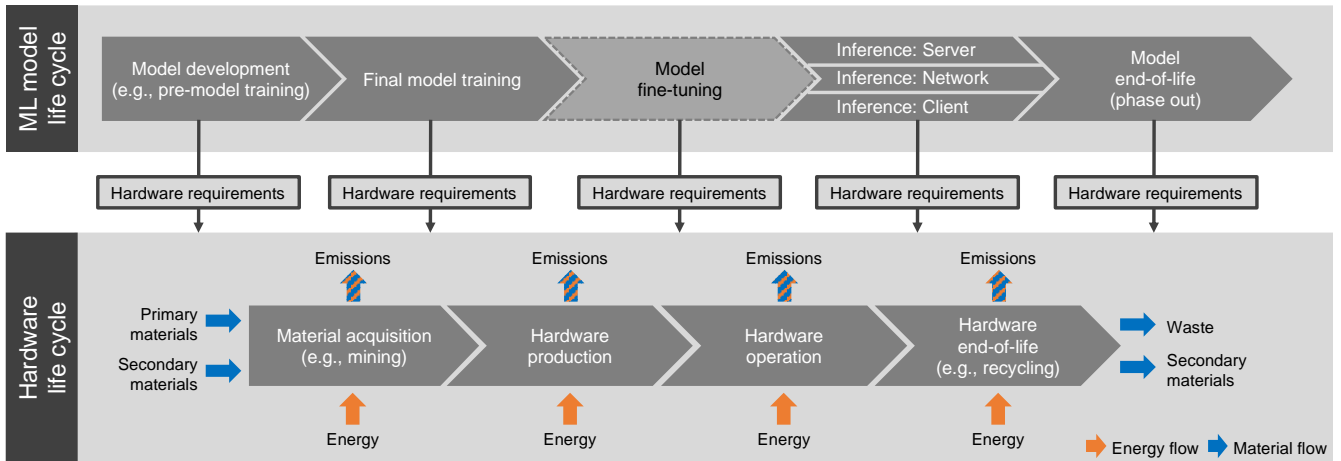
Fig. 1. The life cycle of ML models and the environmental impacts (material/energy flows, emissions) caused by the required hardware, which can differ by ML model phase based on [12], [37]. Emissions can be physical waste (e.g., chemicals used in hardware production) and gases (e.g., CO_2). Fine-tuning describes adapting a pre-trained model to a specific task. This only takes place for some ML models (e.g., foundation large language models, LLMs) and is therefore illustrated in dashed lines. The inference phase is the live operation of the model, in which clients (users) send requests to the model, which processes them. During inference, energy is required for the operation of the model on the server side, for operating the client's end-user device (e.g., laptop computers), and for the data transmission networks.

parameters affects energy use, and the GHG intensity of the electricity additionally impacts GHG emissions. However, the energy use of GPT-3 and Gopher, as well as OPT and BLOOM, differ significantly, even though the model sizes are of a similar order of magnitude, suggesting that other factors are also important. Other research has also shown that model size alone is not a good indicator of ML energy use. For example, Wu et al. [39, p. 4] showed that (considering training and inference) "the Switch Transformer model equipped with 1.5 trillion parameters […] produces significantly less carbon emission than that of GPT-3 (750 billion parameters)." Thus, the energy use of ML model training depends on other factors that need further exploration, such as the model architecture, the frequency of re-training, or the type of hardware equipment used and its utilization rate [12], [31], [34], [39], [10]. Various research efforts are ongoing in this field [34], [40], [18], [17], [41]. In general, it should be noted that many of the statements about ML's rising energy use refer primarily to very large models such as LLMs. ML includes many more applications that are significantly less energy-intensive.

In addition to the final training phase, other ML life cycle phases are crucial for energy use and GHG emissions. First, *model development*, which involves the step-wise development, evaluation, and improvement of (pre-)models, can be more energy-intensive than final model training if the number of pre-models and training runs is high [31]. Specifically for neural network models, the number of possible model configurations and, thus, of training runs for pre-models can be very high [31]. A lot of research effort goes into strategies to optimize model development, e.g., by

improving hyperparameter tuning and reducing the number of pre-model trainings [34], [42].

Second, even though one user request in the *inference phase* does not require a lot of energy, the energy consumption in absolute terms can exceed that of the final training phase if the number of user requests is very high [31]. For example, De Vries [15] estimated that if Google would include "Generative AI" capabilities in its standard keyword search, server-side energy use during inference alone could amount to up to 29.3 TWh annually in a worst-case scenario, more than the electricity consumption of countries like Portugal or Hungary. However, this scenario is unlikely in the near future due to limited GPU availability and high cost [15]. Desislavov et al. [17] provide further analysis and trends of ML inference energy use and argue for a "softer growth" in energy use than other studies.

Third, *material acquisition and hardware production* are resource and energy-intensive processes that cause large amounts of GHG emissions, as modern ICT hardware contains more than 50% of the elements in the periodic table [43]. These emissions are often called embodied emissions. Wu et al. [39] showed that embodied emissions account for roughly 30% of emissions compared to 70% for training and inference across six ML models from Facebook. Luccioni et al. [12] came to similar results. Beyond, training and inference are often conducted in modern data centers powered by renewable, low-carbon electricity. In that case, embodied emissions account for an even higher share of life cycle emissions; however, they are not higher in absolute terms [39].

TABLE I. ENERGY USE AND GHG EMISSIONS CAUSED BY THE FINAL TRAINING OF FOUR LLM MODELS [12], [13], [14], [38] (COMPARISON CITED FROM [12]). DATA CENTER PUE (POWER USAGE EFFECTIVENESS) IS AN INDICATOR OF THE ENERGY EFFICIENCY OF THE HEATING, VENTILATION, AND COOLING (HVAC) SYSTEMS IN DATA CENTERS. THE LOWER THE VALUE, THE MORE EFFICIENT THE HVAC SYSTEM IS, WITH 1 BEING THE BEST POSSIBLE FACTOR.

Model name	Model size [Billions of parameters]	Datacenter PUE factor	GHG intensity of electricity mix [g CO2e/kWh]	Training energy consumption [MWh]	GHG emissions [tons CO2e]
GPT-3	175	1.1	429	1,287	552
Gopher	280	1.08	330	1,066	380
OPT	175	1.09	231	324	76
BLOOM	176	1.2	57	433	30

Hardware end-of-life is usually not considered [18] but may reduce life cycle GHG emissions if the hardware is recycled and the recovered materials are used to reduce demand for primary materials [37]. A challenge here is to increase recovery rates and reduce recycling costs further to increase the economic attractiveness of recycling compared to the acquisition of primary materials [44].

Berthelot et al. [10] also estimated emissions (and energy use) caused by *clients' end-user devices* and *data transmission networks* during inference for the text-to-image generator Stable Diffusion, and showed that these emissions matter (however, their results do not allow to distinguish between emissions caused by the production and operation of client devices). They further argue that breakthroughs in AI could increase ICT usage in general and thus the technology effects caused by the entire ICT sector.

We could not find a study that considered the impacts of *disposing of ML software at its end of life*. Disposing of ML models (e.g., creating back-ups and phasing them out of inference) likely causes little GHG emissions because it is less computing-intensive than other phases and only needs to be conducted once. However, as more and more models and training data are produced, energy use for data storage will further increase. Thus, practical criteria about when to delete models and data are required for ML and digital data in general [45].

ML models can have further unintended environmental impacts that are underexplored. A multi-year project has identified 11 criteria for sustainable software products (Table II), out of which only three concern the direct energy and resource requirements of the hardware used to run the software [46]. Four criteria address the impact of the software on the service lifetime and renewal cycles of hardware (e.g., due to a lack of backward compatibility or platform independence), and four criteria concern the software user's autonomy concerning the future use of software products and data (e.g., lock-in effects and the user's freedom to decide how collected data is used). ML models can have different design characteristics regarding these criteria and unintended consequences for GHG emissions. For example, if a hospital wants to use a new ML software to diagnose diseases, it is not clear what impact the design of the ML software, its interfaces, and its hardware requirements will have on the lifespan of the hospital's existing ICT infrastructure. In a worst-case scenario, an unfavourable design could result in the hospital being forced to renew large parts of its ICT infrastructure. For the climate-friendly provision of ML models, it is essential to understand the characteristics and

TABLE II. CRITERIA FOR SUSTAINABLE SOFTWARE PRODUCTS [46].

Resource efficiency	Hardware lifetime	User autonomy
• Hardware efficiency • Energy efficiency • Resource management	• Downward compatibility • Platform independence and portability • Hardware sufficiency	• Transparency and interoperability • Deinstallability • Warranty functions • Independence of other resources • Quality of product information

design options of ML models concerning these criteria and what significance they have for total GHG emissions.

Berthelot et al. [10] also showed at the example of a text-to-image generator that environmental impact categories beyond GHG emissions and energy use are relevant, such as abiotic depletion: Client, network and servers have an "impact on metal scarcity equivalent to the production of 5659 smartphones" (p. 6). Li et al. [47] estimate "that training GPT-3 in Microsoft's state-of-the-art U.S. data centers can consume a total of 5.4 million liters of water […]. Additionally, GPT-3 needs to "drink" (i.e., consume) a 500ml bottle of water for roughly 10-50 responses, depending on when and where it is deployed" (p. 3).

IV. APPLICATION EFFECTS

Application effects describe the changes triggered by individual ML applications in other sectors and their consequences for GHG emissions. While technology effects per se increase GHG emissions, application effects can increase or decrease them [31]. For example, ML simulation models can be used to accelerate the development of renewable energy technologies or to increase the efficiency in oil and gas exploration [31]. Table III provides a simplified overview of the impact mechanisms through which ML applications can increase or reduce GHG emissions.

To date, academic literature primarily focuses on the GHG-reducing effects of ML applications [18]. Rolnick et al. [8] conducted a comprehensive review of ML applications with the potential for GHG reductions and for facilitating climate change adaptation. They identified and described applications in 13 domains, such as transportation, buildings and cities, farms and forests, education, and finance. Kaack et al. [31] identify certain capabilities or roles that ML applications can provide to enable GHG reductions in other sectors. These include data mining and remote sensing, accelerated experimentation, fast approximation, forecasting, systems optimization and control, and predictive maintenance. These capabilities can be used to design and monitor climate

TABLE III. GHG IMPACT MECHANISMS OF ML APPLICATIONS BASED ON [25], [31]. PLEASE NOTE THAT OTHER CONCEPTUALIZATIONS OF THE IMPACT MECHANISMS EXIST. REBOUND EFFECTS IN PARTICULAR CAN BE FURTHER DIFFERENTIATED.

ML impact mechanism	Description	GHG-reducing example application	GHG-increasing example application
Substitution	Replacement of a conventional service with an ML-based service	ML-based on-demand public transport substituting car transport	ML-based web search replacing the conventional keyword search
Optimization	ML-based process optimization	ML simulation models accelerating the development of low-carbon materials	ML-based optimization and acceleration of oil extraction
Induction	ML use stimulating the use of another resource	n/a by definition	ML-optimized personalized advertisement increasing consumption (and thus production) of physical products
Rebound	ML-induced efficiency increases and price reductions of a service increasing consumption of the same or other services	n/a by definition	ML-based optimization of the fuel and cost efficiency of cars increasing total car transport

policies, accelerate the development of low-carbon technologies, and efficiently plan, design, and operate (real-world) systems in ways that lead to GHG reductions.

These studies also address the fact that ML applications can lead to an increase in GHG emissions. Rolnick et al. [8] state that "ML is only one part of the solution; it is a tool that enables other tools across fields." (p. 59), implicitly recognizing that ML applications are just a means to an end, whereas the end can also foster GHG-intensive activities. Kaack et al. [31] explicitly describe that ML capabilities can be used to accelerate GHG-intensive activities. For example, remote sensing and systems management can also be used to increase yield in cattle farming or intensify oil and gas exploration.

Industrial studies mainly focus on the GHG-reducing potential of ML applications. Some of them even quantify it. For example, a study commissioned by Microsoft and conducted by PwC predicted that by 2030, ML applications will help avoid more than 4% of global GHG emissions [21]; BCG estimates that ML could avoid 5 to 10% of global GHG emissions already today [48]; the Capgemini Research Institute [22] survey showed that companies in 13 countries already achieved average GHG reductions of 13% between 2017 and 2020 with ML. Please note that these studies use the more general term artificial intelligence (AI) but also include ML.

Although ML has great climate protection potential, the results of these studies do not reflect the actual GHG effects of ML applications in a real-life setting [18]. This is because they largely ignore GHG-increasing applications, and the assessments face many methodological challenges and limitations, which were outlined (for digital applications in general) by Bieser and Hilty [49], Coroamă et al. [50], and Bergmark et al. [51]. These include, for example:

- *Selection of applications*: It is impossible to analyse all ML applications that are potentially relevant. Thus, studies can only investigate a subset of ML applications. The choice of applications influences the results [49].

- *Baseline estimation*: To estimate the impact of an ML application, two scenarios, before (baseline) and after the introduction of the ML application, must be compared. Once the ML application has been introduced, the baseline scenario can no longer be empirically observed; it is hypothetical [50]. The more an ML application has penetrated society, the harder it is to isolate its effects and estimate the baseline without ML adoption [50]. For example, if ML-based personalized shopping assistants take over consumer purchase decisions, it will be very difficult to imagine how shopping behaviour would have evolved without it.

- *Impact estimation*: There are only a few empirical case studies about the GHG impacts of specific ML applications. Determining the actual GHG impacts in a real-life setting is challenging because "theoretical potentials materialize only under specific conditions" [52, p. 1]. For example, an ML-based bus service that dynamically adjusts routes to demand might replace private car trips in rural areas and conventional public transport or bike trips in urban areas.

- *Extrapolation*: Even if an empirical case study is available, its results usually cannot simply be extrapolated to a larger population because the case study might not be representative [49]. For example, Malmodin and Coroamă [53] showed that in case studies of energy savings through smart meters, lower savings are observed in case studies with larger populations.

- *Unintended side effects*: Even if a ML application reduces GHG emissions, savings can be offset by unintended side effects such as rebound effects. Due to ICT's "exceptional dynamics of innovation and diffusion", "diverse and complex impact patterns" [54, p. 826] and "the complexity of social and ecological systems" [49, p. 77] such effects are difficult to predict, and even harder to quantify. For example, automated driving technology increases GHG efficiency in road transport. However, it could also increase congestion if people switch from public transport to car transport or let cars circulate when they cannot find a parking space [55].

For these reasons, it is very challenging to determine the actual GHG effect of ML applications, and existing quantitative studies provide only a partial view of the overall effects. Most importantly, since GHG-increasing effects are systematically underexplored, ML's climate protection potential is overestimated. This can even hinder climate protection, as the results of overoptimistic studies nurture the belief that ML applications will automatically lead to GHG reductions without targeted measures [25]. Today, 87% of "climate and AI leaders" expect ML to help combat the climate crisis [56].

The ICT and Sustainability research community was at a similar point concerning the climate protection potential of digital applications (without ML functionality) about ten years ago. ICT industry associations such as the Global Enabling Initiative (GeSI) predicted that digital applications will avoid up to 20% of global GHG emissions by 2030 [57], supporting the notion that digitalization is a silver bullet for climate protection. Today, it is known that these impacts have not materialized as the climate-protecting impacts of digital applications were overestimated [23], [49].

If we do not put the GHG impact assessment of ML applications on a more solid methodological basis, there is a high risk that we fall into the same trap again and do not systematically exploit ML's potential for climate protection. To avoid such a scenario, we must improve our understanding of the diverse, often unintended impact mechanisms of ML applications on GHG emissions. Only then can we deliberately pinpoint climate-protecting and -damaging applications and systematically steer them toward GHG reductions. Researchers have developed frameworks for assessing the positive and negative effects of digital applications [58], [59] and have started applying them to ML applications [18].

V. SYSTEMIC EFFECTS

Systemic effects describe fundamental changes in economic structures and lifestyles due to widespread ML adoption and the consequences for GHG emissions [28], [31]. Just like application effects, these can increase and reduce GHG emissions. For example, ML-supported driving automation, optimization of traffic flows, and mobility

services such as automatic on-demand buses could fundamentally change travel habits, the transport system, and its GHG footprint [60]. Such systemic effects emerge through the causal relationship between many variables and can only be observed from a higher system level over a longer time period [49]. They are thus hard to quantify but can significantly outweigh ML application effects and "are extremely important to consider when evaluating ML use cases" [31, p. 522].

To date, there are no comprehensive studies on the systemic GHG effects of ML. First, there are no *empirical observations* yet because ML adoption has only increased recently. Second, *predictive quantitative studies* of systemic GHG effects exist only for a few ML applications. For example, some studies have assessed the interaction between ML-based automated driving technology, transport mode choice, and transport demand, and warn about its climate-damaging effects if the technology is used primarily to increase comfort in individual motorized (car) transport and not public transport [20]. Third, there are a some *qualitative descriptions* of the systemic GHG effects of ML. For example, ML-based technology for automated driving could strengthen the position of heavy trucks and make it more difficult for other low-carbon technologies to gain traction in the freight transport market [31], [61]. However, such studies do not indicate the sizes of these effects in terms of GHG emissions.

For digital applications in general (without specific ML functionality) some studies exist, most of which applied regression analysis to investigate the relationship between ICT adoption and environmental indicators [24], [62], [63]. For example, Clausen et al. [64] showed that in nine OECD countries, a one percent increase in ICT capital was associated with a 0.110% decrease in energy demand between 2000 and 2014. Schulte et al. [65] observed a decrease of 0.235% in a similar study. The results of these studies indicate that digitalization has, at best, only marginally reduced energy use or GHG emissions. The inherent limitation of these studies is that they treat the causal chains through which digital applications lead to systemic effects as a black box [62]. To intentionally guide systemic effects toward GHG reductions, it is required to uncover the causal chains between ML adoption and GHG emissions, e.g., with complex system modelling techniques [66]. One study tried to do so with a system dynamics model of the prospective effects of ICT on environmental sustainability in the EU for the year 2020 from a 2006 perspective [66]. For example, the model revealed that "ICT applications that make passenger transport more time efficient [...] will create a rebound effect leading to more traffic and possibly more energy consumption" or that "an ICT-supported product-to-service shift" provides high theoretical climate protection potentials (p. 1628). It should be considered that the study is outdated, and a similar study has never been conducted again, neither for digital technologies in general nor for ML specifically.

VI. CONCLUSION

A. Summary of research gaps

We are still far away from holistically understanding the drivers and sizes of ML effects on GHG emissions and suitable GHG reduction measures in practice. The validity of the results of existing studies is limited due to narrow system boundaries and unaddressed methodological challenges. Table IV summarizes the most important research gaps by effect category.

Closing these research gaps is essential for mindfully steering the use of ML toward GHG reductions. Otherwise, there is a high risk that the GHG footprint of ML will skyrocket, that ML applications will primarily accelerate GHG-intensive activities, and that an unprecedented opportunity for decoupling (economic) growth and GHG emissions will be missed.

In fact, it is very likely that the capabilities provided by ML further elevate the role of digital technology in society, increasing its usage and, thus, the GHG footprint of the entire ICT sector. First, ML accelerates the development speed of ICT applications (e.g., by using LLM-based support tools in software development). The number of applications available and the problems they can solve will likely increase. Second, ML expands the human-machine interface, making digital technology usable in more everyday situations (e.g., while driving or cooking). For example, people can now interact with digital technology not only via screens but also through other interfaces (e.g., language and gestures). However, not addressing the energy and GHG challenge of ML might also slow down ML progress due to high energy costs and potentially increasing carbon prices.

TABLE IV. OVERVIEW OF MAIN RESEARCH GAPS BY EFFECT CATEGORY

Effect category	Research gaps
Technology effects	• No study has considered all ML software and hardware life cycle stages. Specifically, model development, fine-tuning, inference, software, and hardware end-of-life are underexplored. • There is no comprehensive overview of factors, as well as hardware and software design options, that impact the GHG emissions caused by ML models. • Unintended ML impacts on hardware lifetime (and user autonomy) are underexplored.
Application effects	• There is only little knowledge about the (prospective) ML applications with the highest potential for GHG reductions, and those with the highest risks for GHG increases. The most effective measures to promote GHG-reducing and mitigate GHG-increasing effects are not known. • Existing quantitative GHG assessments of ML applications tend to focus on their potential to reduce GHG emissions while neglecting the possibility that ML applications also increase emissions. • Quantitative GHG assessments also suffer from methodological limitations, including issues with baseline and impact estimation, extrapolation of case study results, and failure to account for rebound effects.
Systemic effects	• There is minimal knowledge about the causal mechanisms through which ML influences the GHG intensity of each economic sector and how these effects can be steered (e.g., with political measures) in a desired direction. • The little available knowledge on these effects primarily relies on qualitative descriptions. The sizes of these effects that need to be determined using quantitative methods are not known.

B. Fields for further research

To avoid the scenario described above, it is crucial that future research on the GHG impacts of ML improves the methodological foundation of assessments, conducts comprehensive GHG assessments of ML models and applications, and explores how ML developers and users can be incentivized to use ML for GHG reductions. These three research directions that apply to all effect categories are described in more detail below.

Creating sound and robust GHG assessment methods and reporting guidelines: The methodological basis for assessing technology, application, and systemic effects must be advanced so that developers and researchers have the tools to systematically assess and optimize the GHG impacts of ML models and applications. The methods must include techniques for identifying unintended GHG effects (such as rebound effects), which occur at the technology, application, and system effects levels. These guidelines should also suggest measures to deal with uncertainty and consider that the level of accuracy required depends on the assessment's purpose. For example, optimizing an existing ML model usually requires more granular data than estimating a theoretical future potential to avoid GHG emissions with an ML application.

Additionally, standardized KPIs for reporting GHG impacts are required to enable systematic comparisons of GHG impacts and the targeted identification of best practices [40]. Suggestions are already being developed [31]. This groundwork is particularly important as ML technology, its capabilities, and applications are rapidly evolving.

Conducting sound GHG assessments: Actual GHG assessments need to be carried out at every effect level. Concerning technology effects, further end-to-end life cycle GHG assessments of ML models are required to identify those factors that contribute most to ML model GHG emissions and design options to reduce them. These studies also need to consider target conflicts between GHG-reducing design options and their impact on model performance or precision [34].

Technology effects can also be reduced if ML is not used at all in the first place. As De Vries [15, p. 2194] suggested, to reduce ML's energy use, developers should not only "focus on optimizing AI, but also [...] critically consider the necessity of using AI in the first place, as it is unlikely that all applications will benefit from AI [...]." Thus, developers should also have guidelines to recognize when the application of ML is valuable and when it is superfluous for a particular application.

Regarding application effects, assessments should focus on identifying those ML applications with high potential for reducing and increasing GHG emissions, including the conditions under which GHG reductions materialize and increases can be mitigated. Concerning systemic effects, the assessments should focus on understanding the dynamic interaction between ML, the socio-economic system, and GHG emissions to derive effective (political) measures to steer them in a desired direction.

Even though this article primarily focuses on GHG emissions and partly on energy use, ML also affects other environmental indicators that should be addressed as well (e.g., material use for producing server equipment or water use for cooling data centres [10], [47]).

Climate-oriented ML business models: A promising approach is also to incentivize ML developers and users to use ML in ways that reduce GHG emissions or at least avoid incentives that increase them. For instance, many digital platforms' business models contain incentive structures that increase consumption and GHG emissions [23]. The reason is that revenues depend on the number of user clicks on advertisements. Consequently, platform providers strive to extend user engagement through manipulative techniques such as addictive design [67] and enhance advertising effectiveness via personalization [68]. To date, it is not clear what future ML business models will look like and what consequences the choice of business model will have for GHG emissions. Thus, it is necessary to analyse potential ML business models regarding their incentive structure and propose (new) business models and other measures that provide incentives for providers and users for GHG reductions.

Addressing all research needs is challenging because assessing the effects on each level requires the development and application of distinct assessment methods and knowledge from various disciplines. Investigating technology effects requires technical knowledge about ML methods and the capability to abstract them into process models that can be investigated with so-called environmental life cycle assessments (eLCAs). Assessing application effects requires knowledge about the drivers of GHG emissions, about (prospective) ML applications, and their impacts in each economic sector. Investigating systemic effects demands expertise in both ML technology and application effects, coupled with the ability to integrate them using complex systems modelling techniques. Given the rapid advancements in ML technology and applications, all assessments need to account for uncertainty about future developments and for data unavailability. The diversity of the challenges requires interdisciplinary collaboration because they can only be solved by combining expertise and methods from various specialist areas.

REFERENCES

[1] ITU and UNDP, "SDG Digital Acceleration Agenda," 2023. Accessed: Nov. 30, 2023. [Online]. Available: https://www.undp.org/sites/g/files/zskgke326/files/2023-09/SDG%20Digital%20Acceleration%20Agenda_2.pdf

[2] Our World in Data, "Share of United States households using specific technologies." Accessed: Nov. 30, 2023. [Online]. Available: https://ourworldindata.org/grapher/technology-adoption-by-households-in-the-united-states

[3] A. L. Samuel, "Some Studies in Machine Learning Using the Game of Checkers," IBM Journal of Research and Development, vol. 3, no. 3, pp. 210–229, 1959, doi: 10.1147/rd.33.0210

[4] J. Bieser, "Creative through AI," Gottlieb Duttweiler Institute, 51, 2022, doi: 10.59986/CCHA2271

[5] S. Clarke and J. Whittlestone, "A Survey of the Potential Long-term Impacts of AI: How AI Could Lead to Long-term Changes in Science, Cooperation, Power, Epistemics and Values," in Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society, in AIES '22. New York, USA, 2022, pp. 192–202, doi: 10.1145/3514094.3534131

[6] IPCC, "Summary for Policymakers," in Climate Change 2023 Synthesis Report, pp. 1–34, 2023, doi: 10.59327/IPCC/AR6-9789291691647.001

[7] R. Agrawal, V. A. Wankhede, A. Kumar, S. Luthra, A. Majumdar, and Y. Kazancoglu, "An Exploratory State-of-the-Art Review of Artificial Intelligence Applications in Circular Economy using Structural Topic Modeling," Oper Manag Res, vol. 15, no. 3, pp. 609–626, 2022, doi: 10.1007/s12063-021-00212-0

[8] D. Rolnick et al., "Tackling Climate Change with Machine Learning." arXiv, 2019, doi: 10.48550/arXiv.1906.05433

[9] European Commission, "2023 State of the Union Address by President von der Leyen." Accessed: Dec. 08, 2023. [Online]. Available: https://ec.europa.eu/commission/presscorner/detail/ov/speech_23_4426

[10] A. Berthelot, E. Caron, M. Jay, and L. Lefèvre, "Estimating the environmental impact of Generative-AI services using an LCA-based methodology," presented at the CIRP LCE 2024 - 31st Conference on Life Cycle Engineering, 2024, p. 1. Available: https://inria.hal.science/hal-04346102

[11] A. Lacoste, A. Luccioni, V. Schmidt, and T. Dandres, "Quantifying the Carbon Emissions of Machine Learning." arXiv, 2019, doi: 10.48550/arXiv.1910.09700

[12] A. Luccioni, S. Viguier, and A.-L. Ligozat, "Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model," Journal of Machine Learning Research, vol. 24, no. 253, pp. 1–15, 2023. Available: https://jmlr.org/papers/v24/23-0069.html

[13] D. Patterson et al., "Carbon Emissions and Large Neural Network Training." arXiv, 2021, doi: 10.48550/arXiv.2104.10350

[14] J. W. Rae et al., "Scaling Language Models: Methods, Analysis & Insights from Training Gopher." arXiv, 2022, doi: 10.48550/arXiv.2112.11446

[15] A. De Vries, "The growing energy footprint of artificial intelligence," Joule, vol. 7, no. 10, pp. 2191–2194, 2023, doi: 10.1016/j.joule.2023.09.004.

[16] A. Mehonic and A. J. Kenyon, "Brain-inspired computing needs a master plan," Nature, vol. 604, no. 7905, Art. no. 7905, 2022, doi: 10.1038/s41586-021-04362-w

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

[18] A.-L. Ligozat, J. Lefevre, A. Bugeau, and J. Combaz, "Unraveling the Hidden Environmental Impacts of AI Solutions for Environment Life Cycle Assessment of AI Solutions," Sustainability, vol. 14, no. 9, Art. no. 9, 2022, doi: 10.3390/su14095172

[19] S. Hörl, F. Becker, T. Dubernet, and K. Axhausen, "Induzierter Verkehr durch autonome Fahrzeuge: Eine Abschätzung," FEDRO, ETH Zurich, SVI 2016/001, 2019.

[20] M. Taiebat, A. L. Brown, H. R. Safford, S. Qu, and M. Xu, "A Review on Energy, Environmental, and Sustainability Implications of Connected and Automated Vehicles," Environ. Sci. Technol., vol. 52, no. 20, pp. 11449–11465, 2018, doi: 10.1021/acs.est.8b00127

[21] C. Herweijer, B. Combes, and J. Jonathan, "How AI can enable a Sustainable Future," Pwc, Microsoft, 2019. Available: https://www.pwc.co.uk/sustainability-climate-change/assets/pdf/how-ai-can-enable-a-sustainable-future.pdf

[22] Capgemini Research Institute, "Climate AI: How artificial intelligence can power your climate action strategy," 2020. Available: https://www.capgemini.com/gb-en/wp-content/uploads/sites/5/2022/05/Climate-AI-Report-1.pdf

[23] Digitalization for Sustainability, "Digital Reset: Redirecting Technologies for the Deep Sustainability Transformation," Oekom, 2023. Available: https://digitalization-for-sustainability.com/digital-reset/

[24] S. Lange, J. Pohl, and T. Santarius, "Digitalization and energy consumption. Does ICT reduce energy demand?," Ecological Economics, vol. 176, p. 106760, 2020, doi: 10.1016/j.ecolecon.2020.106760

[25] J. Bieser, R. Hintemann, L. M. Hilty, and S. Beucker, "A review of assessments of the greenhouse gas footprint and abatement potential of information and communication technology," Environmental Impact Assessment Review, vol. 99, p. 107033, 2023, doi: 10.1016/j.eiar.2022.107033

[26] R. Vinuesa et al., "The role of artificial intelligence in achieving the Sustainable Development Goals," Nat Commun, vol. 11, no. 1, p. 233, 2020, doi: 10.1038/s41467-019-14108-y

[27] F. Berkhout and J. Hertin, "Impacts of Information and Communication Technologies on Environmental Sustainability: speculations and evidence," Brighton, United Kingdom, 2001.

[28] L. Hilty and B. Aebischer, "ICT for Sustainability: An Emerging Research Field," in ICT Innovations for Sustainability, in Advances in Intelligent Systems and Computing, no. 310. , Springer International Publishing, 2015, pp. 3–36, doi: 10.1007/978-3-319-09228-7_1

[29] N. C. Horner, A. Shehabi, and I. L. Azevedo, "Known unknowns: indirect energy effects of information and communication technology," Environ. Res. Lett., vol. 11, no. 10, p. 103001, 2016, doi: 10.1088/1748-9326/11/10/103001

[30] J. Pohl, L. M. Hilty, and M. Finkbeiner, "How LCA contributes to the environmental assessment of higher order effects of ICT application: A review of different approaches," Journal of Cleaner Production, vol. 219, pp. 698–712, 2019, doi: 10.1016/j.jclepro.2019.02.018

[31] L. H. Kaack, P. L. Donti, E. Strubell, G. Kamiya, F. Creutzig, and D. Rolnick, "Aligning artificial intelligence with climate change mitigation," Nat. Clim. Chang., vol. 12, no. 6, Art. no. 6, 2022, doi: 10.1038/s41558-022-01377-7

[32] S. Hall, C. Moskovitz, and M. Pemberton, "Understanding Text Recycling: A Guide for Researchers," Text Recycling Research Project, 2021. Available: textrecycling.org

[33] A. Martín-Martín, M. Thelwall, E. Orduna-Malea, and E. Delgado López-Cózar, "Google Scholar, Microsoft Academic, Scopus, Dimensions, Web of Science, and OpenCitations' COCI: a multidisciplinary comparison of coverage via citations," Scientometrics, vol. 126, no. 1, pp. 871–906, 2021, doi: 10.1007/s11192-020-03690-4

[34] R. Verdecchia, J. Sallou, and L. Cruz, "A systematic review of Green AI," WIREs Data Mining and Knowledge Discovery, vol. 13, no. 4, p. e1507, 2023, doi: 10.1002/widm.1507

[35] D. Amodei and D. Hernandez, "AI and compute," Open AI, 2018. Accessed: Nov. 27, 2023. [Online]. Available: https://openai.com/research/ai-and-compute

[36] D. Hernandez and T. B. Brown, "Measuring the Algorithmic Efficiency of Neural Networks." arXiv, 2020, doi: 10.48550/arXiv.2005.04305.

[37] L. Hilty and J. Bieser, "Opportunities and Risks of Digitalization for Climate Protection in Switzerland," University of Zurich, Zurich, Switzerland 2017, doi: 10.5167/uzh-141128

[38] S. Zhang et al., "OPT: Open Pre-trained Transformer Language Models." arXiv, 2022, doi: 10.48550/arXiv.2205.01068

[39] C.-J. Wu et al., "Sustainable AI: Environmental Implications, Challenges and Opportunities." arXiv, 2022, doi: 10.48550/arXiv.2111.00364

[40] OECD, "Measuring the environmental impacts of artificial intelligence compute and applications: The AI footprint," OECD, 341, 2022.

[41] N. C. Thompson, K. Greenewald, K. Lee, and G. F. Manso, "The Computational Limits of Deep Learning." arXiv, 2022. doi: 10.48550/arXiv.2007.05558

[42] D. Stamoulis, E. Cai, D.-C. Juan, and D. Marculescu, "HyperPower: Power- and memory-constrained hyper-parameter optimization for neural networks," in 2018 Design, Automation & Test in Europe Conference & Exhibition (DATE), 2018, pp. 19–24, doi: 10.23919/DATE.2018.8341973

[43] B. Bookhagen et al., "Metallic resources in smartphones," Resources Policy, vol. 68, p. 101750, 2020, doi: 10.1016/j.resourpol.2020.101750

[44] C. Achillas, D. Aidonis, C. Vlachokostas, A. Karagiannidis, N. Moussiopoulos, and V. Loulos, "Depth of manual dismantling analysis: A cost–benefit approach," Waste Management, vol. 33, no. 4, pp. 948–956, 2013, doi: 10.1016/j.wasman.2012.12.024

[45] T. Milo, "Getting Rid of Data," J. Data and Information Quality, vol. 12, no. 1, p. 1:1-1:7, 2019, doi: 10.1145/3326920

[46] S. Naumann, E. Kern, A. Guldner, and J. Gröger, "Umweltzeichen Blauer Engel für ressourcen- und energieeffiziente Softwareprodukte," Umweltbundesamt, 119/2021, 2021. Available:

https://www.umweltbundesamt.de/publikationen/umweltzeichen-blauer-engel-fuer-ressourcen

[47] P. Li, J. Yang, M. A. Islam, and S. Ren, "Making AI Less 'Thirsty': Uncovering and Addressing the Secret Water Footprint of AI Models." arXiv, 2023, doi: 10.48550/arXiv.2304.03271

[48] S. Duranton, M. Frédeau, and R. Hutchinson, "Reduce Carbon and Costs with the Power of AI," BCG, 2021. Accessed: Nov. 27, 2023. [Online]. Available: https://www.bcg.com/publications/2021/ai-to-reduce-carbon-emissions

[49] J. Bieser and L. Hilty, "Indirect Effects of the Digital Transformation on Environmental Sustainability: Methodological Challenges in Assessing the Greenhouse Gas Abatement Potential of ICT," in Proceedings of the 5th International Conference on Information and Communication Technology for Sustainability, in EPiC Series in Computing, vol. 52. Toronto, Canada: EasyChair, pp. 68–81, 2018, doi: https://doi.org/10.29007/lx7q

[50] V. Coroamă, P. Bergmark, M. Höjer, and J. Malmodin, "A Methodology for Assessing the Environmental Effects Induced by ICT Services. Part I: Single services," in Proceedings of the 7th International Conference on ICT for Sustainability, Bristol, United Kingdom: ACM, pp. 36-45, 2020, doi: 10.1145/3401335.3401716

[51] P. Bergmark, V. Coroamă, M. Höjer, and C. Donovan, "A Methodology for Assessing the Environmental Effects Induced by ICT Services. Part II: Multiple Services and Companies," in Proceedings of the 7th International Conference on ICT for Sustainability, Bristol, United Kingdom: ACM, pp. 46-55, 2020, doi: 10.1145/3401335.3401711

[52] L. Hilty, B. Aebischer, and A. E. Rizzoli, "Modeling and evaluating the sustainability of smart solutions," Environmental Modelling & Software, vol. 56, pp. 1–5, 2014, doi: 10.1016/j.envsoft.2014.04.001

[53] J. Malmodin and V. Coroamă, "Assessing ICT's enabling effect through case study extrapolation. The example of smart metering," in Proceedings of Electronics Goes Green 2016+ (EGG), Berlin, Germany: IEEE, pp. 1–9, 2016, doi: 10.1109/EGG.2016.7829814

[54] L. Erdmann and L. Hilty, "Scenario Analysis: Exploring the Macroeconomic Impacts of Information and Communication Technologies on Greenhouse Gas Emissions," Journal of Industrial Ecology, vol. 14, no. 5, pp. 824–841, 2010, doi: 10.1111/j.1530-9290.2010.00277.x

[55] J. Bieser and E. Kriukelyte, "The digitalization of passenger transport : Technologies, applications and potential implications for greenhouse gas emissions," KTH Royal Institute of Technology Stockholm, 2021. Available: https://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-305425

[56] BCG, "87% of Climate and AI Leaders Believe That AI Is Critical in the Fight Against Climate Change." Accessed: Nov. 27, 2023. [Online]. Available: https://www.bcg.com/press/7july2022-ai-is-critical-in-fight-against-climate-change

[57] GeSI and Accenture Strategy, "#SMARTer2030. ICT Solutions for 21st Century Challenges," Brussels, Belgium, 2015. Available: https://smarter2030.gesi.org/

[58] ITU, "ITU-T L.1480 (12/2022): Enabling the Net Zero transition: Assessing how the use of information and communication technology solutions impact greenhouse gas emissions of other sectors," 2022. Available: https://www.itu.int/rec/T-REC-L.1480-202212-I

[59] C. Bremer, G. Kamiya, P. Bergmark, V. C. Coroama, E. Masanet, and R. Lifset, "Assessing Energy and Climate Effects of Digitalization: Methodological Challenges and Key Recommendations." nDEE Framing Paper Series , Rochester, NY, 2023. doi: 10.2139/ssrn.4459526

[60] R. Abduljabbar, H. Dia, S. Liyanage, and S. A. Bagloee, "Applications of Artificial Intelligence in Transport: An Overview," Sustainability, vol. 11, no. 1, Art. no. 1, 2019, doi: 10.3390/su11010189

[61] F. Creutzig et al., "Leveraging digitalization for sustainability in urban transport," Global Sustainability, vol. 2, p. e14, 2019, doi: 10.1017/sus.2019.11

[62] J. Bieser and L. Hilty, "Assessing Indirect Environmental Effects of Information and Communication Technology (ICT): A Systematic Literature Review," Sustainability, vol. 10, no. 8, p. 2662, 2018, doi: 10.3390/su10082662

[63] J. A. "Skip" Laitner, C. P. Knight, V. L. McKinney, and K. Ehrhardt-Martinez, "Semiconductor technologies: The potential to revolutionize U.S. energy productivity (Part III)," Environ. Qual. Manage., vol. 19, no. 4, pp. 29–50, 2010, doi: 10.1002/tqem.20261

[64] J. Clausen, T. Niebel, R. Hintemann, S. Schramm, J. Axenbeck, and S. Ifflländer, "Klimaschutz durch digitale Transformation: Realistische Perspektive oder Mythos?," Borderstep Institute, Berlin, CliDiTrans Endbericht, 2022. Available: https://www.borderstep.de/wp-content/uploads/2022/01/AP4_CliDiTrans_Endbericht_20220202.pdf

[65] P. Schulte, H. Welsch, and S. Rexhäuser, "ICT and the Demand for Energy: Evidence from OECD Countries," Environ Resource Econ, vol. 63, no. 1, pp. 119–146, 2016, doi: 10.1007/s10640-014-9844-2

[66] L. Hilty, P. Arnfalk, L. Erdmann, J. Goodman, M. Lehmann, and P. Wäger, "The relevance of information and communication technologies for environmental sustainability – A prospective simulation study," Environmental Modelling & Software, vol. 21, no. 11, pp. 1618–1629, 2006, doi: 10.1016/j.envsoft.2006.05.007

[67] M. Flayelle, D. Brevers, D. L. King, P. Maurage, J. C. Perales, and J. Billieux, "A taxonomy of technology design features that promote potentially addictive online behaviours," Nat Rev Psychol, vol. 2, no. 3, Art. no. 3, 2023, doi: 10.1038/s44159-023-00153-4

[68] A. Bleier and M. Eisenbeiss, "Personalized Online Advertising Effectiveness: The Interplay of What, When, and Where," Marketing Science, vol. 34, no. 5, pp. 669–688, 2015, doi: 10.1287/mksc.2015.0930