



Application Aware, Life-Cycle Oriented Model-Hardware Co-Design Framework for Sustainable, Energy Efficient ML Systems

Exploring new ML models

Deliverable D4.1

WP4 - Interaction and User Studies



This project has received funding from the European Union's Horizon Europe research and innovation programme (HORIZON-CL4-2021-HUMAN-01) under grant agreement No 101070408

Project

Title: SustainML: Application Aware, Life-Cycle Oriented Model-Hardware Co-Design Framework for Sustainable, Energy Efficient ML Systems

Acronym: SustainML

Coordinator: eProsima

Grant agreement ID: 101070408

Call: HORIZON-CL4-2021-HUMAN-01

Program: Horizon Europe

Start: 01 October 2022

Duration: 36 months

Website: <https://sustainml.eu>

E-mail: sustainml@eprosima.com

Consortium: **eProsima (EPROS)**, Spain
DFKI, Germany
Rheinland-Pfälzische Technische Universität Kaiserslautern-Landau (RPTU), Germany
University of Copenhagen (KU), Denmark
National Institute for Research in Digital Science and Technology (INRIA), France
IBM Research GmbH, Switzerland
UPMEM, France

Deliverable

Number: **D4.1**

Title: **Exploring new ML models**

Month: 12

Work Package: WP4 - Interaction and User Studies

Work Package leader: Inria

Deliverable leader: Inria

Deliverable type: R

Dissemination level: Confidential (CO)

Date of submission: 2023-09-31

Version: v2.0

Status: Finished

Version history

Version	Date	Responsible	Author/Reviewer	Comments
v1.0	18-09-2023	Inria	Inria	
v2.0	28-09-2023	Inria	Inria	

Executive summary

SustainML project aims to develop a design framework and an associated toolkit, so-called SustainML, that will foster energy efficiency throughout the whole life-cycle of Machine Learning (ML) applications: from the design and exploration phase that includes exploratory iterations of training, testing and optimizing different system versions through the final training of the production systems (which often involves huge amounts of data, computation and epochs) and (where appropriate) continuous online re-training during deployment for the inference process. The framework will optimize the ML solutions based on the application tasks, across levels from hardware to model architecture. It will also collect both previously scattered efficiency-oriented research, as well as novel Green-AI methods. Artificial Intelligence (AI) developers from all experience levels can make use of the framework through its emphasis on human-centric interactive transparent design and functional knowledge cores, instead of the common blackbox and fully automated optimization approaches.

This report corresponds to *Deliverable D4.1 - Exploring new ML models* of the SustainML project and is structured into four parts: First, we report qualitative studies to better understand the current awareness of ML and HCI experts on their impact on sustainability (section 2). We then present a framework (section 3) that structures the different intersections between sustainability with ML and HCI and describe the resulting research areas based on recent work. In the third part (section 4), we focus on HCI for sustainable ML and how human-centered design practices can be used to support more sustainable ML development throughout the entire life-cycle of systems. We define various impacts of ML systems and propose strategies to mitigate these impacts, including data collection, model training, and recycling. To anticipate later stages and impacts in the life-cycle, however, requires ML experts to express realistic project requirements before starting a ML project, which can be challenging. Following a more human-centered design approach, we examine how current decisions in the early stages are made based on interviews with eight ML experts (section 5).

In order to promote energy efficiency and avoid AI-waste throughout the entire life-cycle of ML applications, however, developers need the ability to make environmentally conscious decisions from the start. Our analysis of current sustainability awareness and existing work at the intersection of HCI and ML shows the potential for using HCI methods to redefine the life-cycle of ML models by support sustainable actions through human-computer collaboration. Our final objective is to create tools that help ML practitioners understand how a certain decision affects the environment and guide them in identifying suitable, more environmentally friendly models for their projects. This deliverable provides the background for developing such tools, which will be presented in later deliverables.

Contents

Executive Summary	3
Contents	5
Acronyms	6
1 Introduction	7
2 Understanding the Intersection of Sustainability with HCI and ML	7
2.1 Methodology	8
2.1.1 Interviews	8
2.1.2 Questionnaire	9
2.2 Awareness of Sustainability in ML and HCI	9
2.2.1 Understanding the Impact of Machine Learning on Sustainability	9
2.2.2 Understanding the Impact of Human-Computer Interaction on Sustainability	11
2.3 Conclusion	12
3 Framing the Intersections of Sustainability with HCI and ML	13
3.1 Methodology	13
3.2 The intersection of HCI and Sustainability	14
3.2.1 HCI → Sustainability	14
3.2.2 Sustainability → HCI	14
3.3 The intersection of ML and Sustainability	15
3.3.1 ML → Sustainability	15
3.3.2 Sustainability → ML	15
3.4 Exploring Research Areas across Research Fields	15
3.4.1 Sustainability → HCI → ML	15
3.4.2 Sustainability → ML → HCI	15
3.4.3 HCI → Sustainability → ML	16
3.4.4 HCI → ML → Sustainability	16
3.4.5 ML → HCI → Sustainability	16
3.4.6 ML → Sustainability → HCI	16
3.5 Conclusion	16
4 Human-Centered Design for Sustainable ML Life-Cycles	17
4.1 The perspective of HCI for Sustainability	17
4.2 The perspective of Sustainable of ML	17
4.3 Supporting decision making for Sustainable Life-cycles in ML	19
4.4 Role of HCI for Sustainable Machine Learning	20
4.5 Conclusion	21
5 Understanding Existing ML Workflows	21
5.1 Interviews	22
5.2 Current ML practice	22
5.2.1 Using data formats and benchmark examples alongside problem definitions to improve understanding	22
5.2.2 Insufficiency of training and evaluation data challenges model selection	23
5.2.3 Combining Literature, LLM's and AutoML to select a suitable model	23
5.2.4 Balancing Performance, user-understandability, and Trust in AI Tools to improve Model Selection.	24



5.2.5	Comparing ML models against existing and self-defined metrics to evaluate suitability	24
5.2.6	Prioritizing Interpretability and Explainability for Model Viability	25
5.3	Conclusion	25
6	Summary	25
	References	27

Acronyms

AI	Artificial Intelligence.
HCI	human-computer interaction.
ML	Machine Learning.
SHCI	Sustainable HCI.

1 Introduction

Sustainability has emerged as an overall concern for society [1]. While research on the intersection of human-computer interaction (human-computer interaction (HCI)) and sustainability as well as machine learning (ML) and sustainability began already more than two decades ago [2, 3], the rapid increase in the complexity of the computational systems we use on a daily basis created a new awareness for this pressing issue. According to researchers, ML can be a double-edged sword that has the potential to address sustainability issues while also contributing to the problem in a growing way [4]. For example DeepLearning, Generative AI, and other similar technologies have raised the bar in terms of data storage, computation time, and environmental impact while an increasing amount of energy [5], water [6] and ecological crisis management [7] is supported by ML applications.

Several research fields are involved in this effort. However, research in ML and HCI, in particular, is at an important moment, with the potential to worsen or ease these serious concerns. While researchers in ML [8] and HCI [9] paid increasing attention to sustainability in the last decade, how sustainability is defined, addressed and evaluated still varies within and across fields [10]. This is due in part to the multifaceted sustainability impacts, which range from environmental to economic to societal effects [11], but it is also due to a lack of structure and comparison across fields.

Over the years, the HCI community considered various approaches to addressing sustainability, beginning with working on individual behavior [12], then applying persuasive technologies [13], and finally arriving at a call for 'Green Policy informatics' that enables sustainable HCI to leverage a more traditional HCI skillset in addressing sustainability issues [9]. At the same time, the majority of work on sustainable ML and artificial intelligence (AI) addresses how to improve data collection, power sources, and infrastructures, as well as how to quantify and lower the carbon footprint [4] associated with developing and fine-tuning an algorithm [14]. Structuring the phrasing around sustainability would therefore help the different communities better explore the potential research space as well as better compare their results within and across communities in order to create a larger successful effort [13]. We approach this topic from within the communities. We began by conducting two qualitative studies with ML and HCI experts to assess general understanding and concerns about sustainability (section 2). We then present a framework (section 3) that structures the different intersections between sustainability with ML and HCI and describe the resulting research areas based on recent work.

Our ultimate objective is to create tools that helps ML practitioners understand how a certain decision affects the environment and guide them in identifying suitable, more environmentally friendly models for their projects. We then continued to focus on how human-centered design practise can support more sustainable ML development across the entire system life-cycle. We discuss the various negative effects of ML systems on sustainability and offer human-centered design approaches to partially rethink the life cycle of ML models (section 4).

To anticipate later stages, use cases and impacts requires ML experts to express realistic project requirements before starting a ML project, which can be challenging. Following a more human-centered design approach, we examine how current decisions in the early stages are made based on interviews with eight ML experts (section 5).

2 Understanding the Intersection of Sustainability with HCI and ML

Before investigating the intersections of HCI and ML with Sustainability, we must first define what we mean by these communities. Understanding the different points of view is critical, and it necessitates taking into account the various community backgrounds. Hence we will briefly introduce the three fields.

Sustainability generally refers to “the development that meets the needs of the present generation without compromising the ability of future generations to meet their own needs” [15]. Other researchers attempted to define sustainability in the context of their application area; for example, in engineering and business studies, sustainability is commonly addressed by its environmental impact, such as total energy consumption [16]. Researchers in design and human-computer interaction highlighted early on that sustainability necessitates a multifaceted strategy [11], ranging from environmental to economical and societal effects [9]. This includes sustainability aspects related to product life-cycle, energy efficiency, and socially responsible design practices. Understanding and achieving sustainability requires a shift to a meaning that is specific and can be quantified, in order to monitor the progress towards achieving the sustainability goals [17], which can be particular challenging across research fields.

Human-Computer Interaction is the study of how people design, implement and use interactive computer systems and how computers affect individuals, organizations and societies [17]. With that it studies how people interact with and through computers and aims to reduce the complexity and/or increase the power of interactions with them, to empower users to focus on their tasks, needs and habits [18, 19].

Machine Learning is a series of algorithms that can automatically improve over time, along with the scientific study of such algorithms. Finding the “fundamental laws that govern all learning processes” is one of its main objectives [20]. As such, it aims to automate the task of analytical model building in order to perform cognitive tasks such as object detection or natural language translation by detecting patterns in data to predict future data or outcomes [21, 22].

2.1 Methodology

Our goal is to develop a better understanding how sustainability is currently perceived in ML and HCI research and the factors researchers are aware of that impact it. Therefor we evaluated 8 interview with ML experts and 10 questionnaire responses by HCI researchers regarding their perception of sustainability. We first introduce our methodology before outlining our results in detail.

2.1.1 Interviews

We were particularly interested in investigating ML experts’ understanding of sustainability and how it affects and is addressed in their daily work.

Participants: We recruited eight ML experts (8 male, avg. 10.3 years of experience [SD=3.1]) through emails to local research facilities or direct contact (details in table 2). While we aimed for an equal gender balance among the participants and tried to recruit female ML experts by personally addressing them and inviting them to our study, unfortunately three female candidates could not find the time or canceled last minute. All participants are researchers who specialise in a particular ML application, such as biology, medicine, or CPUs. Participants were not paid and were informed about their rights and data usage via an informed consent form in accordance with GDPR.

Procedure: Each interview lasted one hour. All interviews were conducted in English by one of the researchers, either in person or online. Participants first confirmed the consent form, before the interviewer briefly introduced the context. In a semi-structured interview, interviewees presented their most recent work before discussing how their work affects sustainability, important sustainability factors, and how they address sustainability in their work.

Data collection and analysis: We ran a mixed-approach thematic analysis [23]. The top-down themes included ‘Sustainability awareness’, ‘Sustainability Factors’ and ‘Impact on Sustainability’ based on literature and interview questions. One author coded the participants’ answers based on both top-down themes and themes that emerged (bottom-up). The same author summarized closely related themes and identified patterns; a second author who was not involved in the interviews re-coded the data using the final set of themes.

ID	Age range	Gender	Years of experience
1	30–40	M	10
2	20–30	M	6.5
3	20–30	M	10
4	40–50	M	15
5	20–30	M	6
6	30–40	M	10
7	30–40	M	11
8	30–40	M	14

ID	Age range	Gender	Years of experience
11	40–50	M	15
12	30–40	M	7
13	40–50	M	20
14	60–70	M	40
15	30–40	M	14
16	30–40	M	9
17	30–40	F	8
18	40–50	M	11
19	30–40	F	11
20	30–40	F	8

Table 2: Demographics of the interviewed ML experts

Table 3: Demographics of the responding HCI experts

2.1.2 Questionnaire

We disseminated a digital survey to assess HCI experts’ understanding of sustainability and how it affects their everyday work.

Participants: We received answers from ten HCI experts (7 men and 3 women, avg.: 14.3 years of experience [SD = 9.9]) primarily working in cross-domains with HCI such as dance, gaming, finance, and AI. All participants are currently working in academic or industry research. The questionnaire was distributed through personal networks to HCI experts with at least four years of experience, and participants were not compensated.

Questionnaire: The questionnaire contained three main questions: 1) regarding their general understanding of the impact of HCI on sustainability and how they keep informed about it; 2) What factors they are aware of that could have an impact on sustainability in their personal work; and 3) If they have/do considered sustainability in their work and which measures they have taken and why. Additionally we collected demographic data.

Data collection and analysis: This time we ran a top-down thematic analysis on the data. Our themes included again ‘Sustainability awareness’, ‘Sustainability Factors’ and ‘Impact on Sustainability’ from literature. One author coded the participants’ answers and identified patterns; a second author reviewed the themes.

2.2 Awareness of Sustainability in ML and HCI

We started analysing current awareness at the intersection of machine learning and human-computer interaction with sustainability through interviews and questionnaires. This provided valuable insights into concerns, current practice and highlighted the different perspectives of these intersecting fields.

2.2.1 Understanding the Impact of Machine Learning on Sustainability

Most of our interviewees (7/8) agreed that machine learning has an impact of sustainability. However, the scale of consideration varied across participants. Some actively work on reducing the impact of models (P4, P5, P6, P7), some collect information about their consumption and report them (P3, P4) and others are aware of them but do not actively consider them in their work (P1, P2, P8).

Reasons for not considering sustainability in their work were due to the perceived small impact of their work (3/8) or priorities of result quality over sustainability impact (1/8). Common arguments among

those answer was that “the models that I am building are nothing in comparison to the big models” (P2) in industry or commercial products and hence academic ML model dimensions are too small to have a real impact (P1, P3). One participant highlighted the importance to balance the quality of the outcome with the sustainability impacts, arguing that “this model is working, maybe I’ll try to do something more frugal” (P2) relies on a satisfying model in the first place.

Evaluating the energy and infrastructure impact of a model is limited Our interviewees distinguished between the impact of the models and the infrastructure needed to deploy them.

Evaluating the model impact: The main impact of ML models is the amount of energy used and the associated carbon footprint (5/8). This is mainly connected with the training time (P6), making deep learning approaches especially costly (P3). Three of our interviewees (P3, P4, P7) reported to track their impact using e.g. existing toolboxes and carbon tracker. However, P4 points out that the results of existing carbon trackers can differ significantly, reducing the comparability of approaches. While measuring carbon emissions is currently the only real indicator of potential impact on the environment in terms of the climate and open question remains “what aspects of machine learning model computation to consider. Like do you consider network requests between different nodes in a data cluster that’s running your model?” (P7) Answering this question can be challenging as data are not always available. Especially when working on or with non-public commercial models e.g. from OpenAI, “both the electricity cost and the general cost of using those models” (P5) stay hidden and make it hard to relate to a single application. Only one interviewee(P4) mentioned to consider the overall impact of his work, including “all the emails, the platform we use, the video conferences[...], if I had to travel to meet with colleagues” to create a complete project budget.

Evaluating the infrastructure impact: Further considerations are related to the infrastructure mainly *hardware* (3/8) but also water consumption in the process (2/6). Interviewees highlight that using PC computers (ca. 800 Watt) often exceeds mobile chips by a 10-fold (e.g. Apple M2 Pro chip 50 Watt), “even in this case, even if the MacBook runs a little slower, but still it saves a lot of energy” (P6). There was a general consensus that working on larger clusters where “maybe 200 people uses one cluster” (P6) should be preferred compared to running a dedicated PC computer(P4). Another factor are the production costs of CPU and GPU chips, which are often used to run large ML models, such as GPT. While there are few data available about the mining and production costs of CPU’s (P4, P7) there is hardly any about GPU production. P4 explained that from a sustainability standpoint it “is completely unbelievable that in 2023 we don’t even know how much it costs in terms of environmental impact to produce a GPU”, even though foundation models increasingly rely on them. A further concern was the amount of water and cooling costs of data center, which are often hidden when computation is outsourced to larger infrastructures (P7).

The Trend to large ML models has a negative effect on sustainability: One recurring topic throughout the interviews was the complexity of ML models and their impact on sustainability. Our experts highlighted that the size and complexity has a large impact on the computation time and infrastructure required, so if an algorithm just gain a few tens of percent of accuracy the carbon footprint increase is very large(P4). Experts further worried about the increasing trend to apply larger complex models, like deep learning, which have an especially large carbon footprint (P1, P3). Such models take “more like one month or hundreds of hours of training on hundreds of GPUs”(P2) to train which has a large cost attached to it.

Some have called to reconsider the impact of this trend in education, academia and industry, reporting that students “forget about all the other simple approaches that can be better, sometimes orders of magnitude better, when you think about inference time” and similar factors (P1). This requires a rethinking about the actually needs of a model, for example “if you have 90% accuracy compared to 95% accuracy, but the 90% accuracy uses only 10% of the footprint, then is the 90% already enough for the actual use case?”(P6). This could also results in refraining at all from machine learning models and

instead turn to “simpler algorithm or simpler decision processes” (P4).

Different perspectives on the intersection of ML and Sustainability: Beyond environmental impacts interviewees also highlighted ‘social factors like algorithmic bias [or] economic factors such as how do we democratize AI and increase access to compute infrastructure to people who may not be able to afford it’. Supporting economic sustainability requires “opening both the code of training, how it is trained, on which data, what is the result you have so that you can have people that test fairness, privacy” (P5). Work on sustainability hence means also to provide models and code along with ML models (P5). Only one of our interviewee reported to work on sustainability concerns by using machine learning models to generating guidelines for pesticide spreading and safe housing locations (P5).

Taking a larger perspective, sustainability can be also understood across the whole machine learning life-cycle, balancing “the training, the footprint cost during the training, and the longer time deployment”, including the used hardware in the later stages or if a model could be fine-tuned or need to be retrained after some time (P6).

In general, we found differences in how our interviewees approached sustainability in their work. Some concentrated on the environmental implications of hardware and models, as well as how to evaluate them. Others were more concerned with how the ML community as a whole can address sustainability, such as changes in education/ML methods that reduce (unnecessary) computing and work that focuses on the entire life-cycle and deployment of ML models in a sustainability context.

2.2.2 Understanding the Impact of Human-Computer Interaction on Sustainability

Most participants(8/10) were aware of the impact of computer systems and HCI on sustainability through articles and recent publications. Some are actively involved in university (P14) and community (P18) structured to reduce the impact of HCI research on the environment. Two participants answered having no information about the particular impact of HCI on sustainability (P13, P17) beyond the general concerns.

Main sustainability impact from computation, hardware and end-user materials Our respondents agreed on three major aspects that influence HCI sustainability: energy consumption with the attributed carbon footprint (7/10); hardware used in (interactive) computing systems (3/10); and the impact of end-user materials on the environment (4/10). The increasing energy intensity of computing systems, especially with ML systems, is becoming a topic of increased concern (P14, P16, P18, P20). One participant highlighted that “User interfaces and the whole software stack are extraordinarily wasteful of computing resources if you consider that they are ~50000 times faster than 30 years ago and in many cases don’t do that much more” (P14). The increases energy consumption by current system “is in most countries still not sustainable”, also reducing the accessibility of advancements (P16). Manufacturing of computational equipment was seen as another impact source, especially in regard of electronic components (P14, P18, P19). The third factor is driven by “more general capitalistic concern for the proliferation of relatively cheap (but not always!) and ultimately disposable electronics” (P18) resulting in the expectation of replacing VR headsets, mobile phones or headsets every few years (P11, P15, P18).

Hidden negative factors encourage by HCI work and how to address them: Respondents reported concerns about the impact of novelty on the sustainability of new technology, as P11 stated: “HCI promoting and pushing for ”innovation” with new devices, gadgets, features that feed into a consumerist system and foster more discarding and more consumption.” Beyond innovations, work in HCI can also create a “so-called rebound effect where more efficient technology ends up increasing emissions because it triggers an increase in the use of that technology.” (P14). Especially with the objective to make technology more accessible to a larger audiences creates a certain responsibility by the HCI community for the increase of technology use overall (P15).

One respondent suggested to rethink “the user needs and the scenarios that we decide as researchers to focus on; the technologies what we employ; our design solutions (to what extent do we consider sustainability as a design criterion, constraint or objective?); our research methodologies; and our everyday practices as researchers.” (P13). Some respondents stated that they actively address some of these factors in their work by focusing more on “on making existing devices last longer and become easier to maintain/take care of” (P11) or to ensure “that the system we’re building accommodates older devices” (P19). Others have started adapting their design process to reduce the amount of waste in fabrication (P18) or started integrating more sustainable options as default in their system (P20).

Different perspectives on the intersection of HCI and Sustainability: The responses provided above demonstrate that sustainability was viewed not only from a computational standpoint, but also as something to promote throughout the design cycle and throughout the research life itself. This included reducing personal negative impacts through less travel (P13, P14, P15), reduced use of consumables overall (P13, P15, P18).

Other participants suggested work that encourage behaviour change towards more sustainable lifestyles and technology use (P16) or informing users of the long-term consequences of their actions through new visualization (P15).

We discovered differences in how respondents framed and addressed sustainability issues in HCI. More emphasis was placed on the researcher’s responsibility to decide how sustainability should be considered and addressed in their work: as an objective, as part of their method, or in everyday practise. It was also highlighted that more sustainable design is required to reduce the negative impacts of HCI research in general.

2.3 Conclusion

Our interview and questionnaire revealed a common understanding in the research community on the impact of technology on the environment. We identified that HCI researchers were more aware of the sustainable impacts of the whole technological life-cycle from material-waste during the design to the increased consumption of end-user hardware. This is in line with current research trends on sustainable HCI and sustainable interaction design presented in section 3. We also identified a stronger sense of perceived individual responsibility by researchers, e.g. in form of consumption or travel, but also on the hidden impacts of HCI research has on the end user behavior. This awareness is an opportunity to reflect and rethink our objectives, methods and artifacts we create. Work presented in 6.6. ML → Sustainability → HCI, show examples how this could be addressed in the future.

The history of sustainable ML is traditionally more focused on the impact of models and hardware on the environment than the individual or the life-cycle. Nevertheless, we saw in our interviews that this has shifted and ML experts are aware and address a diversity of sustainable impacts in regards to model accessibility, or improving the ML life-cycle, which is supported by literature that we highlight in section 3. We would like to emphasize the discussion about the necessity of model complexity, and the responsibility ML experts expressed to advocate for reasonable use of model, including pruning models, simpler models or even no ML models to solve a task. However, this requires as much an academic shift in how students are taught model diversity as a better public discussion on necessary capabilities instead of aiming for always 100% accuracy for every situation. The research area HCI → Sustainability → ML opens up a potential research space to explore how HCI methodology could help to integrate and improve such discussions to design of more sustainable ML models.

3 Framing the Intersections of Sustainability with HCI and ML

We developed a framework to better describe and outline research areas at the intersection of sustainability, ML, and HCI based on the qualitative data described above and our observations while investigating existing literature. Sustainability research in HCI and ML is quite diverse, and we notice that many participants say “we do sustainability and HCI” or “we do sustainability and ML”, but we realised that there are many different ways to do each of these things. In accordance with our previously described findings, we can identify two factors that primarily influence the research conducted at these intersections: the research stance and the relation direction to the addressed field.

We propose that the research stance or perspective, which frequently stems from the context of the primary research field, influences the type of contributions, methods, and evaluations applied. We further propose that the relation between the fields is directional and implies the focus of the research area. To illustrate these aspects let us look at the intersection of *Sustainability* and *Machine Learning (ML)*.

- *Sustainability* → *ML* (reads ‘Sustainability for Machine Learning’)
It uses methods and evaluations from a sustainability perspective e.g. reducing energy, hardware materials uses in/by and for ML applications. For example reducing the energy consumption of a model by reducing its training time significantly. The evaluation methods used are derived from the original research stance in this case the sustainability domain, such as reduced carbon emissions.
- *ML* → *Sustainability* (reads ‘Machine Learning for Sustainability’)
It uses methods and evaluations from a machine learning perspective, e.g. building ML models, to reveal, predict or reduce sustainability impacts. Examples are ML applications to optimize water management, or predict impacts of global warming. The evaluation methods stems from the ML domain e.g. prediction accuracy.

We used this method to provide more structured insights into existing work at the intersections of sustainability with HCI and ML (see following section 3) and to provide insights into the cross-sections of all three fields and the research areas it opens (section 3.4).

3.1 Methodology

We combined our qualitative results with a structured literature analysis and present our insights in form of the proposed framework. Each resulting research area is illustrated by representative work in this field. We conclude with a short discussion about opportunities for research based on the proposed framework.




Figure 1: Papers at the intersection of Sustainability, HCI and ML in the ACM library (2018-2023)

Sampling: In order to get a more recent picture of the discussions with the HCI and ML community in regards to sustainability we focus on publications in the ACM (Association for Computing Machinery) digital library that address at least one intersection between these three fields. For locating relevant papers, the following steps were taken:

1. We limited the publishing years to the past five years (2018 (June) - 2023 (June)) to focus on recent work.
2. We limited the contribution type to *Research Article* to find more advanced work.
3. We searched focused on relevant terms in the abstract or title:
 - For ML - Sustainability: the keywords “ML” OR “machine learning” AND “sustainability” resulting in 428 full publications
 - For HCI - Sustainability: the keywords “HCI” AND “sustainability” resulting in 92 full publications;
 - For HCI - ML - Sustainability: the keywords “ML” OR “machine learning” AND “HCI” AND “sustainability” resulting in 9 full publications

A year distribution of the resulting paper is shown in Fig. 1.

Analysis: We begin our investigation by reading the abstracts of the resulting papers and extracting the main statements relevant to the intersection and shortlisting the most relevant papers. Data was extracted from those revealing essential ideas, methodology, and discoveries that helped in the convergence of the core research areas. We chose a representative sample of these for the proposed research area in sections 3.?? to highlight the difference within these research areas and make no claim of completeness.

3.2 The intersection of HCI and Sustainability

Based on our framework we will outline work in the different resulting research areas and provide representative example of existing work.

3.2.1 HCI → Sustainability

This research area uses methods and evaluations from HCI and design research to address sustainability issues, which is often reference as Sustainable HCI (Sustainable HCI (SHCI)). SHCI started on focusing on increasing sustainability awareness in individual behaviors [12]. Disalvo and colleagues, for instance, advocated for persuasive technologies that use eco-feedback to provide users with information about their energy consumption [13], as suggested by some of our respondents in our earlier study. Some argue that when developing interactive systems, HCI methodologies such as prototyping could be used to increase sustainability feedback [24]. However, the focus of technology as a solution and the focus on the individual is sometimes criticised [9]. For more work related to SHCI and this discussion we suggest Hansson et al.’s review [2] as well as Bremer et al.’s critical review [9].

3.2.2 Sustainability → HCI

This research areas stems from a sustainability stance including evaluation methods and takes an interest in HCI practices often attributed to SHCI or sustainable interaction design (SID). Preist et al.’s work, for example, show how sustainable interaction design can be quantitatively evaluated on the example of Youtube [25] revealing to users hidden ecological costs of such platforms. Another form of work falling into this category is using compostable material to build interactive systems as suggested by Song et al. [26]. In the same category falls the reduction of fabrication waste suggested by one of our respondents (P18). Finally, we would like to highlight work addressing the sustainability of the entire life cycle of interactive system. Vasquez et al. showed in their work how a sustainable prototyping life cycle allows designers to consider the environmental impact of digital fabrication throughout the artifact’s life resulting in more sustainable decision making along the way [27].

3.3 The intersection of ML and Sustainability

3.3.1 ML → Sustainability

This research area includes work that uses ML methods and evaluations e.g. building ML models to reveal, predict or reduce sustainability impacts. ML models for better understanding ecosystems and advancing conservation strategies are examples of such work [28]. An example is P5's work on generating guidelines for pesticide spreading and safe housing locations, mentioned during our studies. Another aspect of this research area is e.g. improving water [6] or energy management [5]. This can be especially helpful in critical situations such as anticipating forest fires to improve reaction time and preservation [7]. For more work in this area we suggest Rolnick et al.'s review on addressing climate change with machine learning [3].

3.3.2 Sustainability → ML

Finally, this research area is defined by methods and evaluations from a sustainability perspective e.g. reducing energy, hardware materials uses in/by and for ML applications. The evaluation methods applied stem from the sustainability domain e.g. reduced carbon emissions. Work in this research are often emphasizing new models that uses less energy during training and with comparable outcomes [29]. Some research emphasises the carbon footprint associated with training large ML models, as well as the importance of prudence in such resource-intensive activities [4], as suggested by our experts. A broader interpretation suggest that sustainable ML is a combination of its environmental footprint combined with ethical concerns, model reliability, and considerations of longevity [4]. This is consistent with work on improving the ML life cycle, as mentioned by P6, which proposes a comprehensive approach from data acquisition to model retirement, with the goal of ensuring holistic sustainability across environmental, social, and economic dimensions. [30].

3.4 Exploring Research Areas across Research Fields

This section discusses the various intersections between all three domains, on the same idea of 'stance domain → target domain' but applied in sequence.

3.4.1 Sustainability → HCI → ML

The closest term related to this research area is sustainable interactive Machine Learning (iML) which frequently refers to models that require fewer iterations or less user feedback, making the interaction process less taxing and more efficient for users [31]. Another perspective on this research area is the emphasis on the sustainability of the computational process itself when real-time feedback is involved [32]. Work that considers the ethical and societal implications of iML systems, ensuring that they remain inclusive, fair, and bias-free when adapting based on user interactions [33] can be seen as another sub-area.

3.4.2 Sustainability → ML → HCI

Integrating sustainability concepts into ML technologies for the HCI domain has become a significant goal in recent research. A number of research addresses how ML algorithms can be beneficial for HCI experiences through e.g. intelligent agents. Recent research has gone deeply into frameworks that help promote more efficient, sustainable ML algorithms, which boosts user experience in HCI [34]. Energy-efficient ML techniques in the context of awareness services and Human Activity Recognition [35] could be further placed within this research area.

3.4.3 HCI → Sustainability → ML

In this research area HCI's user-centric principles are used to address sustainability issues with advancing ML technology. For example, user-friendly ML model interfaces with sustainability-aware feedback loops may aid in the development of more effective and energy-efficient models [36]. Another part of this area of study would be to use HCI techniques to create sophisticated AI/ML-based systems that are recognised as responsible, accessible and safe [37]. While some work exists in this field, it is still rare.

3.4.4 HCI → ML → Sustainability

The intersection of interactive machine learning and sustainability has gained interest in recent years [38]. Another line of research at this intersection is improving interpretability in ML, which was shown crucial for users to understand and trust automated conclusions in regards of sustainability [39]. HCI's strength in including users in the process, for example, for interactive ML techniques, enables faster and better prediction by relying on human experience e.g. to anticipate crop illnesses [40, 41]. Relevant work in this research area is outlined in [42].

3.4.5 ML → HCI → Sustainability

One interpretation of this area is intelligent user interfaces for sustainability. A variety of studies use a human-centered design approach that prioritizes needs and limitations, resulting in a multi-disciplinary analysis to build sustainable adaptable technology solutions that contribute to a viable world with fewer materials and energy consumption [43]. Another example in this study area is EnergyLife, a system that uses wireless sensors, mobile and ambient interfaces to turn energy consumers into active players in a game, promoting more sustainable activities by linking players inside and across homes [44]. More research is needed to address current sustainability challenges by using the full potential of ML and design.

3.4.6 ML → Sustainability → HCI

ML for Sustainable HCI is still a less represented research area, that aims to use model-driven insights provided by ML to develop more sustainable HCI practices and solutions. Work examples are optimization algorithm for CAD system that help designer reduce waste material in manufacturing (see review [45]), or using advanced computational models to improve sustainable building design (see review [46]). Despite these contributions, this research area is still in its early stages, with great potential for future research on ML for sustainable HCI.

3.5 Conclusion

Our proposed framework can help shape the focus of sustainability work in a more multidisciplinary world. Distinguishing academic work by research stance and research relation to the objective aids in determining which objectives, design, and evaluation methods are reasonable, applicable and understudied. This more distinct classification can assist researchers in better identifying relevant related work, identifying research gaps, exploring available methods, and comparing their work to existing publications. It can also provide a framework for investigating design and evaluation methods across domains in order to develop an agreed-upon standard for cross-disciplinary research in sustainable HCI and ML. Using our framing, can also provide outlooks for research areas at the intersection of all three fields. There has been a lot of research done in some of the presented cross-disciplinary research areas, each with their own summaries and discussions, but some areas of research appear to receive less attention. The proposed research structure highlights future research opportunities at these intersections.

4 Human-Centered Design for Sustainable ML Life-Cycles

We locate our work at the intersection of HCI for Sustainable ML ($HCI \rightarrow Sustainability \rightarrow ML$). Our ultimate goal is to develop tools that assist machine learning practitioners in understanding how a particular decision affects the environment and guiding them in identifying more environmentally friendly models throughout the whole ML life-cycle. This requires us to take a closer look at the overall research on ML life-cycles and how design practice can be applied. In the following section we therefore first elaborate on the two intersections $HCI \rightarrow Sustainability$ and $Sustainability \rightarrow ML$ before presenting potential strategies and considerations of sustainability within ML life-cycles. This section presents our perspective on $HCI \rightarrow Sustainability \rightarrow ML$ and layouts potential research directions that could result in more sustainable ML life-cycles.

4.1 The perspective of HCI for Sustainability

Sustainable HCI (SHCI) focused initially on individual behavior and follows a strategy based on how to the user aware about his actions [12]. DiSalvo and colleagues, for example, advocated for a persuasive strategy based on technology that uses eco-feedback to provide users with information about their energy consumption [13]. Many have objected to this approach on the grounds that users are rational consumers [9] and that methodologies commonly used in HCI, such as rapid prototyping, user involvement, and iterative design, can rather allow for early and increasing feedback when developing interactive systems [24]. In a recent review on existing work on sustainable HCI and its critiques, the authors concluded that the HCI community “should challenge the narrative that we can rely on technology to save us, just as we challenged the narrative that climate change is an individual behaviour change problem” [9]. Among the suggested directions emphasized by this review is that technology may give clarity to the complicated dynamics of multiple emission sources, including that of technology itself.

In line with such work, we suggest using a human-centered approach we want to address a more holistic picture of sustainable development and the ML life-cycle. While the sustainable ML community frequently addresses sustainable development, we believe there is a major potential to apply such techniques to improve decisions made by developers when designing systems, such as the requirement for and utilization of data and the algorithms applied.

The collaboration between Human-Computer Interaction (HCI) and Machine Learning (ML) can effectively address sustainability challenges by developing intelligent systems that enhance user behavior and reduce environmental impact. While previous work has often focused on the end-user of a system or the system itself, we would like to emphasize the developer’s role in making more sustainable decisions across the entire ML life-cycle and to support democratization of ML approaches using human centered design methods. We discuss relevant aspects that developers should consider at each stage to make more sustainable decisions. With this paper, we aim to develop a discussion on building and developing a more sustainable future.

4.2 The perspective of Sustainable of ML

Within the sustainable ML and artificial intelligence (AI) community there are two main branches of work: ML for sustainability and sustainability of ML [47]. The majority of work on ML for sustainable focuses on building ML models to predict the impact of technology, such as renewable energy [48] or electric cars [49], on the environment. On the other hand, sustainability of ML addresses how to enhance data collection, power sources, and infrastructures as well as how to quantify and lower the carbon footprint associated with developing and fine-tuning an algorithm [14]. Strubell et al.’s work [4] in particular has pushed the ML community to prioritize the carbon footprint of ML models on their research agenda. Their carbon footprint analysis of training their own models led them to the conclusion that we need to lower the carbon footprint of building and operating ML models, which received large echo throughout the

community [50]. In line with this perspective, Wynsberghe in a recent article defined Sustainable AI/ML as: “a movement to foster change in the entire life-cycle of AI products (i.e. idea generation, training, re-tuning, implementation, governance) towards greater ecological integrity and social justice” [47]. He further emphasizes that sustaining environmental resources for present and future generations, societal economic models, and core societal values should all be compatible with developing ML systems. We see several overlaps in such a perspective with the previously outlined goals of sustainable HCI, and would like to highlight the role the HCI community may play in such an endeavor through our work.




Figure 2: The original CRISP-DM Life-cycle [51] (left), adapted Life-cycle from CRISP-DM (right)

Perspectives on ML Life-cycles: Our work aims to raise awareness of sustainable decisions and their consequences throughout the ML life-cycle. Among the first ML life-cycle model developed was the CRISP-DM in 1999 (Fig. 2.left) [51]. It is based on the premise that training and data management, in particular, significantly add to the overall carbon footprint of ML applications [52]. While it was primarily developed for Data Mining processes, it has gained wider acceptance for representing ML life-cycles in general [51].

It contains six distinct phases: The fundamental problem and user demands, as well as the project’s objectives and requirements, are determined during the **Business understanding** phase. **Data understanding** refers to collecting data and identifying data quality issues. In the **Data preparation** phase the improved dataset is prepared, which often includes data analysis, data cleaning, and data transformation. In the **Monitoring** phase different modelling techniques are explored, selected and calibrated. The model is then examined during the **Evaluation** phase to ensure that it meets the established objectives. Finally, the resulting model is applied in the context of use during the **Deployment** phase.

Later work has adapted the CRISP-DM model [53] and reduced it to four phases (for comparison we marked these phases within the initial CRISP-DM model in Fig. 2.right). It contains a more general **Data extraction** phase, which covers the identification of the main problem, the collection of necessary data and the transformation into a dataset. The second **build and train the model** phase comprises selecting the appropriate algorithm and train it on the prepared data. The third **Run the model** phase involves all actions required to integrate the model into the production system. These may include evaluating parameters, interpreting the results, and identifying any areas for improvement. In the fourth **End-of-Life management** phase, different modelling techniques are tried out to reuse the final results, metrics and model’s performance. We continue to discuss potential methods and approaches to assist developers at each stage of making more sustainable decisions.

4.3 Supporting decision making for Sustainable Life-cycles in ML

Supporting more sustainable decision making for ML life-cycles entails a number of considerations, including ethical data collection, energy-efficient training and deployment, ongoing monitoring, and responsible retirement of models and data. From one step to the other during the ML life-cycle, the developer should be aware of the impact of their action in either the model building or using. This includes comparing model alternatives based on their impact on the environment and collecting only relevant data necessary for the task at hand.

Developing ML algorithms that prioritize sustainability as a key objective requires new approaches. Our focus is to develop tools that explain to ML practitioners how a given algorithm impacts the environment and assist them in identifying appropriate, more sustainable models for their projects. We hope that by doing so, we can encourage the use of more energy-efficient models, as well as more data-efficient training and optimization strategies for algorithms, thereby lowering overall power consumption. To illustrate our vision, we modified the previously mentioned ML life-cycles (Fig. 3) to address more sustainable concerns.

Project understanding: Describing desired ML model characteristics and expressing the project's goals, before choosing a model requires users to express possibly ambiguous or contradictory objectives. Therefore, developers usually depend on their own personal experience or on existing well-documented ML models and toolkits. As a result, supporting developers to find more suitable, energy-efficient alternatives is critical to achieve more sustainable ML practice overall. We emphasize the importance for enabling developers to better compare project goals and constraints, such as maximum training time, amount of training data, or favoring speed over accuracy, in order to assess trade-offs. The HCI community has previously demonstrated in other domains how iterative processes can help users explore and narrow down their needs and ideas to achieve a specific goal. In a similar way, we intend to gather more information by focusing on the user's current understanding, knowledge, and perspective through interviews and observations in order to design products that allow them to iteratively express their needs. We suggest more research to identify the understanding of ML experts before selecting a model to build tools encouraging suitable and more sustainable models.

Data Extraction: Data management, which includes data storage, processing, and sharing of datasets can have a high impact on the environmental impact of a ML application. In order to avoid energy waste, developers need to be made aware of the minimum criteria for the collected data relevant to their current problem to avoid over-collection. Depending on the task, this could include the change from a deep learning approach to a model requiring less data or even a human-in-the-loop approach. This entails tools to aid developers to better understand the trade-off between accuracy, amount of data needed, and means to explore the quality of the data based on their current scenario. Further, sustainable data management also has to address potential issues with data ownership and governance. Addressing such issues early-on, requires redefining of what data quality means including the impact on the environment, as much as on the human and society it reflects. This could be addressed by developing new data visualization, data labeling, or data cleaning techniques that take into account ethical and societal concerns. In addition to addressing significant ethical issues, this could help guarantee that the data is both useable and effective for its intended users.

Building and Training the model: Using energy-efficient hardware, optimizing algorithms to reduce the number of iterations required, and using pre-trained models or transfer learning to reduce the amount of training required are all ways to optimize the training process to reduce energy consumption. However, ML experts may be unaware of such alternatives, and how to explain and visualize such differences remains an open question. Given the circular relationship between ML algorithm design, development, and deployment, it is critical to consider the re-usability and recyclability of hardware and infrastructure. Furthermore, the development of ML models that can be repurposed and adapted for different use cases. Using the prototyping concept, we see a need to create digital abstract models that represent the final ML

algorithm. It enables developers to test and refine their ideas before finalizing the system by identifying reuse opportunities while also easier considering alternatives. In terms of technological systems, this is still a neglected topic in sustainability research.




Figure 3: Suggestion of a sustainable ML Life-cycle

Deployment: It is critical to better understand the needs and limitations of both users and stakeholders involved in the deployment process before deploying ML models. Conducting usability testing on the ML model and interface, for example, can assist in identifying any issues related to ease of use, user satisfaction, or efficiency. This may entail testing with a small group of users and incorporating feedback into the design before releasing models that would otherwise have to be retrained or adapted later due to incorrect configuration or interaction design.

Monitoring: Legacy software and hardware are a large source of energy waste. While most system monitoring observes and addresses errors, it ignores the potential of improving ML model performance. We believe that additional monitoring during the life of an algorithm, with a particular emphasis on energy reduction, can aid in identifying issues or inefficiencies that may be affecting the overall sustainability in the long-term. Reducing the gap between ML model developers and users by collecting regular feedback on user needs as well as technological advancements allows existing models to remain relevant for longer while avoiding unnecessary re-runs of existing model solutions.

Recycling: We think it is important to consider several aspects in this phase: 1) Establish guidelines for retiring ML models that are transparent in terms of the system's purpose/use as well as any data related with these models, which should be handled responsibly. This may involve deleting data or securely storing it to prevent unauthorized access. 2) Enable the reuse of the results and the data of the previous running model to avoid unnecessary retraining and deployment.

4.4 Role of HCI for Sustainable Machine Learning

To address current ML models' sustainability issues, we must reconsider our understanding, selection, and maintenance of these models. There is a number of common assumptions regarding successful ML models, which are (1) a model has to fit all possible cases; (2) more data equals more precision; and finally (3) higher precision and recall are always better.

We see two major challenges with (1): developers and stakeholders are quite often unaware of a model's long-term use and application, and this ignores the constantly changing environments and needs of users,

resulting in constant adaptation and retraining throughout the entire ML life cycle. Instead, focusing on a specific use case, including end-users during the development and allowing models to learn over time or connect to other models may help to avoid such wasteful efforts. The general drive for more data (2) has resulted in an increase in the amount of training data, which necessitates extensive computing resources and can have a significant carbon footprint. We believe that data quality and relevance should be prioritized, which necessitates a better understanding of the data source, for example, by investigating alternative approaches such as human-in-the-loop data labeling. Aiming for high precision and recall in ML models (3) favors the use of complex and computationally intensive models that are difficult to scale and maintain. Instead, involving the user in model training may reduce the need for large amounts of data and computation. Overall, we think that it is not always necessary to create and train new models because improved recycling infrastructure may enable the reuse of models that have already been trained. Better documentation, explainability, and transparency during model development, on the other hand, are required to enable such practice on a broad scale.

To reduce the size and carbon footprint of ML models, some of the underlying assumptions and approaches that have traditionally guided ML development must be reconsidered. By focusing on more tailored, efficient, and sustainable approaches to ML development, we can help to reduce the environmental impact of these technologies.

Supporting democratization of ML approaches: We hypothesise that taking a more iterative human-centered approach that helps developers express their goals, explore alternative models, and explains and guides developers through the ML development process and beyond has the potential to make ML models more accessible and easier to use overall. Such democratization of ML approaches refers to the idea of providing a broader community with access to knowledge, tools and resources. This includes also promoting more ethical and responsible use of ML models. As ML becomes more widespread, we think that guiding such development in a more sustainable way and aiding user to better understand the limitations and potential biases of these models would benefit society overall. Increased understanding and accessibility would also enable users to recognize when and how to use ML models appropriately, as well as interpret and communicate the models' results to others. More research and development of new and innovative approaches to problem-solving and decision-making, however, is required to enable such future.

4.5 Conclusion

To create sustainable approaches for ML life-cycles, multidisciplinary collaboration and a multi-stakeholder strategy that considers the social, economic, and environmental implications of ML algorithms are required. As our understanding of sustainable ML approaches continues to develop, it is critical to integrate constant learning and adaptation in the ML life-cycles itself. This involves taking into account the energy efficiency of hardware and infrastructure as well as the carbon impact of ML algorithms.

Our work shows opportunities how to rethink the current ML life-cycle in a more sustainable way. We also highlight relevant aspects that developers should consider at each stage to make more sustainable decisions. In accordance with human-centered design, we will observe current practise of each of these steps in order to better understand current needs and how to incorporate more sustainability awareness into every decision.

5 Understanding Existing ML Workflows

In today's fast changing technological environment, understanding existing machine learning (ML) activities is critical. These workflows, which consist of a number of structured phases ranging from problem understanding through model deployment, provide insights on best practices, optimization strategies, and the complicated process of building efficient ML models. By delving into standard processes, one

may learn about the complexities of the ML life-cycle, allowing for more informed decisions, improving model efficiency, and enabling the effective adoption of ML solutions in a variety of applications.

5.1 Interviews

In this study, we examine how ML experts find, choose, use, and understand models in their everyday work. In particular, we are interested in specific stories about how models are selected for different problems. Our goal is to better understand the ML tasks from the user description and their methods to select ML models. With this knowledge, we intend to identify relevant characteristics of a ML project in terms of constraints and context to differentiate between different models.

Participants: We recruited eight ML experts (8 male, avg. 10.3 years of experience [SD=3.1]) through emails to local research facilities or direct contact (details in table 2). All participants are researchers who specialize in a particular ML application, such as biology, medicine, or CPUs. Participants were not paid and were informed about their rights and data usage via an informed consent form in accordance with GDPR.

Procedure: Each interview lasted one hour. All interviews were conducted in English by one of the researchers, either in person or online. Participants first confirmed the consent form, before the interviewer briefly introduced the context. In a semi-structured interview, interviewees presented their specific stories about how models are selected for different problems. Our goal is to better understand the ML tasks from the user description and their methods to select ML models.

Data collection and analysis: We ran a mixed-approach thematic analysis [23]. The top-down themes included ‘Exploring Various Approaches to Tackle the Initial Problem’, ‘Discovering multiple challenges in the learning process’, ‘Identifying Recognized Gaps through Research and Collaboration’, ‘Identifying model selection parameters during the learning process’, ‘Using Multiple Evaluation Metrics’ and ‘Prioritizing Interpretability and Explainability for Model Viability’ based on literature and interview questions. One author coded the participants’ answers based on both top-down themes and themes that emerged (bottom-up). The same author summarized closely related themes and identified patterns; a second author who was not involved in the interviews re-coded the data using the final set of themes.

5.2 Current ML practice

5.2.1 Using data formats and benchmark examples alongside problem definitions to improve understanding

The majority of our interviewees (7/8) agreed that a clear starting point is required for the learning process. However, the starting point differed between participants and can be combined. Most of them (6/8) began with understanding the problem, half of them considered the data as an a starting point (P6, P2, P3, P4). Moreover, most of them focuses on defining goals and metrics. From another side, six participants emphasize on existing knowledge and approaches. This does not require that each participant explores only one approach. In order to justify their decisions, they test with multiple techniques and switch from one technique to another. For example, P8 begin with a clear understanding of the problem, emphasize on goals and metrics and then he leverages existing work and approaches.

Understanding the problem is crucial: Most participants (6/8) agree that before going into model selection, it is necessary to first understand the problem. “If the definition of the problem is well, then it’s very easy to select the model” highlighted P8. This however requires experts to have a clear understanding of their target audience and application domain (P4). To support this process some participants start with outlining different input factors in form of sketches or presentations(P1), or by investigating the source problem further(P6). Participants frequently begin by framing the situation as a narrative, highlighting its origins and conditions (P6, P8). They then identify and interpret the type of solution that would

best solve the situation at hand based on this full description. Additionally, one participant start by dividing the problem into steps to make it more easy. The task of " From where you start when you have a Learning Problem " requires different answers. it requires the use of diagrams and mathematical representation (P1), the definition of their priorities regarding the problem (P4, P3, P5)

Starting with data and format consideration: Half of the participants rely on Data and format consideration to solve their learning problem. So, analyzing the data is an important starting point for them (P3, P4, P6). This analysis is mainly by observing the input "I think what I need to understand is what are the inputs, the data, what kind of output."highlighted (P4), its format and how it is linked to the output. Even P3 test a simple model on the dataset, to see what it needs and what kind of models can highly perform with it. In the other side, one participant started with searching for literature but based on the data observation(P2).

Leveraging existing knowledge and approaches: A small proportion (3/8) of the participants are actively looking for solutions or current achievements by referring to literature. Some participants (P2,P6), in particular, go into existing literature to find identical challenges and develop solutions. One Participant (P3), on the other hand, is primarily concerned with identifying impressive features through literature review. An interesting point has been raised, particularly by P8 and P1, is that the literature frequently presents different solutions to all difficulties and allows the extraction of varied information to develop a clearly interpretable model.

Focusing on goals and metrics: Some participants (P8, P1) underlined the need of setting a clear goal before jumping into coding. Clarifying the end aim is crucial for them. In contrast, some participants (P1, P5), prefer to prioritize the establishment of specific evaluation metrics. This strategy is extremely common when the user has a high level of experience and can explicitly define an established statistic. The fundamental metric, whether it is centered on accuracy or the need for real-time processing, is determined from the start(P8). Meanwhile, P7 emphasizes the significance of calibrated log likelihood, proposing a novel strategy that values this specific strategy. Taking a somewhat divergent approach, P2 believes in constructing an initial test model.

5.2.2 Insufficiency of training and evaluation data challenges model selection

Half of our participants faced challenges during their learning process due to a lack of data. Participants highlighted the problem of data availability issues, describing the difficulty in obtaining essential data for robust analysis or model training (P1, P7). This might be also due to the scarcity of available data in distinct domains, which probes especially challenging for interdisciplinary collaborations (P4). One participants also highlighted the lack of data and support to evaluate the outcome of a model at the point of selecting it (P3).

5.2.3 Combining Literature, LLM's and AutoML to select a suitable model

Most of interviewees in the study (5/8), stated that they used literature and other collaborative tools or features in their work, " the academic literature is usually the first thing. Twitter is another source that I use."highlighted P7, demonstrating the importance of existing knowledge and collaborative advances in their different interests. P4, for example, highlights the value of shared resources in extending and improving techniques to reach his final goal and find the good model. This shows a purposeful combination of personal and collective work to improve decision making quality. Other participants (P4, P7), on the other hand, prefer to base their work on pre-existing answers and beginning points obtained from literature. These participants can build on a strong basis by delving into a repository of current knowledge, making their work more informed and strong. P7, in particular, stands out for identifying ignored research needs in this rich variety of literature and collaboration during his model selection.

5.2.4 Balancing Performance, user-understandability, and Trust in AI Tools to improve Model Selection.

Considering user understanding of model results: When deciding which model to use, some participants (2/8) place a high value on the understanding of the models by the end-user. Notably, P5 prefers models that have essential clarity in their decision-making systems. By using such models, he hopes to ensure that end-users understand the reasoning behind the model's conclusions, making the results more transparent and trustworthy. In a different way, P1 takes a strategic approach, beginning with the selection of a baseline model. This serves as a reference or a benchmark, allowing the end user to compare the performance of the potential models to this baseline.

The selected model performs well with data: Participants (6/8) were mostly in agreement on the importance of data compatibility in the model selection process. For most of them, the most important criteria is that a model correlates well with their specific dataset. Several individuals, including P6, P7, P2, and P3, highlight the necessity of completely knowing the structure of their data. They feel that understanding the complexities and nuances of their data helps in selecting models that are able to use the knowledge contained within. Meanwhile, P1 takes an experimental position. This participant tries to find which models work best in practice by running several models on the same dataset. This repeated method provides insights into model performance under consistent settings, enabling improved decisions. P8 focuses on this concept by emphasizing the extraction of model selection parameters directly from data. This participant's approach indicates a data-driven methodology in which the properties of the data itself determine the model selection criteria.

Identifying Selection Parameters: "And how do you go about actually testing different algorithms to see which one works better?" Half of participants agreed that this is a question that primarily concerns the identification and use of particular metrics or parameters to select their model. P8's approach is fundamental, with accuracy as the primary objective. This participant considered accuracy to be a critical metric. P1's solution, on the other hand, requires a delicate balancing act. The participant is looking for models that find an appropriate balance between the complexity (or number) of parameters and overall performance. This implies that for P1, both efficiency and effectiveness are critical. P2 wanted to create uniformity to the evaluation process by defining a standard set of parameters for each challenge. Finally, highlighting that different challenges may entail different evaluation measures, P3 is willing to recalibrate and redefine the selection criteria based on the specific needs of each problem.

Trusting AI tools in decision making: Some of participants (2/8), displayed a considerable degree of confidence in AI tools, indicating a preference to rely on these systems when making decisions. For instance, P1 often consults ChatGPT, showcasing trust in its ability to provide insights or suggestions. On a similar note, P2 leans towards using AutoML, a tool designed to automate the process of model selection and tuning. This inclination demonstrates P2's belief in the tool's capability to identify optimal machine learning models and configurations for a given problem.

5.2.5 Comparing ML models against existing and self-defined metrics to evaluate suitability

Comparing Against Existing Solutions and Benchmarks: The comparison of the selected models against existing benchmarks or solutions is a key feature among the participants (4/8). This analytical approach assures that their models not only perform well in isolation, but also when compared to existing solutions. P5, for example, makes comparisons by citing academic and professional literature. This strategy emphasizes a dependence on literature and previously reported results, providing a context for evaluating the effectiveness of their model. P3's strategy include comparing the findings of their model to those of other existing solutions. He identified the relative strengths and weaknesses of their selected model by contrasting it with the results of existing models or solutions. Meanwhile, P6 focuses on real results in prediction, comparing the predicted results of his model to real values. P8 begins with a foundational

methodology, establishing a baseline model. This baseline serves as a standard or reference against which other models, including the one chosen by the researchers, are evaluated. Finally, P5 examined and compared the efficacy of each model in the same settings by running many models on the same dataset. This head-to-head comparison provides insight into each model's relative performance, ensuring a full evaluation.

Defining standard evaluation metrics: All participants agreed that model evaluation requires the use of specific metrics based on the problem and personal preferences. For example, P6 and P2 emphasize loss functions as an evaluation metric. P4, on the other hand, uses the Pearson correlation as an evaluation metric. P3's training time appears as a significant indicator, demonstrating model's efficiency. P7 uses both the F1 score and the calibrated log probability as evaluation measures in a more holistic approach. P8 identified the evaluation metrics as an implication of several aspects of model performance at the same time. While admitting the relevance of different metrics, P1 believes that the precise strategy chosen should be intimately linked to the nature and requirements of the initial problem.

5.2.6 Prioritizing Interpretability and Explainability for Model Viability

The majority of participants (5/ 8) emphasize the importance of explainability and interpretability in the model selection process. For these people, it's not just about how effectively a model works, but also about how clearly it reaches its conclusions. P2, for example, valued models that not only perform well but also provide good explanations and reasons for their results. In keeping with the interpretability theme, both P2 and P3 preferred models that developed justifications and explanations, particularly when it comes to decision-making. The goal is to create models that not only make correct decisions but also explain why they are correct. For one participant(P5), simplicity and explainability go hand in hand. This method reinforces the concept that the simpler and less complicated a model is, the easier it is to understand and explain its operations and conclusions. P4 takes a different but related approach, focusing on explainability through incremental improvement. This implies that P4 may begin with a basic model and then make incremental changes, each time improving its explainability, until an ideal balance of performance and transparency is obtained.

5.3 Conclusion

This study brings light on the field's complicated and multifaceted aspects. Experts are influenced by both established agreements and changing perspectives while dealing with varied difficulties, such as the importance of interpretability and the complexities of model selection. Furthermore, the importance of research, collaboration, and a clear starting point emphasizes the constant interaction of underlying knowledge and creativity. As the field of machine learning grows, these themes not only provide insight into existing methods, but also open the way for future discoveries and breakthroughs. We will use these insights as a baseline for our next deliverable D4.2 – Explainable AI for decision making, in which we will provide examples how to support more sustainable decision making within some of these processes.

6 Summary

This report summarises HCI and ML researchers' current understanding of sustainability regarding their community and work. We discuss how their point of view affects research concerns, methods, and objectives. It emphasises the potential for combining both fields to more holistically address sustainable machine learning goals. Following that, we propose a framework for providing a structured lens on the intersections of HCI, machine learning, and sustainability. We present representative work from the last five years to demonstrate the impact of these intersections and outline a structured framing of research areas along those intersections. Focusing on a better understanding of the sustainable ML life-cycle, we

then outline relevant aspects and considerations that developers should address, as well as how human-centered design can support such efforts. Finally, in order to lay the groundwork for creating systems that can assist ML experts in making more environmentally conscious decisions in the future, we examine current practise within these ML life-cycle steps.

References

- [1] Jörg Kienzle et al. “Toward model-driven sustainability evaluation”. In: *Communications of the ACM* 63.3 (2020), pp. 80–91.
- [2] Lon Åke Erni Johannes Hansson, Teresa Cerratto Pargman, and Daniel Sapiens Pargman. “A Decade of Sustainable HCI: Connecting SHCI to the Sustainable Development Goals”. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. Yokohama, Japan: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445069. URL: <https://doi.org/10.1145/3411764.3445069>.
- [3] David Rolnick et al. “Tackling climate change with machine learning”. In: *ACM Computing Surveys (CSUR)* 55.2 (2022), pp. 1–96.
- [4] Emma Strubell, Ananya Ganesh, and Andrew McCallum. “Energy and policy considerations for deep learning in NLP”. In: *arXiv preprint arXiv:1906.02243* (2019).
- [5] Yuanlong Li et al. “Transforming cooling optimization for green data center via deep reinforcement learning”. In: *IEEE transactions on cybernetics* 50.5 (2019), pp. 2002–2013.
- [6] Gissella Bejarano et al. “SWaP: Probabilistic Graphical and Deep Learning Models for Water Consumption Prediction”. In: *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. BuildSys '19. New York, NY, USA: Association for Computing Machinery, 2019, pp. 233–242. ISBN: 9781450370059. DOI: 10.1145/3360322.3360846. URL: <https://doi.org/10.1145/3360322.3360846>.
- [7] Soumik Saha et al. “Prediction of forest fire susceptibility applying machine and deep learning algorithms for conservation priorities of forest resources”. In: *Remote Sensing Applications: Society and Environment* 29 (2023), p. 100917.
- [8] Randy H Katz et al. “Sustainable Computing: Informatics and Systems”. In: (2010).
- [9] Christina Bremer, Bran Knowles, and Adrian Friday. “Have We Taken On Too Much?: A Critical Review of the Sustainable HCI Landscape”. In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 2022, pp. 1–11.
- [10] Eli Blevis. “Sustainable interaction design: invention & disposal, renewal & reuse”. In: *Proceedings of the SIGCHI conference on Human factors in computing systems*. 2007, pp. 503–512.
- [11] Paul Dourish. “HCI and environmental sustainability: the politics of design and the design of politics”. In: *Proceedings of the 8th ACM conference on designing interactive systems*. 2010, pp. 1–10.
- [12] Bran Knowles, Oliver Bates, and Maria Håkansson. “This changes sustainable HCI”. In: *Proceedings of the 2018 CHI Conference on human factors in computing systems*. 2018, pp. 1–12.
- [13] Carl DiSalvo, Phoebe Sengers, and Hrönn Brynjarsdóttir. “Mapping the landscape of sustainable HCI”. In: *Proceedings of the SIGCHI conference on human factors in computing systems*. 2010, pp. 1975–1984.
- [14] Alexandre Lacoste et al. “Quantifying the carbon emissions of machine learning”. In: *arXiv preprint arXiv:1910.09700* (2019).
- [15] Gro Harlem Brundtland. “Our common future—Call for action”. In: *Environmental conservation* 14.4 (1987), pp. 291–294.
- [16] Tobias Welz and Matthias Stuermer. “Sustainability of ICT hardware procurement in Switzerland: A status-quo analysis of the public procurement sector”. In: *Proceedings of the 7th International Conference on ICT for Sustainability*. 2020, pp. 158–169.
- [17] Kuntal Saroha, Sheela Sharma, and Gurpreet Bhatia. “Human computer interaction: An intellectual approach”. In: *IJCSMS International Journal of Computer Science and Management Studies* 11.02 (2011), pp. 147–154.
- [18] Philip Brey and Johnny Hartz Søraker. “Philosophy of computing and information technology”. In: *Philosophy of technology and engineering sciences*. Elsevier, 2009, pp. 1341–1407.

- [19] Gaurav Sinha, Rahul Shahi, and Mani Shankar. “Human computer interaction”. In: *2010 3rd International Conference on Emerging Trends in Engineering and Technology*. IEEE. 2010, pp. 1–4.
- [20] Tom Michael Mitchell. *The discipline of machine learning*. Vol. 9. Carnegie Mellon University, School of Computer Science, Machine Learning ..., 2006.
- [21] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [22] Michael I Jordan and Tom M Mitchell. “Machine learning: Trends, perspectives, and prospects”. In: *Science* 349.6245 (2015), pp. 255–260.
- [23] Virginia Braun and Victoria Clarke. “Using thematic analysis in psychology”. In: *Qualitative research in psychology* 3.2 (2006), pp. 77–101.
- [24] Jon Froehlich, Leah Findlater, and James Landay. “The design of eco-feedback technology”. In: *Proceedings of the SIGCHI conference on human factors in computing systems*. 2010, pp. 1999–2008.
- [25] Chris Preist, Daniel Schien, and Paul Shabajee. “Evaluating Sustainable Interaction Design of Digital Services: The Case of YouTube”. In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI ’19. Glasgow, Scotland Uk: Association for Computing Machinery, 2019, pp. 1–12. ISBN: 9781450359702. DOI: 10.1145/3290605.3300627. URL: <https://doi.org/10.1145/3290605.3300627>.
- [26] Katherine W Song et al. “Towards Decomposable Interactive Systems: Design of a Backyard-Degradable Wireless Heating Interface”. In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI ’22. New Orleans, LA, USA: Association for Computing Machinery, 2022. ISBN: 9781450391573. DOI: 10.1145/3491102.3502007. URL: <https://doi.org/10.1145/3491102.3502007>.
- [27] Eldy S. Lazaro Vasquez, Hao-Chuan Wang, and Katia Vega. “Introducing the Sustainable Prototyping Life Cycle for Digital Fabrication to Designers”. In: *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. DIS ’20. Eindhoven, Netherlands: Association for Computing Machinery, 2020, pp. 1301–1312. ISBN: 9781450369749. DOI: 10.1145/3357236.3395510. URL: <https://doi.org/10.1145/3357236.3395510>.
- [28] Thomas G Dietterich. “Steps toward robust artificial intelligence”. In: *Ai Magazine* 38.3 (2017), pp. 3–24.
- [29] Roy Schwartz et al. “Green ai”. In: *Communications of the ACM* 63.12 (2020), pp. 54–63.
- [30] Peter Henderson et al. “Towards the systematic reporting of the energy and carbon footprints of machine learning”. In: *The Journal of Machine Learning Research* 21.1 (2020), pp. 10039–10081.
- [31] Todd Kulesza et al. “Why-oriented end-user debugging of naive Bayes text classification”. In: *ACM Transactions on Interactive Intelligent Systems (TiiS)* 1.1 (2011), pp. 1–31.
- [32] Jianbo Chen et al. “Learning to explain: An information-theoretic perspective on model interpretation”. In: *International conference on machine learning*. PMLR. 2018, pp. 883–892.
- [33] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. “” Why should i trust you?” Explaining the predictions of any classifier”. In: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 2016, pp. 1135–1144.
- [34] Anja Thieme, Danielle Belgrave, and Gavin Doherty. “Machine learning in mental health: A systematic review of the HCI literature to support the development of effective and implementable ML systems”. In: *ACM Transactions on Computer-Human Interaction (TOCHI)* 27.5 (2020), pp. 1–53.
- [35] Chi Yoon Jeong and Mooseop Kim. “An energy-efficient method for human activity recognition with segment-level change detection and deep learning”. In: *Sensors* 19.17 (2019), p. 3688.
- [36] Haomin Chen et al. “Explainable medical imaging AI needs human-centered design: guidelines and evidence from a systematic review”. In: *NPJ digital medicine* 5.1 (2022), p. 156.
- [37] Albrecht Schmidt et al. “Artificial intelligence for humankind: a panel on how to create truly interactive and Human-Centered AI for the benefit of individuals and Society”. In: *IFIP Conference on Human-Computer Interaction*. Springer. 2021, pp. 335–339.

- [38] Andreas Holzinger. “Interactive machine learning for health informatics: when do we need the human-in-the-loop?” In: *Brain Informatics* 3.2 (2016), pp. 119–131.
- [39] Finale Doshi-Velez and Been Kim. “Towards a rigorous science of interpretable machine learning”. In: *arXiv preprint arXiv:1702.08608* (2017).
- [40] Xu Ma et al. “Fully convolutional network for rice seedling and weed image segmentation at the seedling stage in paddy fields”. In: *PloS one* 14.4 (2019), e0215676.
- [41] Jerry Alan Fails and Dan R Olsen Jr. “Interactive machine learning”. In: *Proceedings of the 8th international conference on Intelligent user interfaces*. 2003, pp. 39–45.
- [42] Damian A Tamburri. “Sustainable mlops: Trends and challenges”. In: *2020 22nd international symposium on symbolic and numeric algorithms for scientific computing (SYNASC)*. IEEE. 2020, pp. 17–23.
- [43] Claudia De Los Rios Perez et al. “Holistic Approach for Sustainable Adaptable User Interfaces for People with Autism Spectrum Disorder”. In: *Proceedings of the 26th International Conference on World Wide Web Companion*. 2017, pp. 1553–1556.
- [44] Christoffer A Björkskog et al. “EnergyLife: pervasive energy awareness for households”. In: *Proceedings of the 12th ACM international conference adjunct papers on Ubiquitous computing-Adjunct*. 2010, pp. 361–362.
- [45] Hussien Hegab et al. “Design for sustainable additive manufacturing: A review”. In: *Sustainable Materials and Technologies* (2023), e00576.
- [46] Paul Westermann and Ralph Evins. “Surrogate modelling for sustainable building design – A review”. In: *Energy and Buildings* 198 (2019), pp. 170–186. ISSN: 0378-7788. DOI: <https://doi.org/10.1016/j.enbuild.2019.05.057>. URL: <https://www.sciencedirect.com/science/article/pii/S0378778819302877>.
- [47] Aimee Van Wynsberghe. “Sustainable AI: AI for sustainability and the sustainability of AI”. In: *AI and Ethics* 1.3 (2021), pp. 213–218.
- [48] Jung-Pin Lai et al. “A survey of machine learning models in renewable energy predictions”. In: *Applied Sciences* 10.17 (2020), p. 5975.
- [49] Irfan Ullah et al. “A comparative performance of machine learning algorithm to predict electric vehicles energy consumption: A path towards sustainability”. In: *Energy & Environment* 33.8 (2022), pp. 1583–1612.
- [50] Roberto Verdecchia, June Sallou, and Luís Cruz. “A Systematic Review of Green AI”. In: *arXiv preprint arXiv:2301.11047* (2023).
- [51] Robert Philipp et al. “Machine learning as a service: Challenges in research and applications”. In: *Proceedings of the 22nd International Conference on Information Integration and Web-based Applications & Services*. 2020, pp. 396–406.
- [52] Carole-Jean Wu et al. “Sustainable ai: Environmental implications, challenges and opportunities”. In: *Proceedings of Machine Learning and Systems* 4 (2022), pp. 795–813.
- [53] Jason Brownlee. *Machine learning mastery with Python: understand your data, create accurate models, and work projects end-to-end*. Machine Learning Mastery, 2016.