HHAI 2024: Hybrid Human AI Systems for the Social Good
F. Lorig et al. (Eds.)

doi:10.3233/FAIA240209

Exploring Human-AI Collaboration and Explainability for Sustainable ML

EYA BEN CHAABEN [a,1]
[a] *Université Paris-Saclay, CNRS, Inria, LISN*

Abstract. 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. In order to promote energy efficiency and avoid AI-waste throughout the entire life-cycle of ML applications, developers need the ability to make environmentally conscious decisions from the start. In facts, we aim 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. By participating to the doctoral consortium, we contribute to develop a discussion on building and developing a more sustainable future.

Keywords. Human-Computer-Interaction, Sustainability, Machine Learning

1. Introduction

Sustainability has emerged as an overall concern for society [4]. While research on the intersection of Machine Learning and sustainability began already more than two decades ago [8], 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 [11].

Several research fields are involved in this effort. However, research in HCI, in particular, is at an important moment, with the potential to worsen or ease these serious concerns.

Over the years, the HCI community considered various approaches to addressing sustainability, beginning with working on individual behavior [5], then applying persuasive technologies [3], 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 [1]. 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 [11] associated with developing and fine-tuning an algorithm [6]. We explore this subject from within communities, aiming to integrate the two fields to foster more sustainable decision-making. Specifically, we employ Human-Computer Interaction (HCI) approaches to incorporate

[1]Corresponding Author: Eya Ben chaaben, eya.ben-chaaben@inria.fr

sustainable practices throughout the Machine Learning (ML) model selection and ML life cycles.

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 through these steps:

1. Develop a better understanding of how users plan, search for, and select ML models.
2. Develop new interactive methods for users to express, define, and develop relevant project descriptions using a human-computer partnership approach.
3. Develop new explainability approaches to allow intelligent systems to communicate alternative ML models in a more context-dependent manner, while taking resource cost estimates into account.
4. Develop and apply evaluation methods to determine the efficacy of the interaction from a human perspective.

2. Related Work

We start with briefly introducing current work in HCI and sustainable ML.

The perspective of HCI: Human-Computer Interaction is the study of how people design, implement and use interactive computer systems and how computers affect individuals, organizations and societies [9]. 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 [2,10].

The perspective of Sustainable ML: Within the sustainable ML and artificial intelligence (AI) community there are two main branches of work: ML for sustainability and sustainability of ML [13]. The majority of work on ML for sustainable focuses on building ML models to predict the impact of technology, such as renewable energy [7] or electric cars [12], 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 [6]. For this project, our focus is on the sustainability of machine learning. To achieve this, we employ Human-Computer Interaction (HCI) methods to promote more sustainable ML practices.

3. Research Questions and Challenges

Our work uses ML-experts processes and workflows as a baseline to develop decision support system. While the ultimate goal is to aid sustainable decision making for the general ML developer community, i.e. experts and non-experts, our first steps aims to understand and reflect expert knowledge.

The first part of the project address the challenge to assist ML practitioners in expressing the desired objectives and limitations that an ML model should address in a specific project. We aim to provide a guide to explore these alternatives based on benchmarks and simulations to reduce the SustainML's overall impact, avoiding unnecessary network architecture searches and simulations.

Our work therefore addresses three main questions in this context:

1. Does providing energy consumption feedback for each contributing factor of the model selection process reduce the energy consumption of the final model training?
2. Does providing recommendations ,e.g. model parameters or domain data, support users to align their model better to their use case?
3. Does providing energy consumption feedback for each contributing factor of the model selection process impact ML experts behavior when selecting ML models?
4. Does providing energy consumption feedback for each contributing factor of the model selection process improve the sustainability awareness of ML experts in the short-term and long-term?

4. Framing the Intersection of Sustainability with HCI and Machine Learning

As the thesis title suggests, we've identified three primary themes: ”HCI,” ”ML,” and ”sustainability.” As a result, it is critical to first identify the various points of intersection among these themes, frame them, and establish the context in which my thesis fits.

4.1. Methodology

To create a framework that outlines different research areas where sustainability intersects with HCI and ML, we gather data from questionnaires with eight HCI experts, interviews with thirteen ML experts, and existing literature. This enables us to identify current trends and areas of focus that already exist and propose new directions sparsely represented across domains.

4.2. Results

According to the study, researchers are generally aware of the environmental impact of technology. HCI researchers demonstrated a thorough understanding of sustainability throughout the technology life cycle, focusing on material waste and end-user hardware consumption. This is consistent with current trends in sustainable HCI, emphasizing individual researcher responsibility and prompting a reevaluation of research objectives and methodologies. Traditionally, sustainable ML focused on model and hardware impacts, but recent developments show that ML experts are addressing broader sustainability issues. Classifying academic work by research stance helps to identify gaps and standardize approaches across disciplines. While some research areas have received extensive attention, others remain underexplored.

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.

5.1. Interviews

In this study, we examine how ML experts find, choose, use, and understand models in their everyday work. 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.

Study procedure:
Thirteen ML experts (11 male, 2 female, avg. 8 years of experience [SD=3.1]) were recruited via email or direct contact from local research facilities. They participated voluntarily, acknowledging their rights and data usage under GDPR through informed consent. Data collection involved mixed-approach thematic analysis, with top-down themes derived from literature and interview questions.

5.2. Current ML practice

Regarding the impact of ML on sustainability, our findings indicate varied considerations among participants, with some actively mitigating environmental impacts, others monitoring their model's carbon footprint, and some not prioritizing sustainability due to perceived insignificance. Evaluation of both model and infrastructure impact highlighted challenges in measuring carbon emissions and emphasized the need for comprehensive assessment, including hardware and water consumption in data centers. The trend towards larger ML models emerged as a sustainability concern, raising questions about the necessity and environmental impact of increasingly complex models. Additionally, the study explored different perspectives on the intersection of ML and sustainability, encompassing environmental, social, and economic factors.

The study on model selection strategies revealed various approaches among participants. A clear starting point was crucial, though participants differed in whether they focused on understanding the problem, analyzing data, defining goals, or leveraging existing knowledge. Insufficient training and evaluation data posed challenges for half of the participants. The integration of literature, Large Language Models (LLMs), and AutoML was vital, blending personal and collective resources for decision-making. Participants balanced performance with user understandability, trust in AI tools, and data compatibility. Interpretability and explainability were critical for model viability, with most participants valuing transparency alongside performance.

6. Supporting Sustainable Decision Making in ML Practice

In this section, we go deeper into the insights extracted from our previous qualitative studies and discuss the implications they have for design.

6.1. Design Implications

Based on our thematic analysis of the interviews, we extracted design implications for supporting ML developers to make more sustainable decisions throughout the ML-learning life cycle. In Fig. 1 we have included design considerations aligning with the ML-life cycle.

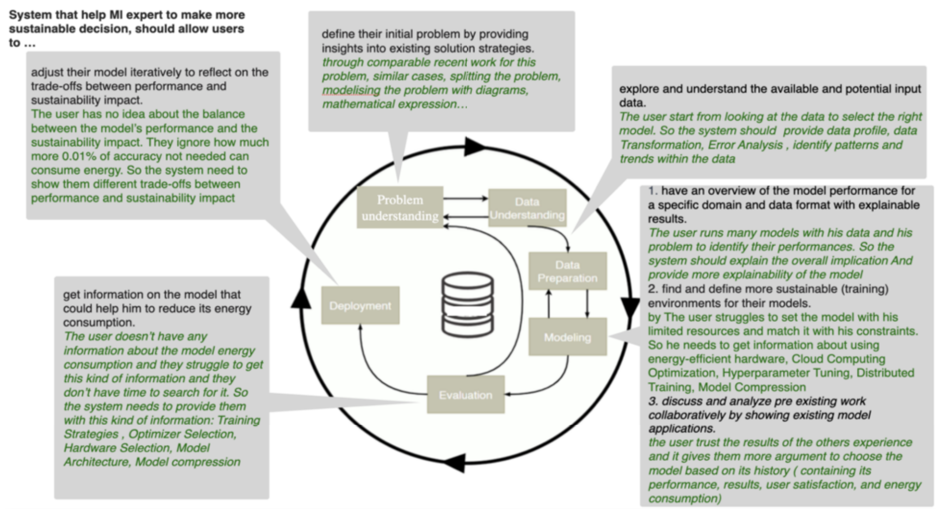

Figure 1. Overview of Design Implications based on our qualitative investigations

Problem understanding
DI 1: System that help ML expert to make more sustainable decision, should allow users to define their project needs and constraints. This can be achieved by presenting relevant recent work related to the problem, showing similar case studies for reference, decomposition of the problem into simple tasks and facilitating the visualization of the problem through diagrams or mathematical expressions.

Data understanding & preparation
DI 2: System that help ML expert to make more sustainable decision, should allow users to explore and understand the available and potential input data. Our interviewed ML experts usually begin their model selection process by examining the data. Such system should offer essential features like data profiling, data transformation capabilities, error analysis tools, and functionality to identify patterns and trends within the data. We need to develop visualizations and exploratory data analysis tools that generate at the same time feedback of key sustainability indicators, such as energy consumption patterns, resource usage trends, or environmental impact assessments.

Modeling
DI 3: System that helps ML expert to make more sustainable decision, should allow users to have an overview of the model performance for a specific domain and data format with explainable results. Systems should display different key performance metrics for each tested model to let the user compare and evaluate the effectiveness of different models. It

should also offer detailed explanations for model predictions and outcomes, highlighting the factors influencing model decisions.

DI 4: System that helps ML expert to make more sustainable decision, should allow users to find and define more sustainable computing environments for their models. Users have to consider constraints such as limited resources and the need to optimize model performance within these limitations. To address this, the system should provide information about the energy efficient hardware options, optimizations strategies, hyperparameters tuning techniques distributed training capabilities, and model compression methods.

DI 5: System that helps ML expert to make more sustainable decision, should allow users to analyze pre-existing work (collaboratively). Our interviewees expressed the practice to review results and experiences of others, which gives them valuable insights and additional arguments for choosing a model based on its historical performance, results, user satisfaction, and energy consumption. To address this, the system should offer community forums and discussion boards where the users can share their experiences, insights, and recommendations regarding model applications and their sustainability implications or establish model repository with user ratings and reviews.

Evaluation
DI 6: System that helps ML expert to make more sustainable decision, should allow users to get relevant information related to their needs and goals of models that could help him to reduce its energy consumption. Users are often unable to obtain detailed insights into a model's energy consumption without investing large efforts and often do not have the time to conduct deep studies. System should therefore help developers to access this information quickly. This includes information on a variety of factors that influence energy consumption, such as training strategies, optimizer selection, hardware compatibility, model architecture, and model compression techniques.

Deployment
DI 7: System that helps ML expert to make more sustainable decision, should allow users to adjust their model descriptions iteratively to reflect on the trade-offs between performance and sustainability impact. Users are often flexible about some model specifications, which however can be different depending on the project, application domain or outside parameters such as human supervision. Demonstrating various trade-offs between performance metrics and sustainable impact such as providing energy consumption implications of small improvements in accuracy, e.g. a 0.01% gain, can help developers to make more informed decisions.

7. Discussion and Future work

First, we want to gain a better understanding of ML developers' workflows, specifically planning, searching, and selecting ML models. To help developers make environmentally conscious decisions from the start, we conducted interviews and observational studies. These insights are used used to develop tools, strategies, or interventions to help ML developers effectively integrate sustainability considerations into their work processes. Furthermore, we may discover areas where additional research or resources are required to address knowledge or practice gaps related to long-term machine learning development. Then we work on new explainability approaches for intelligent systems that communicate alternative ML models contextually while taking resource costs into account. Additionally, we want to create interactive visualizations for developers and intelligent systems to explore existing and energy-efficient ML models, thereby encouraging the use of environmentally friendly algorithms.

We are currently in the implementation phase of the prototype. We focus on establishing the frameworks for both the frontend and backend components for our system. In order to determine if and how our prototype supports and changes developers' decision-making processes in comparison to current processes, a structured observation with 12 ML developers will be conducted to evaluate the final prototype.

8. Acknowledgements

I express my deepest gratitude to my doctoral supervisors, Prof. Wendy Mackay and Dr. Janin Koch for their invaluable guidance throughout my research journey. Lastly, I acknowledge the generous financial support provided by European Union's Horizon Europe research and innovation programme (HORIZON-CL4-2021-HUMAN-01) under grant agreement No 101070408, project SustainML (Application Aware, Life-Cycle Oriented Model-Hardware Co-Design Framework for Sustainable, Energy Efficient ML Systems).

References

[1] 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*, pages 1–11, 2022.

[2] Philip Brey and Johnny Hartz Søraker. Philosophy of computing and information technology. In *Philosophy of technology and engineering sciences*, pages 1341–1407. Elsevier, 2009.

[3] 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*, pages 1975–1984, 2010.

[4] Jörg Kienzle, Gunter Mussbacher, Benoit Combemale, Lucy Bastin, Nelly Bencomo, Jean-Michel Bruel, Christoph Becker, Stefanie Betz, Ruzanna Chitchyan, Betty HC Cheng, et al. Toward model-driven sustainability evaluation. *Communications of the ACM*, 63(3):80–91, 2020.

[5] Bran Knowles, Oliver Bates, and Maria Håkansson. This changes sustainable hci. In *Proceedings of the 2018 CHI Conference on human factors in computing systems*, pages 1–12, 2018.

[6] Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. Quantifying the carbon emissions of machine learning. *arXiv preprint arXiv:1910.09700*, 2019.

[7] Jung-Pin Lai, Yu-Ming Chang, Chieh-Huang Chen, and Ping-Feng Pai. A survey of machine learning models in renewable energy predictions. *Applied Sciences*, 10(17):5975, 2020.

[8] David Rolnick, Priya L Donti, Lynn H Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, et al. Tackling climate change with machine learning. *ACM Computing Surveys (CSUR)*, 55(2):1–96, 2022.

[9] Kuntal Saroha, Sheela Sharma, and Gurpreet Bhatia. Human computer interaction: An intellectual approach. *IJCSMS International Journal of Computer Science and Management Studies*, 11(02):147–154, 2011.

[10] Gaurav Sinha, Rahul Shahi, and Mani Shankar. Human computer interaction. In *2010 3rd International Conference on Emerging Trends in Engineering and Technology*, pages 1–4. IEEE, 2010.

[11] Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for deep learning in nlp. *arXiv preprint arXiv:1906.02243*, 2019.

[12] Irfan Ullah, Kai Liu, Toshiyuki Yamamoto, Rabia Emhamed Al Mamlook, and Arshad Jamal. A comparative performance of machine learning algorithm to predict electric vehicles energy consumption: A path towards sustainability. *Energy & Environment*, 33(8):1583–1612, 2022.

[13] Aimee Van Wynsberghe. Sustainable ai: Ai for sustainability and the sustainability of ai. *AI and Ethics*, 1(3):213–218, 2021.