SustainML

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

Explainable AI for decision making (Initial)

Deliverable D4.2

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

eProsima The Middleware Experts

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:	**D4.2**
Title:	**Explainable AI for decision making (Initial)**
Month:	18
Work Package:	WP4 - Interaction and User Studies
Work Package leader:	Inria
Deliverable leader:	Inria
Deliverable type:	R
Dissemination level:	Confidential (CO)
Date of submission:	2024-04-01
Version:	v2.0
Status:	Finished

Version history

Version	Date	Responsible	Author/Reviewer	Comments
v1.0	18-03-2024	Inria	Inria	
v2.0	01-04-2024	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 relates to *Deliverable D4.2 - Explainable AI for decision making (Initial)* of the SustainML project, which is the first of two deliverables concentrating on the creation of an interactive system that is more context-dependent. This will help ML professionals to better comprehend trade-offs relevant to the current task. The ultimately objective of this work package is to deliver a design, working prototype, and online tutorial that explain our process so that other AI researchers can apply this type of human-computer collaboration approach.

This deliverable expands on our previous work (see *D4.1 - Exploring new ML models*), where we reported qualitative studies to better understand the awareness of ML and HCI experts on their impact on sustainability, as well as their existing workflows. The report is organized into four parts: First, we discuss our research question that guides the work presented in this report, expanding on the insights and design implications we derived based on our previous qualitative work (section 2). In section 3, we explain how we used these requirements to create design rationales for the future prototype. In the third part (section 4), we discuss how this prototype is embedded and linked to the SustainML platform. We conclude with an outlook (section 5) on the current development status and discuss its relevance to the proposed tasks in WP4 overall.

This deliverable provides the first step towards developing such tools, which will be described, evaluated and applied in later deliverables.

Contents

Acronyms

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

1 Introduction

Adopting greener, less energy-consuming models by ML practitioners is crucial to effectively reduce overall AI-waste. Nevertheless, there are numerous alternative algorithms available to tackle a specific set of questions and tasks. Selecting a machine learning model for a particular task is hence typically influenced by past experience or the presence of comprehensive documentation or guidance. Embracing novel approaches often requires investing extra time and effort to grasp their purpose, capabilities, features, and suitability for a given task. Our focus is on developing interactive tools that assist ML practitioners in better understanding the impact of algorithms and their work on AI-waste as well as guiding them to find more suitable sustainable models. Two key challenges arise in this context: (1) assisting ML practitioners in expressing the desired objectives and limitations that an ML model should address in a specific project, and (2) aiding them in evaluating and choosing the most environmentally friendly ML model that aligns with those objectives.

(1) When it comes to describing the desired characteristics of a machine learning model and expressing the goals of a project, users often find themselves facing the challenge of articulating potentially ambiguous or contradictory objectives. Our goal is to create a system that enables machine learning experts to efficiently express the objectives and limitations of their projects. This includes parameters like the maximum training time, the amount of training data, and whether to prioritize speed over accuracy. From our interviews and surveys, it became apparent that the information available at this early stage can differ significantly from one project to another. Experts have also noted a lack of consideration for objective necessities, such as attempting for the highest level of accuracy by utilizing the most advanced models, without reflecting on whether it is truly necessary for the specific application at hand. Another main challenges is that these goals and constraints vary depending on the project at hand. Developing a flexible approach for machine learning projects requires acknowledging that objectives can vary depending on the domain and application context. These objectives may not always be describable or predictable in advance. A system designed to facilitate sustainable decision making should assist in the exploration and adjustment of these variables when trying to assess trade-offs and provide context dependent information regarding the task or domain of use.

(2) Understanding the described characteristics and constraints of a given project will assist the user in choosing a suitable ML model. Nevertheless, there are numerous models available for a specific problem, each with distinct features that result in trade-offs, including their CO2 footprint. To assist users in evaluating these trade-offs and making an optimal and energy-efficient decision, we aim to develop interactive systems. This system is designed to assist both experts, and ultimately non-experts, in exploring trade-offs and providing explanations about the implications of their choices within the project's context. Existing AI solutions for 'human-in-the-loop' interaction tend to oversimplify the user's role by reducing it to mere data entry for algorithmic instructions, without considering their technical expertise or leveraging it for learning purposes. We think such a interactive system should rather assist in the exploration and adjustment of variables when trying to assess trade-offs of existing alternatives, instead of providing an "optimal" decisions, empowering users to make the final decision. Our goal is to develop interactive systems that enable developers to explore and understand the trade-offs involved in different machine learning models. This includes considerations such as time, hardware requirements, and estimated CO2 footprint. By working closely with developers, we aim to promote the discovery and adoption of more environmentally friendly alternatives.

This deliverable expands on our previous work (see *D4.1 - Exploring new ML models*), where we reported qualitative studies to better understand the awareness of ML and human-computer interaction (HCI) experts on their impact on sustainability, as well as their existing workflows. In this report, we will discuss the next phase of the design process for the prototype. The following report, *D4.3*, will provide a detailed account of the final implementation and study, along with a prototype of the system. First, we describe the insights and design implications that we have derived from our previous qualitative work (section 2) followed by our research objectives which forms the foundation for the work presented in this report. In section 3, we provide a detailed explanation of how the requirements were utilized to formulate design rationales for the upcoming prototype. In section 4, we will explore the integration of this prototype with the SustainML platform. In section 5, we discuss its relevance to the proposed tasks in WP4 overall and provide an outlook on the current development status.

2 Supporting Sustainable Decision Making in ML Practice

In this section, we will go deeper into the insights extracted from our previous qualitative studies and discuss the implications they have for design. Next, we will outline the subsequent stages of this process and shed light on our research focus within this particular context.

2.1 Design Implications

In *D4.1* we reported our insights from eight interviews with ML-experts. In the meantime we extended this pool to thirteen interviews (11 male, 2 women, avg. 8 years of experience [SD=3.1]) aiming to include a larger diversity including more female developers as well as more diversity of years of experience. All participants are researchers or industry developers who specialize in a particular ML application, such as biology, medicine, or CPUs. We will not go into detail here about the additional collected data, but can share our paper in submission that includes those for further details.

Based on our thematic analysis of these data we extracted design implications for supporting ML developers to make more sustainable decisions throughout the ML-learning life cycle. In Fig. 1 we show an overview of these implications aligning with the ML-life cycle presented earlier. For each stage of the life cycle, we have included design considerations which we will describe in more detail below.

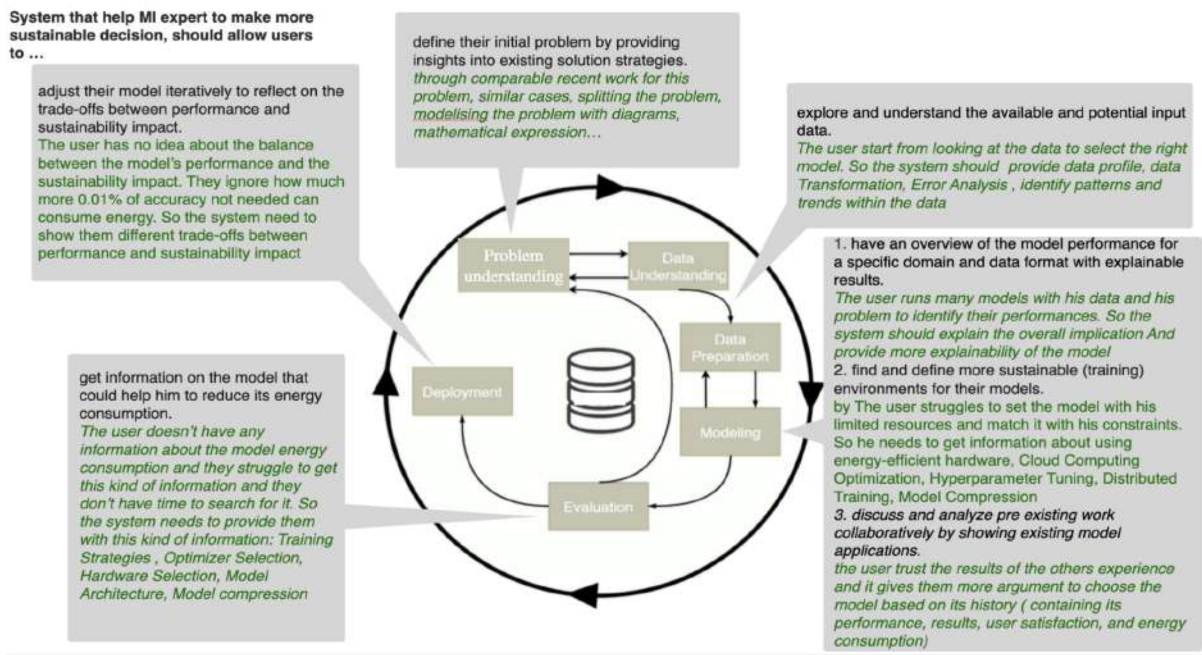

Figure 1: Overview of Design Implications based on our qualitative investigations

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 within the specified domain and data format. 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. By offering this kind of recommendations in relation to their sustainability impact, users can effectively navigate the complexities of configuring sustainable computing environments for their ML models.

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. By providing comprehensive insights into these factors, the system enables users to make informed decisions about energy efficiency in their ML models, ultimately contributing to more sustainable machine learning practices.

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. By providing insights into how changes in model configurations or algorithms impact both performance and energy consumption, users can make more informed decisions to achieve the best balance between model effectiveness and sustainability.

2.2 Research Objective

As a reminder 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. We will evaluate our prototype in later stages with both, experts and non-experts, to gather information about the suitability of the interactive system for the project as well as the further needs of non-experts.

The first part of the project will address the challenge to assist ML practitioners in expressing the desired objectives and limitations that an ML model should address in a specific project. As described in DI1-DI4, these earlier stages of the process require a deep understanding of the available knowledge as well as the needs information available at the time. 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. Based on these initial specified needs the SustainML ecosystem will explore potential models and their sustainability impact.

Our work therefore addresses three main questions in this context:

1. Does providing domain specific baselines help ML experts better define their project needs?
2. Does making the interconnections between data, project goal, and model environments and their impact on sustainability visible encourage ML developers to make more sustainable project definitions?
3. Does providing information regarding the source of energy consumption impact the long-term sustainable behavior of ML experts when selecting models?

In the future part of the project we will focus to assist users in evaluating these trade-offs and making an optimal final energy-efficient decision. This stage will be based on the simulations from the SustainML ecosystem and will allow user to interactively explore existing ML model alternatives for the current task. The ultimate goals is to combine both into one interface using consistent and transparent interaction design.

3 Prototype Design

The overall sustainability impact of a ML project can be divided into: the required hardware e.g. data centers; the model training; and the model inference. Our aim is to assist ML developers in gaining a deeper understanding of specifications and finding sustainable solutions before training the model itself. Systems that support developers with this task inquire about and comprehend project needs, ultimately suggesting suitable algorithms to meet those needs as well as suitable alternatives. In the SustainML proposal, two main challenges are identified in this context: Assisting ML practitioners in articulating the desired objectives and limitations of an ML model for a specific project, as well as aiding them in identifying the most environmentally friendly ML model that aligns with those objectives before starting the training process.

3.1 Structure and Process

One of the primary goals of our conducted interview analysis was to identify what factors ML experts take into account in their initial decision when choosing a model approach, as well as what information is available at the time of this choice. We triangulated this information with relevant literature on sustainable (HCI for) ML and engaged in discussions with other consortium partners involved in the SustainML platform. In order to recommend appropriate models with the SustainML platform, we needed to know which factors ML experts think are crucial when choosing a model, which features of a model affect sustainability, and what details about the project we needed to gather to make these suggestions.

We focused on the three main aspects that have an overall impact on the sustainability evaluation of a ML project [1]: data, problem, environment (see Fig. 2). We will further outline below how we plan to support developers explore these aspects of their project.

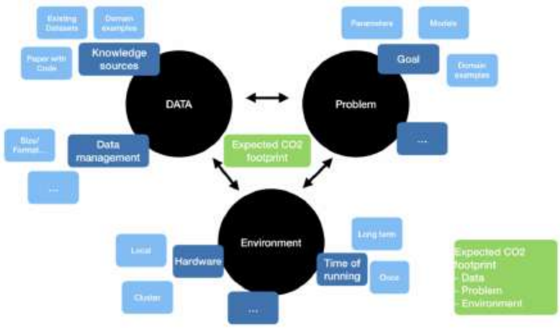

Figure 2: Overview of Design Structure

3.1.1 Data

This part relates to the selection or provision of data for the project. From our interviews, it became evident that the availability of data can vary significantly. To address this, we provide developers with the option to upload their own data and also search for existing datasets based on domain, use case or type of data. This aligns with our suggestion regarding the promotion of sustainable ML life-cycles, where data recycling can play a crucial role. Users should have complete control over the data recording process and the ability to modify, delete, expand, or analyze the data in a manner that suits their needs. The information gathered and analyzed from the data is then utilized to: assess the sustainability impact of these data using common storage structure (if not specified otherwise); and provide information for other aspects of the system such as project description.

3.1.2 Problem

The problem description varies depending on the project's objectives. Although the main objective of the project, such as image classification or data prediction, was usually well-known by our interviewees, the specific parameters and specifications for the final model were often rather ambiguous. Our system will help developers to explore suitable specification by limiting the options based on the overall goal, helping the developer to focus on relevant decisions. The relevant information to specify, depending on the model goal, is based on the knowledge graph built by DFKI as part of the SustainML project. At the same time we also allow access to comparable domain or goal parameters to provide more guidance.

By analyzing the provided information, any inconsistencies in the specifications and discrepancies across the *Data* and *Environments* parts are brought to attention. In such cases, alternative inputs may be recommended. Using the provided information, we can also provide estimates for the energy consumption needed to operate the current model, as well as suggest alternative options based on these estimates. The estimates will be improved by collecting information processed by the SustainML pipeline resulting in a benchmark.

3.1.3 Environment

This part relates to the training environment of the model. Despite the potential significance of hardware and ML environment on the sustainability of a model, only a handful of interviewees acknowledged its impact. There is a wide range of options available nowadays for running ML models, which include different hardware requirements like GPUs. However, not everyone has access to these facilities. To make this prototype more applicable to the general ML developing community, we will provide a number of options that can be simulated by the SustainML pipeline.

The system will operate using three primary parts: data, problem description, and computing environment. The user's interaction will focus on collecting information regarding each of these parts. Throughout each of these parts, we will emphasize the current impact on sustainability and present alternative options for the selected aspects. These alternatives will be supported by information from relevant projects in the domain or task. Users have the freedom to begin with whichever part they find most suitable, enabling them to explore many different project characterizations using the system's suggestions and information.

3.2 Design rationals

In this section we outline our design rationales to better support these aspects, to be precise we focus on (a) *providing and suggesting information*; (b) *making interconnections visible* to address the first challenge. In order to help ML designer explore greener models we explored how (3) *deconstructing sustainability impact*; and (4) *support decision making* through visualizing relevant trade-offs can help ML developers to make more sustainable decisions in the long term.

3.2.1 Providing vs. suggesting Information

During our interviews, machine learning experts shared a wide range of initial requirements for a new project. Some had previously gathered data and knew exactly what strategy they intended to take; others, on the other hand, had more open-ended questions they wanted to investigate without as much defined framing. A system designed to assist ML experts in making informed decisions regarding sustainability prior to project initiation must not only enable them to present established information, but also facilitate interactive exploration of requirements. Hence we think from a design perspective, it is important for such an interactive system to enable users to input their requirements *and* offer requirement suggestions based on similar cases or domains to allow this flexibility in application.

3.2.2 Making interconnections visible

Understanding the connections between the necessary data, the desired goal requirements, and the environment could assist ML developers in exploring the many choices available for a particular task. During our interviews, it became apparent that some ML experts struggled to articulate their evaluation objectives and requirements. Instead, they preferred to experiment with multiple machine learning models in order to discover one that aligned with their needs. However, there reported uncertainty regarding whether these models were truly optimal and energy efficient. Minimizing the environmental impact of ML models extends beyond just training the final model, but also encompasses model comparison and selection. For ML developers looking to make more sustainable decisions, it is crucial to have a tool that can assist in specifying requirements, identifying alternatives, and highlighting conflicting information throughout the project description phase (before training). By comparing selected requirements to existing benchmarks and use cases, one can e.g. analyze potential trade-offs and minimize the trial and error involved in training models that may ultimately be rejected. This could further help to reflect on the necessity of the provided information or their contradiction.

3.2.3 Deconstructing sustainability impact

We observed a wide range of awareness of the environmental effect of machine learning models. Although some ML experts reported to be aware of the negative effects of machine learning on the environment, we observed a lack of knowledge in (a) accurately identifying the source of the issue—such as data center cooling or hardware manufacturing—and (b) quantifying the extent of the problem—such as regionally dependent energy sources. To encourage more environmental awareness among ML experts and achieve a lasting impact, it is crucial to not only assist with immediate decision-making but also educate experts about the complex nature of sustainability impacts. For instance, opting for complex models typically corresponds to handling a larger amount of data, which in turn necessitates storage and cooling capabilities, as well as longer runtime and more advanced hardware requirements. Therefore, it is crucial to consider a multitude of factors when making such decisions. As mentioned in our previous report, there are additional dimensions of sustainability that go beyond energy consumption. Although the current prototype primarily emphasizes this aspect, we envision the potential to expand our approach to encompass other factors in future efforts. Overall empowering ML developers with the necessary knowledge to critically evaluate their work themselves can lead to significant positive impact in the future.

3.2.4 Supporting decision making vs. providing optimized solution

The selection of a ”greener” model is a multi-objective process that results in an optimal Pareto front based on the available data, needed model performances, and existing environment. While many persuasive systems for sustainable technology, particularly in HCI, have attempted to modify human behavior to an ideal option, we instead strive to educate and highlight alternatives on this front, allowing the user to make the decision for themselves. We view the user of our systems as a knowledgeable agent who possesses an extensive amount of information on how to adapt goals and make decisions in this domain, which may not be fully possible or necessary to provide to the system. Reducing accuracy by 5%, for example, can significantly affect data, runtime, and model requirements. However, whether this is appropriate in the context of the task or goal is determined by a variety of criteria, including the ability to include human supervision or the system's safety criticality. The choice of which trade-offs would best satisfy the project objectives while being more sustainable than other alternatives, in our opinion, belongs entirely to the developer.

3.3 Paper Prototypes

In an iterative process we developed a number of paper prototypes based on the users described workflow enriched with the above mentioned design rationals and system requirements. We further stared gathering information regarding the systems needs, such as benchmarks.

Figure 3.left shows some of the interfaces the system will contain. It allows developers to see all three interconnected main components, each representing a distinct aspects of the machine learning (ML) model decision aspect, and the progress of their specification as well as the current level of energy consumption

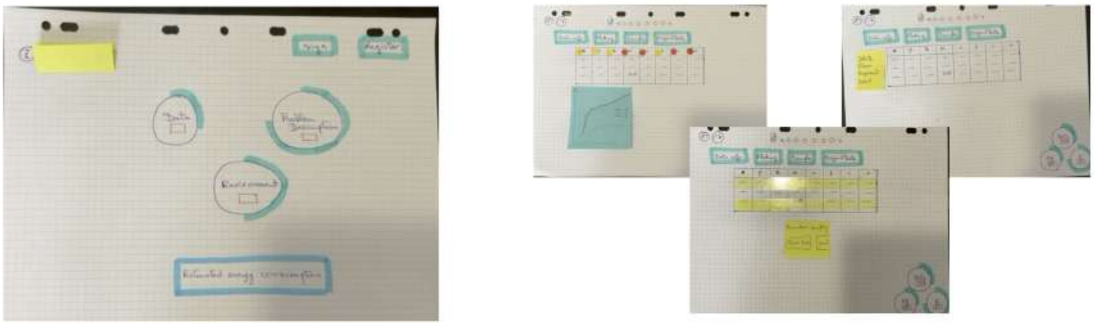

Figure 3: Paper prototypes: overview screen (left), data analysis susection (right)

of the current model and existing alternatives. We anticipate that the selection process will not necessary be done in one session, and our system will support this need by allowing project accounts.

During the data part, the system assists users in manipulating and preparing their data, as illustrated in the Fig. 3.right. We further support this process through recommendations, feedback mechanisms, and ranking techniques related to the data at hand, comparable projects as well as information retrieved through other aspects of the project description such as the project goal description. Similarly, in the problem description part, the system enables users to articulate their initial problems by suggesting a number of required, and optional information. This functionality extends to the computing environment part as well. Each component contains a information about the sustainability impact and provide further explanations in regards of alternatives.

4 Embedding of the prototype in the SustainML platform

The interactive system will serve as the SustainML pipeline's entrance point. This means that we must triangulate the demands of ML developers with insights from research on sustainable ML and HCI, as well as the information requirements of the whole SustainML pipeline.

To facilitate this process, we require more specific information on the data requirements and output generated by each of the SustainML pipeline's subsequent stages. We began a number of exchanges with the various partners and entities in order to obtain an overview of the necessary information as well as information that is currently or will be accessible to us in the near future. We will quickly address the following user-relevant features, whereas EProsima will go into further detail in their deliverable about the links between the various pipeline components (see Fig. 4).

4.1 Input source for defining initial model requirements

The user interface has to serve input information to: the Model Provider to evaluate the current task; the application-level requirements to find the most suitable hardware; and finally to the carbontracker to identify the source of energy used by the model training and application - giving insights beyond the energy consumption.

The user interface will adapt the needs of information to the knowledge graph provided by DFKI for the ML task encoder. To be more precise, depending on the project goal, e.g. classification or prediction, different information are required to make a suitable suggestion. The interface including suggestions, alternatives, and explanations will adapt to these requirements. We further discussed with the hardware providers in more detail which hardware is currently available for evaluation by the pipeline as well as

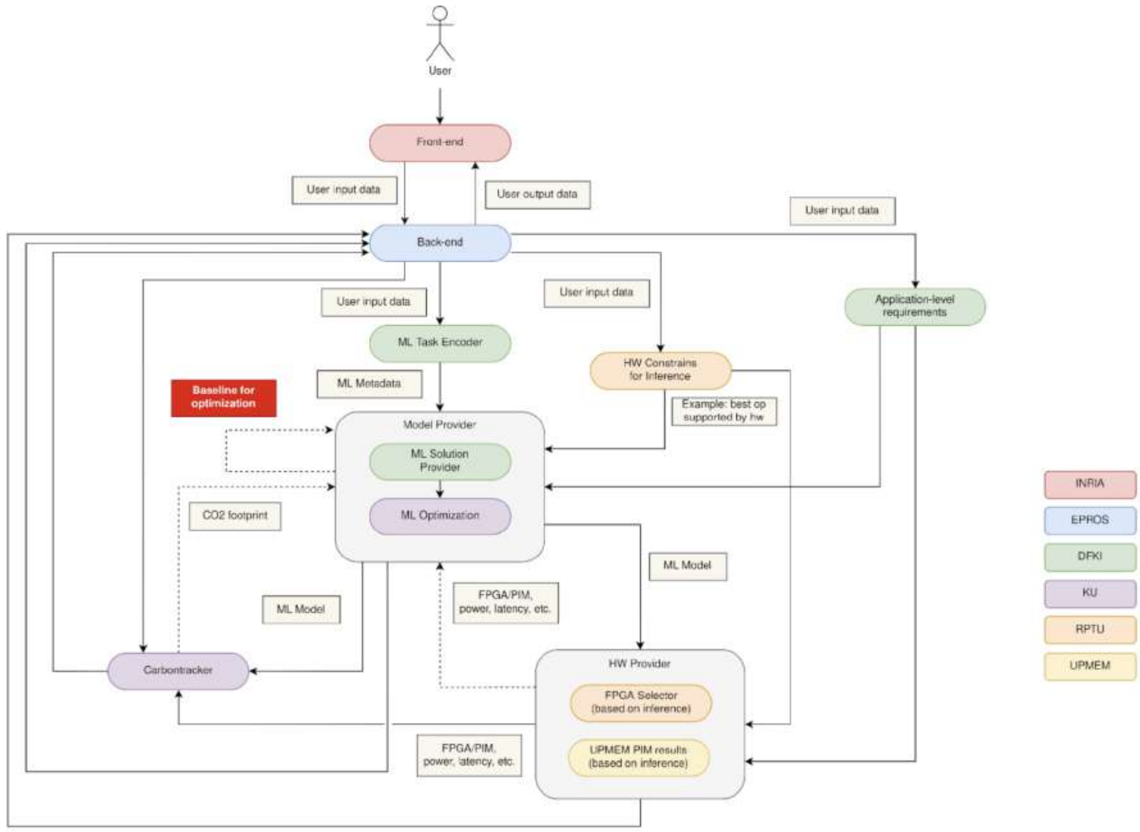

Figure 4: SustainML platform interconnections

how to extend it in the future. In return we will use their collected data of energy consumption of certain hardware as a baseline for our initial prediction of hardware consumption.

4.2 Output source for providing guidance and decision support

As the goal of the SustainML platform is to avoid AI-waste overall, we intend to provide as much information and decisions support before a model is trained and eventually even data are collected. For this reason the user interface will operate based on estimate information we receive from the connected modules, which will be updated over time allowing to provide pre-training information without the need for running full model simulations.

Once the user has provided a thorough description of the project, the platform will then forward the user data to the relevant stakeholders. Their findings will provide a range of recommendations derived from realistic simulations, which can then be presented to the developer for a more well-informed final decision.

5 Discussion

In this section, we will provide a brief overview of the goals of WP4 and explain how our current work aligns with the project goals outlined in each task. Additionally, we provide some information about the ongoing development process.

5.1 Contextualizing Report in WP4

T4.1 Developing ML project descriptions using human-computer partnership [M1-M24] In this task we aim to better identify and understand ML developers existing workflows to plan, search for, and select ML models in their current work process. ML developers need the ability to make environmentally conscious decisions from the start, but often have limited knowledge about all requirements. In order to assist developers with this task we conducted interviews and observational studies of the current approaches of ML professionals which are described in detail in *D4.1*. We extended our interview and literature study which resulted in a submitted publication as well as the source of information for this prototype.

T4.2 AI Explainability for Greener ML models [M1-M24] In this task we aim to create new explainability approaches that will enable intelligent systems to communicate alternative ML models in a more context-dependent manner while also considering resource cost estimates. It will take the initial project description into account and assist the developer in identifying trade-offs for the current project based on previously identified needs (T4.1), in order to make more environmentally conscious decisions. This prototype describes the first step towards such a prototype focusing on expressing the initial project description and exploring the expected trade-offs for such models. In a later step we will expand this prototype to also reflect trade-off of the final suggested alternatives resulting in a consistent, interactive access point for the SustainML pipeline.

T4.3 Exploring ML models using resource footprint estimation [M1-M24] This task aims to design and develop new interactive visualizations in which developers collaborate with intelligent systems to explore the potential space of existing and more energy- efficient ML models, and to identify more environmentally friendly approaches appropriate for their current project. This will aid in the development of new approaches and in the adoption of more environmentally friendly algorithms. This task is in parallel to T4.2, where aimed to get a better understanding on how to compute the estimates resource footprint using data from the partners as well as publicly available data as presented in [1] and similar. We continue building estimates and benchmarks based on simulations of the SustainML pipeline and plan to make them available to a larger audience after anonymizing the data.

T4.4 Design and develop a working prototype [M13-M36] The aim of this task is to incorporate results from T4.2 and T4.3 into concrete interaction and interface concepts. We plan investigate how to provide users with an intuitive, incrementally-developed representation of the available ML model space, visualizing their relationships, benefits, and resource costs in terms that the user can understand and manipulate. The presented design is the first step towards that goal.

5.2 Future Work

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. From the frontend, we use React, a popular JavaScript library known for its component-based architecture and fast rendering capabilities. For backend framework, we chose Node.js, which uses an event-driven, non-blocking I/O model to create scalable and performant server-side applications. For the different functionalities, especially the ones for data pre-processing, we use essential libraries like Pandas, which is well-known for its powerful data manipulation capabilities, Seaborn for data visualization, and PyTorch for machine learning model development. To provide a seamless integration of the gathered data into the pipeline, we work closely with EProsima who are responsible for the backend integration.

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. The study and the prototype will be included in the next deliverable *D4.3 Explainable AI for decision making (Final)*. The interactive systems will ultimately be made available via the online platform as an entry point to the SustainML discipline after being assessed in a subsequent study involving less experienced ML developers, such as ML students.

While the presented prototype mainly focuses on the initial project description, we will expand our work based on the initial design implications also on the later aspects of the ML process. This includes different visualizations and project-depend explanations of the final alternatives presented by the SustainML platform.

6 Summary

In this deliverable we report on our current progress for developing the central interactive systems for supporting more sustainable decisions during the ML development as part of WP4. We describe our design implication based on earlier conducted interviews with ML experts and outline our current research goals. In the later part of the report we discuss our design rationals and decision to develop such a system. We then outline the new SustainML ecosystem and how our work is embedded. Finally, we briefly discuss how this work is placed in the overall WP and outline current and future work on the realization of the prototype.

References

[1] Alexandra Sasha Luccioni, Yacine Jernite, and Emma Strubell. *Power Hungry Processing: Watts Driving the Cost of AI Deployment?* 2023. arXiv: 2311.16863 [cs.LG].