



Living Sustainability: In-Context Interactive Environmental Impact Communication

ZHIHAN ZHANG, University of Washington, USA

PUVARIN THAVIKULWAT and ALEXANDER LE METZGER, University of Washington, USA

YUXUAN MEI, FELIX HÄHNLEIN, and ZACHARY ENGLHARDT, University of Washington, USA

GREGORY D. ABOWD, Northeastern University, USA

SHWETAK PATEL and ADRIANA SCHULZ, University of Washington, USA

TINGYU CHENG, University of Notre Dame, USA

VIKRAM IYER, University of Washington, USA

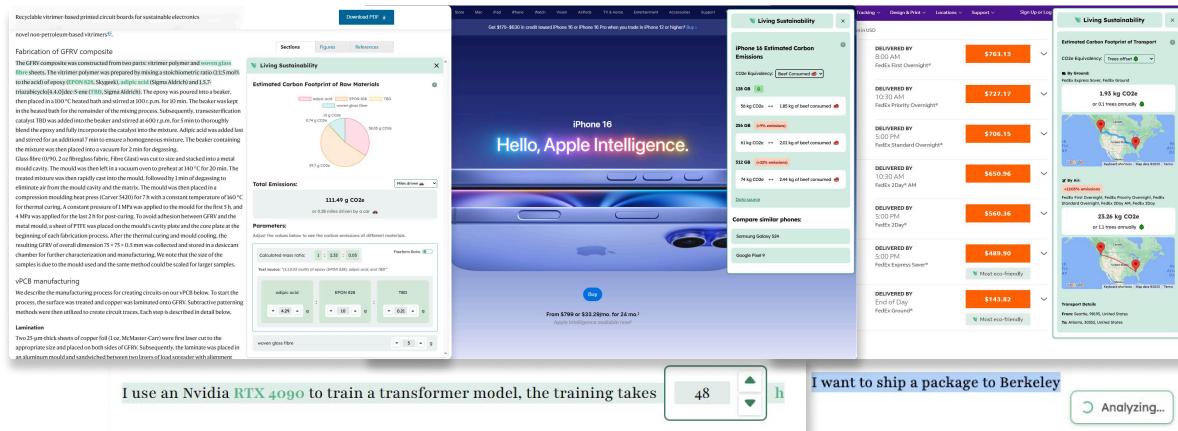


Fig. 1. Understanding the environmental impact of everyday objects and activities typically requires time-consuming expert analysis through Life Cycle Assessment (LCA). Living Sustainability is an autonomous sustainability assessment tool that transforms unstructured natural language into in-context interactive visualizations for various scenarios, including raw materials and manufacturing, embodied carbon of finished products, transport, and use phase.

Authors' Contact Information: [Zhihan Zhang](#), Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA, USA, zhihan@cs.washington.edu; [Puvarin Thavikulwat](#); [Alexander Le Metzger](#), University of Washington, Seattle, WA, USA; [Yuxuan Mei](#); [Felix Hählein](#); [Zachary Englhardt](#), Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA, USA; [Gregory D. Abowd](#), Northeastern University, Boston, MA, USA; [Shwetak Patel](#); [Adriana Schulz](#), Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA, USA; [Tingyu Cheng](#), University of Notre Dame, Notre Dame, IN, USA, tccheng2@nd.edu; [Vikram Iyer](#), Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA, USA, viyier@uw.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2474-9567/2025/9-ART153

<https://doi.org/10.1145/3749488>

Climate change demands urgent action, yet understanding the environmental impact (EI) of everyday objects and activities remains challenging for the general public. While Life Cycle Assessment (LCA) offers a comprehensive framework for EI analysis, its traditional implementation requires extensive domain expertise, structured input data, and significant time investment, creating barriers for non-experts seeking real-time sustainability insights. Here we present the first autonomous sustainability assessment tool that bridges this gap by transforming unstructured natural language descriptions into in-context, interactive EI visualizations. Our approach combines language modeling and AI agents, and achieves >97% accuracy in transforming natural language into a data abstraction designed for simplified LCA modeling. The system employs a non-parametric datastore to integrate proprietary LCA databases while maintaining data source attribution and allowing personalized source management. We demonstrate through case studies that our system achieves results within 11% of traditional full LCA, while accelerating from hours of expert time to real-time. We conducted a formative elicitation study (N=6) to inform the design objectives of such EI communication augmentation tools. We implemented and deployed the tool as a Chromium browser extension and further evaluated it through a user study (N=12). This work represents a significant step toward democratizing access to environmental impact information for the general public with zero LCA expertise.

CCS Concepts: • **Human-centered computing** → **Interactive systems and tools**; • **Software and its engineering** → **Abstraction, modeling and modularity**; **Domain specific languages**; • **Information systems** → **Information retrieval**; • **Computing methodologies** → **Natural language processing**; **Knowledge representation and reasoning**.

Additional Key Words and Phrases: Sustainable Computing, Augmented Communication, AI for Sustainability, Life Cycle Assessment, AI Agents

ACM Reference Format:

Zhihan Zhang, Puvarin Thavikulwat, Alexander Le Metzger, Yuxuan Mei, Felix Hählein, Zachary Englhardt, Gregory D. Abowd, Shwetak Patel, Adriana Schulz, Tingyu Cheng, and Vikram Iyer. 2025. Living Sustainability: In-Context Interactive Environmental Impact Communication. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 9, 3, Article 153 (September 2025), 42 pages. <https://doi.org/10.1145/3749488>

1 INTRODUCTION

Climate change presents an unprecedented global challenge, with greenhouse gas emissions continuing to rise at an alarming rate. Global carbon dioxide emissions reached 41.6 billion metric tons in 2024 [29], and a 42% reduction is needed by 2030 to stay on track for the goal to limit global warming to 1.5°C [79]. As people around the world recognize the need for action [64], there is also a growing willingness to consider environmental impact (EI)¹ in everyday decisions. Recent global surveys show that 70% of consumers are willing to pay a premium for environmentally friendly smartphones [48], and studies show products with sustainability certifications have seen sales increases [65].

Despite consumers' willingness to consider EI, such as carbon, in their decisions, this information is missing for the vast majority of products and services in our daily lives. Aside from direct emission sources such as fossil-fuel burning cars and heating systems, many products and activities cause *indirect emissions* from electricity generated by fossil-fuel power plants or high-emission manufacturing processes [35]. To accurately quantify the EI of an object or activity, one must account for its full life cycle, including the embodied emissions from raw material extraction and production, as well as the energy consumed during operation.

This work asks the question: can we enable anyone—with no background in environmental impact assessment—to interactively analyze EI at scale? Just as nutrition labels have transformed food purchasing by making complex nutritional information both available and understandable to the general public, we envision that by making EI information contextual and accessible, we can similarly guide sustainable decision-making in everyday life.

Achieving this goal is challenging because EI is traditionally modeled through Life Cycle Assessment (LCA), a manual, expert-intensive process that involves analyzing each component and process within a system throughout its life cycle [52]. LCA serves three critical purposes in addressing this challenge: 1) enabling accurate carbon

¹In this work we focus on the specific impact of carbon footprint

accounting across life cycle stages, 2) supporting decarbonization strategies during the design phase, and 3) guiding sustainable choices by surfacing lower-carbon alternatives. Currently, LCA is completely inaccessible to the public, and even infeasible for large companies to perform for every product and service. For non-experts, even seemingly simple activities pose significant modeling challenges. Creating accurate LCA models requires identifying all relevant components and processes, defining appropriate system boundaries, constructing flows, and selecting the closest matching emission factors from LCA databases for each entry. For example, accounting for transport requires knowing the endpoints and specific modes; electricity use depends on regional power grid compositions and their varying carbon intensities. Non-experts lack the structured knowledge needed to decompose such everyday objects and activities into analyzable components.

While domain-specific tools have emerged to provide EI estimates—for flights [54], driving routes [34], mechanical CAD [6], and electronic designs [40, 61, 89]—they require structured input, operate in specialized contexts, and are disconnected from everyday use scenarios. A generalizable tool that can handle diverse queries about daily activities while comprehensively covering different LCA stages remains an open challenge.

In this paper, we challenge the conventional assumption that LCA requires structured input with detailed life cycle inventory to compute EIs such as carbon footprint (CF). We present Living Sustainability, the first autonomous system that enables in-context interactive EI communication by reasoning about and augmenting static natural language content. We implement Living Sustainability as a Chromium browser extension that enables users to seamlessly view and modify sustainability information on arbitrary web content and documents.

To achieve this, we first introduce a data abstraction (DA) that represents LCA topology. This DA for LCA provides multiple benefits for large language model (LLM) agents: it enables the incorporation of expert LCA knowledge, correct by construction, interpretability, and editability. We then introduce an AI agent framework that leverages language modeling, including semantic similarity and LLMs, to transform unstructured natural language input into this DA for further modeling. The agent retrieves the emission factors from multiple sources, and employs an extensible nonparametric datastore that enables secure integration of proprietary LCA databases (e.g., ecoinvent [26]). The datastore allows for 1) the use of proprietary data without training on it, 2) attribution of data at the individual emission factor level, and 3) personalized data management that allows both LCI companies to opt out by removing databases from the store and users to add custom sources. Importantly, because of the DA, emission factors or data sources can be swapped out dynamically, without the need to reconstruct the LCA model. We note that this modularity and data security address key industry concerns about data security and intellectual property regarding design and supply chains [28].

Overall, Living Sustainability creates in-context interactive learning environments through visualizations and explorable explanations [82]. We aim to democratize access to accurate, contextual EI information for everyday objects and activities, empowering the public to make environmentally-conscious choices, while contributing to broader industrial and societal shifts toward more sustainable products and practices. We summarize our contributions as follows:

- Living Sustainability is the first autonomous EI assessment tool that transforms unstructured natural language into in-context interactive visualizations. This empowers users with *zero* LCA expertise to explore the EI of their everyday objects and activities through an intuitive interface.
- We develop an end-to-end agent workflow that transforms unstructured input into structured DA suitable for LCA modeling with >97% accuracy. Our agent can retrieve relevant emission factors from both established LCA databases and online sources such as research papers and environmental reports.
- We introduce an LCA data abstraction for simplified LCA that enables uncertainty estimation, bi-directional parameter visualization, and data source traceability. We programmatically compile LCA flowcharts from the DA and show, through case studies spanning major LCA stages (raw materials, manufacturing, transport, and use),



Fig. 2. (RE) Five life cycle stages and example LCA system boundaries. Diagrams adapted from [89]. (b) Current simplified LCA tools are limited to specific narrow scenarios and require either proprietary BOMs [56, 70] or detailed predefined options [1, 16] as input, limiting accessibility for general users.

that our simplified LCA methodology can achieve results with <11% discrepancy compared to traditional full LCA.

- We contribute empirically-derived design strategies for EI communication, based on thematic analysis from formative studies with six domain experts to gain design insights. We implement these insights in a complete system consisting of a backend server, a database, and a frontend UI with dynamic web integration. We deploy the tool as a Chromium browser extension and evaluate its usability and usefulness through a user study with 12 participants.

2 RELATED WORK

Our work draws on, and contributes to research across three key areas: life cycle assessment methodologies and tools, interactive document augmentation, and sustainable computing in HCI. Below, we present related efforts to promote sustainability awareness and democratize environmental impact assessment and communication.

2.1 Democratizing LCA for Sustainability Awareness

Life cycle assessment (LCA) provides a systematic methodology for evaluating the EIs of objects and processes throughout their existence. Traditional process-based LCA involves defining system boundaries (e.g., “cradle-to-gate” for embodied emissions or “cradle-to-grave” for full life cycle analysis), as illustrated in Figure 2a, followed by detailed inventory analysis of inputs (energy, materials, resources) and outputs (emissions, waste) at each life cycle stage. While LCA can evaluate multiple EIs, we focus specifically on CF or Global Warming Potential (GWP), measured in kilograms of carbon dioxide equivalent (kg CO₂e), as this metric has become a primary sustainability indicator for organizations worldwide [47].

Understanding the EI of everyday activities and objects is crucial for enabling informed sustainability decisions among the general public [25, 50]. Research has shown that providing clear feedback about resource consumption can promote awareness and drive behavior change [31, 32]. Traditional LCA tools and databases like Sphera/GaBi [74] and ecoinvent [26] enable detailed assessments but present significant barriers to widespread adoption. License costs can reach tens of thousands of dollars, and conducting accurate assessments requires both detailed bills of materials (BOMs) and expertise in matching components to appropriate database entries [9]. These challenges make comprehensive LCA prohibitively expensive even for large companies, resulting in limited publicly available EI information for most everyday objects and activities.

In response to these accessibility challenges, researchers and industry practitioners have developed various approaches to streamline and simplify LCA. The Product Attribute to Impact Algorithm (PAIA) [61] simplifies assessment by mapping high-level product attributes to EIs at the module level. Several industry tools [23, 56, 70] have adopted similar component-based models, while specialized solutions like ACT [37] and DeltaLCA [89] focus on specific domains such as computer system design. The Economic Input-Output LCA (EIO-LCA) [8, 40]

takes a different approach, estimating impacts based on market prices, though this often sacrifices precision due to its reliance on aggregated sector data.

Despite these advances in simplifying LCA, existing tools remain inaccessible to the general public. They still require either structured input with detailed life cycle inventory of sub-components and production steps, specialized domain knowledge, or access to proprietary design files (Figure 2b). Simple calculator applications like the Microsoft Sustainability Calculator [23] and the Idemat app [43] attempt to serve less-technical audiences, but they are limited to carbon accounting for Azure cloud workloads and material selection, and thus do not support understanding the EI of everyday activities. This gap between the need for accessible EI information and the limitations of current tools motivates our work on automated LCA.

2.2 Augmenting Communication

The challenge of communicating EI information to unfamiliar audiences extends beyond just generating accurate assessments—it requires creating engaging, understandable, and actionable presentations of that information. The HCI community has extensively explored how interactive augmentation can enhance document comprehension and learning experiences, providing valuable insights for EI communication.

Early work in augmented reading interfaces focused on enriching scholarly communication by providing contextual information about references, technical terms, and mathematical symbols without disrupting the reading flow. ScholarPhi [38] and Kim et al. [49] demonstrate how in-situ annotations and linking can reduce cognitive load by eliminating the need to switch between different parts of a document. These principles are particularly relevant for EI communication, where users often need to understand technical terminology and metrics within their original context.

Explorable explanations [82], originally demonstrated in mathematics [19, 83] and science education [24, 36, 63], allow readers to manipulate parameters and observe outcomes directly within the learning context through interactive media. Such interactive environments have proven especially effective for teaching abstract concepts that are difficult to visualize—a challenge that parallels the communication of EI metrics like carbon footprint.

Recent work in data-driven journalism [22] has shown how interactive articles with dynamic figures and embedded simulations can make complex data more accessible to general audiences. However, creating such interactive experiences traditionally requires significant programming expertise and time investment. To address this challenge, researchers have developed tools to semi-automatically augment existing documents [19, 36, 58] by creating interactive visualizations, but they focus on domains with well-established modeling tools and standardized datasets. EI communication presents unique challenges because there is no standardized format, and it requires both technical accuracy and intuitive presentation of abstract concepts that may be unfamiliar to general audiences.

2.3 Sustainable Computing in HCI and UbiComp

The increasing societal focus on environmental consciousness has resonated strongly within the HCI and UbiComp communities, leading to a growing body of work on this topic over the past decade. This includes the development of sustainable computing hardware, such as unmaking [72, 73], the use of bio-based [51, 60, 80, 81, 90] and recyclable materials [18, 88], with the aim of minimizing the environmental hazards of waste, the processing of traditional plastics, and extending the life cycle of electronic components. Additionally, innovative design tools have been created to reduce EI by integrating sustainability considerations into the early design stages [17, 89] and facilitating reuse at the end-of-life [55].

Building sustainability awareness requires more than just novel materials and end-of-life solutions, it demands tools that help people understand and act upon EIs in their daily lives. A substantial thread of Sustainable HCI (SHCI) research has examined how to effectively provide eco-feedback to users through visualization and

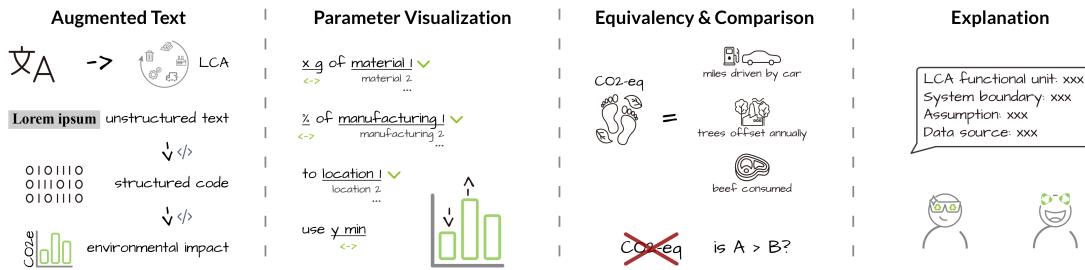


Fig. 3. Sketches of user-elicited ideas for EI communication methods from the formative study. (Left to right) augmented text in which the EI is automatically calculated and shown in context, bi-directional parameter visualization where users can interactively change quantities to see the change in EI, impact equivalency and comparison to relatable metrics like cars driven or meat consumed, as well as methodology explanations and transparent data sources.

design choices. Prior work [30, 32] established design dimensions for eco-feedback technology, emphasizing how information presentation affects user engagement with consumption data. Meanwhile, research into digital interventions [10, 11] questions their effectiveness in realizing savings due to high baseloads and rebound effects. Recent critical reflection on SHCI's trajectory [12] proposes "Green Policy Informatics" as a focused pathway to leverage HCI expertise for urgent climate action. This perspective aligns with identified opportunities for HCI engagement with public policy across domains [77] and advocates for interdisciplinary systemic understanding and the inclusion of nonhumans to broaden SHCI research [59].

EI metrics such as GWP are more abstract than direct measurements of energy and water consumption [30, 33]. Aside from gas-powered vehicles, carbon emissions are physically abstracted away from us at power plants or manufacturing facilities. Moreover, there is currently *zero* accessible mapping for how these large-scale, aggregate emissions translate to individual consumption. Our work aims to fill this gap by developing the first accessible tool that shows the carbon impact of everyday decisions like shopping using our automated, accurate EI assessment system. Our work directly addresses the call within SHCI for "the deconstruction of popular myths about the carbon impacts of everyday life" [12]. We also acknowledge that effective sustainability outcomes require broader sociocultural changes and policy interventions beyond individual decisions [59, 69].

3 FORMATIVE STUDY

To inform the design of our system, we conducted a formative study with domain experts to 1) understand current sustainability behaviors and identify gaps in existing tools and practices, and 2) gather insights into effective communication strategies through design elicitation. Details of the study method are provided in Appendix A.1.

3.1 Elicited Communication Strategies

Our formative study revealed several critical challenges in current sustainability assessment practices. Below, we present four key communication strategies, each addressing specific gaps identified, drawing from participants' experiences and needs (Figure 3).

3.1.1 Augmented Text. Unsurprisingly, all participants expressed frustration with the lack of accessible, credible tools for understanding the EI of their daily choices. While all participants showed strong interest in learning about carbon footprints, they found the current process overwhelming, with one noting, "Environmentalism is so political... I only trust research papers, but they're too time-consuming to read."

To address this accessibility barrier, we want to enable EI assessment information within users' natural reading and browsing context. Text emerged as the most intuitive input method for non-experts, as it aligns with how

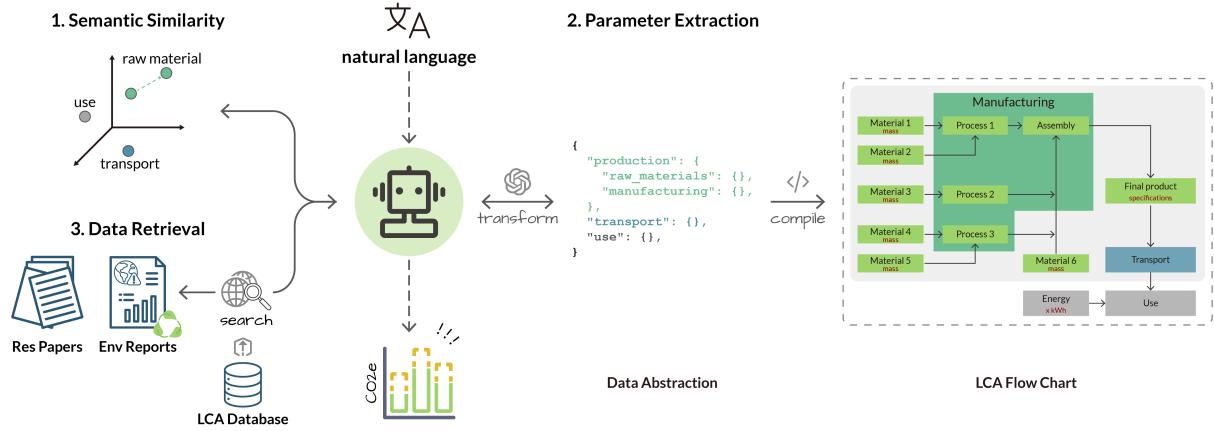


Fig. 4. Overview of Living Sustainability pipeline, the system 1) classifies input natural language description into LCA stages, 2) transforms into LCA data abstraction, 3) retrieves emission factors from databases and online sources, and finally 4) conducts life cycle assessment and creates interactive visualizations.

people naturally express their questions and scenarios. Our design enables users to either highlight text while reading online content or input free-form descriptions of scenarios they want to analyze, similar to interacting with a chatbot. This approach maintains contextual continuity, eliminating the need to switch between multiple sources [39] for understanding EI.

3.1.2 Parameter Visualization. Interactive visualization [36, 39] emerged as a crucial strategy for helping users understand the relationship between their choices and EI. For example, when discussing cloud server usage, participants wanted to see how different configurations (e.g., runtime, hardware specifications, location) affect overall emissions. Our design implements bi-directional binding between parameters and visualizations, enabling users to explore scenarios through direct manipulation. When users adjust parameters like operation time or power consumption, the system automatically updates corresponding EI calculations and visualizations.

3.1.3 Equivalency and Comparison. Participants without LCA expertise consistently struggled to interpret CO₂e metrics meaningfully. The study revealed two complementary approaches to making EI more comprehensible. First, participants strongly preferred relatable equivalencies, such as “miles driven by car”, “food consumed”, and “trees offset”, which translate abstract CO₂e values into familiar terms. Second, they emphasized the value of relative comparisons between options, noting that seeing relative differences often matters more than absolute values when making sustainable choices [89].

3.1.4 Explanation. Concerns about data reliability and greenwashing emerged as significant barriers to trust. While participants expressed strong support for universal EI metrics, our LCA expert highlighted a crucial challenge: carbon footprint numbers from different companies are often incomparable due to varying assumptions in their calculations [5]. To address these trust concerns, our design incorporates detailed breakdowns of calculations, including methodology explanations and clear attribution of data sources [4]. This transparency helps users make informed judgments about the reliability of EI assessments while promoting more sustainable choices.

latitude, longitude	::= float	coordinates ::= (latitude, longitude)
location, origin, destination	::= coordinates	
material	::= amount, unit, <i>ef</i>	
process	::= energy, <i>ef</i> , location, time	
independent_material	::= material, process*	
related_materials	::= (material, ratio)*, process*	
raw_materials	::= independent_material*, related_materials*	
production	::= raw_materials, process*	
transport_mode	::= vehicle vessel airplane	
segment	::= origin, destination, weight, transport_mode, <i>ef</i>	
transport	::= segment*	
operation	::= time, power, energy, <i>ef</i> , location	
use	::= operation*	
life_cycle_stage	::= production transport use	
Simplified LCA	::= life_cycle_stage*	

Table 1. Syntax of the data abstraction for our simplified LCA methodology. * represents the Kleene star operator, denoting variable-length sequences of items of the associated type.

4 SYSTEM OVERVIEW

In this section, we introduce Living Sustainability, an autonomous system that creates interactive, in-context environmental impact visualizations from natural language. The system empowers non-technical users to understand EIIs without requiring LCA expertise. Users simply input text they are interested in, and the system employs multi-step planning, automatically processing the text through four steps that emulate expert LCA workflows: 1) classifies text into relevant life cycle stages, 2) extracts key LCA parameters and transforms them into a data abstraction for LCA, 3) retrieves emission factors from authenticated LCA databases and online sources including academic research papers, and 4) conducts life cycle assessment and generates embedded interactive visualizations (Figure 4). This staged approach combines algorithms and workflows optimized for each step, reducing computation and latency, while constraining hallucinations compared to using a single LLM to directly generate EI values. The resulting visualizations support dynamic parameter manipulation and provide transparency through detailed methodology and functional unit explanations, and data source attribution.

4.1 LCA Data Abstraction

To bridge unstructured natural language and structured LCA modeling, we introduce an LCA data abstraction inspired by LCA-as-Code [75] and domain-specific languages (DSL) for computational fabrication [15, 85]. Our DA captures and transforms essential EI parameters for primary LCA phases into a hierarchical, machine-interpretable format while preserving the LCA topology, such as process sequences and input-output dependencies. The DA serves as both a reasoning scaffold for agents and a flexible runtime representation for real-time LCA. We designed the DA around four core benefits:

- (1) **Expert Knowledge Distillation.** The DA distills our in-depth research on how LCA is modeled and interpreted, and the key elements and structure necessary to create an accurate simplified LCA methodology. This also acts as a constraint for LLMs that bounds outputs within methodologically valid LCA representations.

- (2) **Validation by Construction.** Borrowing from principles in programming languages, our DA enforces validation similar to the idea of being "correct-by-construction": if an input cannot be expressed in the DA, the system cannot proceed with LCA. This enables automatic validation of missing or irrelevant components before modeling.
- (3) **Interpretability.** Since the DA explicitly encodes the LCA model structure, including relationships between elements, processes, and their interdependencies, it supports interpretability as the LCA flowchart can be reconstructed from the DA.
- (4) **Editability.** The DA separates the model structure from the emission factor values, which allows users to dynamically substitute or personalize data sources (e.g., switching between ecoinvent versions, open sources, or proprietary databases) without reconstructing the model. This modularity supports both user editability and industry needs for data security and provenance control.

As shown in [Table 1](#), the Simplified LCA representation is composed of a sequence of life cycle stages, including production, transport, and use. Each life cycle stage contains stage-specific properties and structured nodes with both quantitative parameters and metadata for source tracking, uncertainty annotation, and unit normalization.

The production stage is represented as a list of pairs of *raw materials* and *manufacturing processes*. The *raw materials* branch supports two specification types: *independent material* for directly specified quantities (e.g., 2 kg of plastics) and *related materials* for materials defined through ratios or relationships (e.g., 1:2 mixture of coffee and milk). Each material node captures quantity specifications (amount/ratio value and unit) and associated emission factor (e.g., carbon intensity).

The node structures of process-oriented phases, such as manufacturing, are designed to meet their specific requirements. *Manufacturing* is represented as a transformation *process* applied to raw materials:

$$\text{apply}(\text{process}, \text{raw_materials}^*) \quad (1)$$

Each process is defined as:

$$\text{process} ::= f(\text{raw_materials}) \rightarrow (\text{energy}, ef, \text{output}^*, \text{waste}^*, \text{location}) \quad (2)$$

where the process transforms the raw materials into outputs (the products generated by the process) and waste (the waste produced by the process). The process consumes energy and takes into account the geographic location where it occurs.

Transport nodes capture logistics parameters such as locations, weights, and transport mode. *Use* phase nodes capture energy consumption through power, time, and location. The DA includes source tracking for uncertainty quantification, with each parameter annotated as either *present* (explicitly stated in the input text) or *inferred* (derived from context or defaults).

Our DA is fully representable in JSON format (see [§B](#) for examples), facilitating both processing and serialization in the frontend while maintaining human readability.

4.2 LCA Stage Classification

The first step in our process is identifying text relevant to LCA calculations and identifying the relevant life cycle stage (transport, use, etc.). Unlike previous work in augmenting educational content (e.g., mathematics [\[19\]](#) or physics [\[36\]](#)) where users can select appropriate simulation types based on basic domain knowledge, LCA life cycle stages are often unfamiliar to the general public. For instance, both manufacturing and use phases may involve energy consumption.

To address this challenge, we implemented a semantic similarity-based classification algorithm using Sentence-BERT [\[66\]](#), which is more computationally efficient than LLMs for this task. Our approach compares the embedding of input text to a small set of pre-classified reference examples, the closest match determines the stage classification.

Based on our evaluation (see §6.1.1 for details), we selected MiniLM [84] with 22M parameters for deployment, as it achieved comparable accuracy to larger models in our case while offering significantly faster inference times.

4.3 Parameter Reasoning and Environmental Impact Analysis

Following stage classification, Living Sustainability employs a multi-agent pipeline implemented using GPT-4o-mini (gpt-4o-mini-2024-07-18) as the LLM backbone to extract parameters and retrieve emission factor data.

Parameter Extraction Agent. This agent transforms unstructured text into a structured DA suitable for LCA modeling. Employing chain-of-thought, extra task-related domain information, and few-shot prompting strategies, the agent follows a systematic process to: 1) identify key LCA parameters such as material quantities, manufacturing processes, and energy consumption, 2) classify parameter values as either explicitly stated in the text or reasonably assumed based on domain knowledge (e.g., when material quantities are not specified, the agent assumes a default value of 1 gram to enable baseline calculations as values can be modified by the user), and 3) handle complex parameter relationships (e.g., when materials are described in terms of fixed ratios or proportions).

When encountering detailed chemical descriptions such as stoichiometric relationships of raw materials, the agent can first identify the molecular weight of each component through its knowledge base or online retrieval. It then applies the specified stoichiometric ratios to convert all measurements to a standardized mass-based representation.

Data Retrieval Agent. This agent determines emission factors (e.g., kg CO₂e per kg for material or per kWh for electricity) for each identified parameter through a multi-source retrieval system. To optimize computational efficiency and reduce runtime, we designed the agent to interface directly with APIs rather than a vision-based web-browsing approach. The agent leverages multiple data sources: it queries research papers and public reports through Google Custom Search API, accesses authenticated LCA databases such as Climatiq [20] for transportation emissions, European Commission LCIA [21] and Idemat [43] for raw materials, and utilizes Electricity Map [57] for location-specific grid carbon intensities.

To reduce the token size, we equipped the agent with a custom preprocessing tool that employs regular expressions and semantic cues [86] to filter retrieved content. This tool removes irrelevant elements, such as webpage headers, footnotes, and HTML syntax, to reduce the processed string length. The agent ensures traceability by storing data source URLs within the DA representation.

4.4 Simplified LCA Methodology

After the agents have transformed the input into the DA structure and retrieved emission factors, the EI is computed by aggregating the emissions E across all LCA stages. The EI of an everyday object or activity can be, in a simplified form, expressed as:

$$E_{object/activity} = E_{RM} + E_M + E_T + E_U \quad (3)$$

where E_{RM} represents emissions from raw materials, E_M is manufacturing, E_T is transportation, and E_U is use phase. Each stage is the sum of multiple finer-grained contributing components. For instance, raw materials E_{RM} can be expressed as:

$$E_{RM} = \sum_{i=1}^N E_{material,i} \quad (4)$$

where the emissions of each material i are calculated as $E_{material,i} = \frac{m_i \cdot ef_i}{L_i}$, here m_i is the mass of material i , ef_i is its corresponding emission factor (typically expressed in kg CO₂e per unit mass), and L_i is a loss factor accounting for material yield rate or process waste. Similarly, the total manufacturing is the sum of all individual

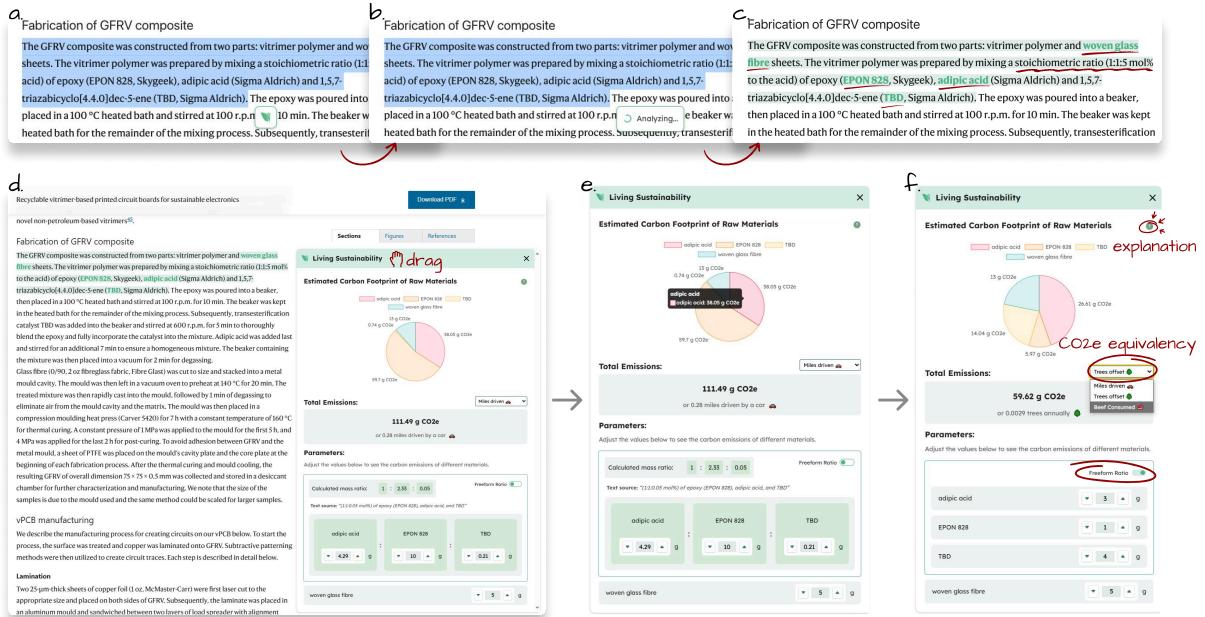


Fig. 5. User flow walkthrough. User highlights text while reading, system highlights relevant parameters and overlays interactive visualization for exploring environmental impact.

processes $E_M = \sum_{j=1}^N E_{process,j}$, transportation $E_T = \sum_{k=1}^N E_{seg,k}$ as the sum of each segment k of the supply chain journey from source to destination.

We evaluate the accuracy of this simplified LCA methodology in §6.1.3 through case studies compared to traditional full LCA.

5 SYSTEM DESIGN

After presenting our backend autonomous pipeline for EI assessment, we now demonstrate how Living Sustainability makes this capability accessible in everyday contexts. Living Sustainability is currently supported as a web browser extension compatible with Chromium-based browsers, as we aim to integrate directly into users' natural browsing and reading experience through simple text selection and automatic webpage parsing. In this section, we walk through the user interaction flow and showcase the common daily scenarios powered by Living Sustainability without requiring technical expertise or specialized tools.

5.1 User Flow Walkthrough

Users interact with Living Sustainability through a simple two-step process: natural language input or automatic webpage parsing, followed by interactive exploration of the results. We demonstrate Living Sustainability's interaction flow through an example of analyzing material synthesis from a sustainable electronics paper [88].

Step 1: Input and Initiate LCA. The user encounters a technical paragraph describing the synthesis of a composite material. Despite the complexity of chemical stoichiometry and material science terminology, the user simply highlights this text to initiate analysis (Figure 5a).

Living Sustainability processes this input through its backend pipeline, typically within 5 seconds, automatically identifying raw materials and their relationships, converting molar ratios to mass-based quantities, and retrieving



Fig. 6. Augmented shopping for an electronic product. System parses product specifications from webpage, estimates the EI, and recommends alternatives with similar features and price points.

emission factor for each component (Figure 5b). The system also highlights relevant parameters in the original text and transforms numerical values into interactive input boxes. This text augmentation works with any webpage content that can be modified through the browser’s DOM (Document Object Model), but not with documents like PDF that are rendered as fixed-layout content in browsers (Figure 5c).

Step 2: Interact and Explore. Living Sustainability overlays an interactive visualization window on the webpage, showing the EI breakdown of each component while maintaining the original browsing context (Figure 5d). By default, the visualization preserves stoichiometric relationships, allowing users to adjust quantities while maintaining the molar ratios from the original text (Figure 5e). As shown in Figure 5f, users can toggle this constraint off to freely enter custom quantities for individual materials. As quantities change, the visualization updates in real-time to show both absolute EI values and relatable equivalency metrics (e.g., miles driven) that help contextualize CO₂e values. Through an information panel, users can access detailed methodology explanations, including LCA functional units and data source attributions.

5.2 Supported Augmentation Scenarios

We further present Living Sustainability through daily scenarios that naturally combine multiple communication strategies highlighted in our formative survey, rather than organizing by individual features. We aim to demonstrate both how features work together in practice and the system’s broad applicability across different daily use cases.

5.2.1 Shopping. When users browse major electronics retailers (e.g., Amazon, Apple, Google, Samsung), Living Sustainability automatically processes product pages. For example, Figure 6 illustrates when visiting an iPhone 16 product page, upon page load, the system 1) parses the product name and basic specifications, and 2) retrieves or estimates the EI. We prioritize data from official product environmental reports, when available, as these are legally verified and endorsed by companies. The system also recommends alternative models with similar features and price points but potentially lower EIs. We hope this tool will seamlessly integrate EI assessment into the everyday shopping experience.

5.2.2 Package Shipping. We demonstrate how EI information can be integrated into existing shipping workflows. For example, on FedEx’s webpage (Figure 7), one of the world’s largest freight carriers, when users enter shipping details and click “show rates”, Living Sustainability automatically calculates the EI of each available shipping method. To provide deeper insight into shipping modes, we augment the interface with an interactive map

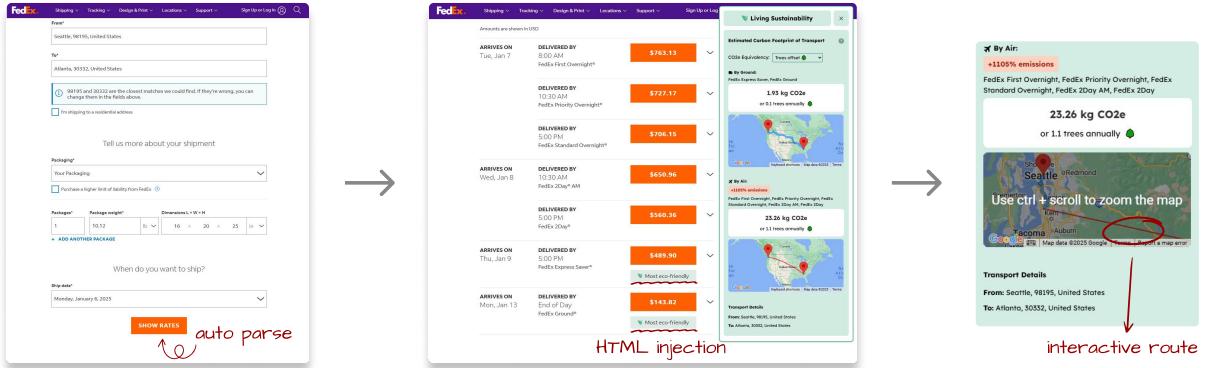


Fig. 7. Augmented transportation for a package shipping service FedEx. System integrates with existing shipping workflow, parses package shipping details, estimates EI for different shipping options, and embeds visual indicators through webpage DOM.

visualization powered by Google Maps and FedEx. These maps show the routes of available transportation modes based on shipping parameters—for instance, automatically excluding air shipping options for nearby destinations. We also developed a special compatibility for FedEx: the tool embeds “*most eco-friendly*” icons next to the shipping options with the least EI through the DOM.

5.2.3 Cloud Computing. For cloud service platforms like Microsoft Azure (Figure 8), Living Sustainability helps users understand the EI of their infrastructure choices. When users select an instance configuration, the system automatically parses the information on the webpage, then prompts users to specify their expected instance usage rate to obtain workload patterns. Using this information, the system calculates emissions by combining instance-specific power consumption data with regional power grid carbon intensities, leveraging data from Microsoft Sustainability Calculator [23], Cloud Carbon Footprint [44], and Climatiq API [20].

5.2.4 General Browsing and Reading. We demonstrate Living Sustainability’s generalizability by supporting EI analysis from any natural language input, including both plain text and voice. The system also handles incomplete and partial information by actively guiding users to provide necessary details for LCA.

For example, given the input “*I want to ship a package to Berkeley*”, the system identifies missing key parameters required for transportation modeling (origin location and package weight). Rather than attempting estimations

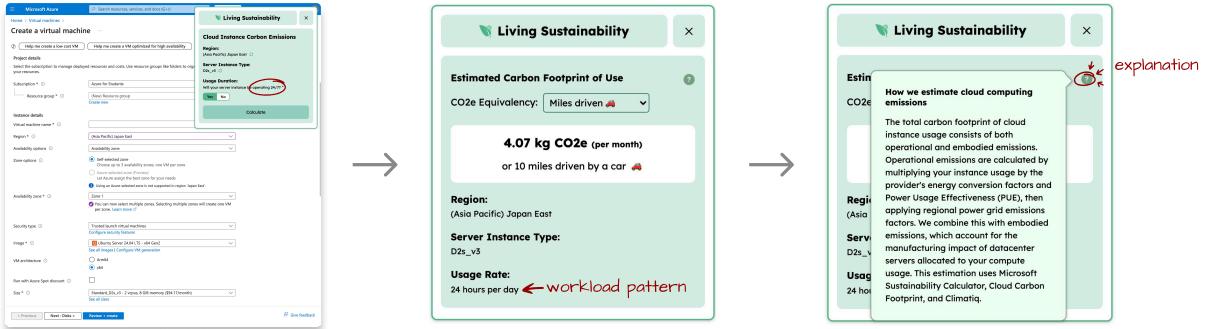


Fig. 8. Augmented use for a cloud service. System parses server information on Azure webpage, prompts users for workload patterns, and estimates EI.

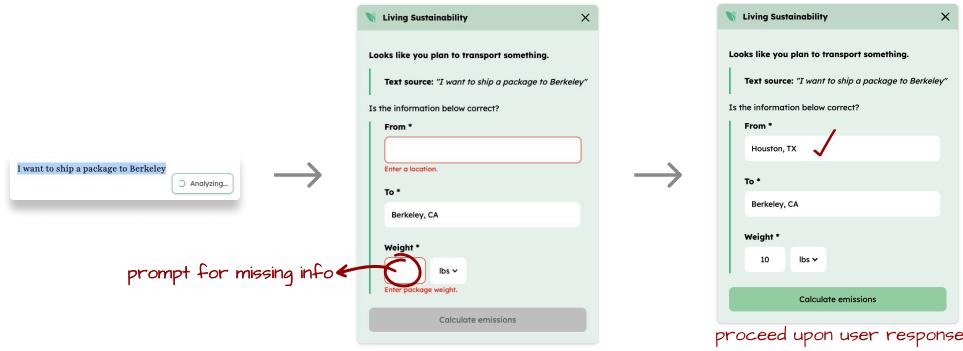


Fig. 9. General text input with incomplete information handling. System identifies missing key parameters, pauses and requests additional input, and proceeds with the calculation upon user response.

with hallucination, it pauses and explicitly requests the missing information. Once the user provides additional details, the system proceeds with the calculation (Figure 9).

The system shows different behaviors based on the type of missing information. Figure 10 demonstrates when processing text input like “*I use an Nvidia RTX 4090 to train a transformer model, the training takes 24h*”, the system detects missing location data. However, unlike transportation scenarios that require complete location information, the system can estimate an EI range using the global power grid emission factor. Alternatively, users may specify a location to reduce uncertainty. Additionally, although the user input in this example does not specify the power consumption of the device, the system autonomously retrieves this deterministic value through an online search as the information is publicly available.

To preserve context during parameter adjustments, the system implements dynamic bi-directional text binding: numerical values in the original text are transformed into interactive input fields by injecting HTML elements into the webpage, enabling users to modify parameters directly within their source context.

5.3 Implementation

Our system consists of a Chromium extension frontend developed in HTML, CSS, and JavaScript, complemented by a Node.js backend server deployed on a virtual machine (Figure 11). We implemented this architecture because Chromium blocks local backend operations and cannot securely store such credentials on the client side. The backend manages secure API communications and coordinates our autonomous EI modeling pipeline.

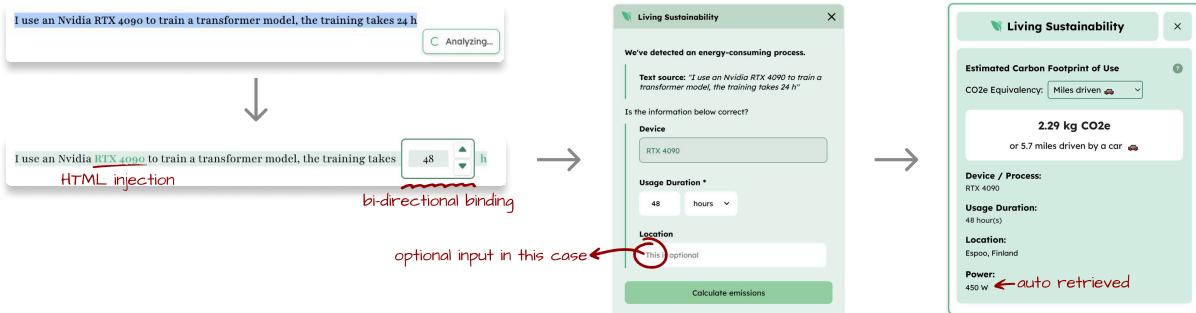


Fig. 10. General text input with partial information. System detects missing location data, calculates EI ranges using the average power grid emission factor if the location is not specified.

5.3.1 Data Management. We implemented a cloud-based database solution using MongoDB Atlas to store carbon emissions data. The database grows dynamically as users interact with the tool, storing previously estimated product carbon footprints and retrieved emission factor data. For instance, when a user queries a phone's EI, the system first checks the database for existing calculations. If no data exists, the system performs a new retrieval or estimation, and then stores the results. This approach reduces runtime and computational overhead by eliminating redundant calculations for frequent queries.

The database utilizes BSON format for its integration with the JSON-based data exchange of our extension. The database schema accommodates key details about each product, including its carbon footprints, specifications, data sources, and LCA methodologies. To facilitate comparative analysis, the database links each product to its competitors, allowing users to evaluate alternatives based on EI. Each entry also includes a method field that distinguishes between directly retrieved from online sources or estimated through our pipeline.

5.3.2 Dynamic Web Integration. Our extension implements HTML parsing and injection mechanisms to analyze either entire webpages or user-selected text. This functionality is enabled by a JavaScript content script that integrates with active webpages to process content and respond to user interactions.

HTML Injection. The extension implements a text highlighting system that monitors user selections through event listeners. When text is highlighted, the system injects a floating action button adjacent to the selection. Upon triggering, the system processes the highlighted text by 1) parsing it into a string, 2) structuring into a JSON

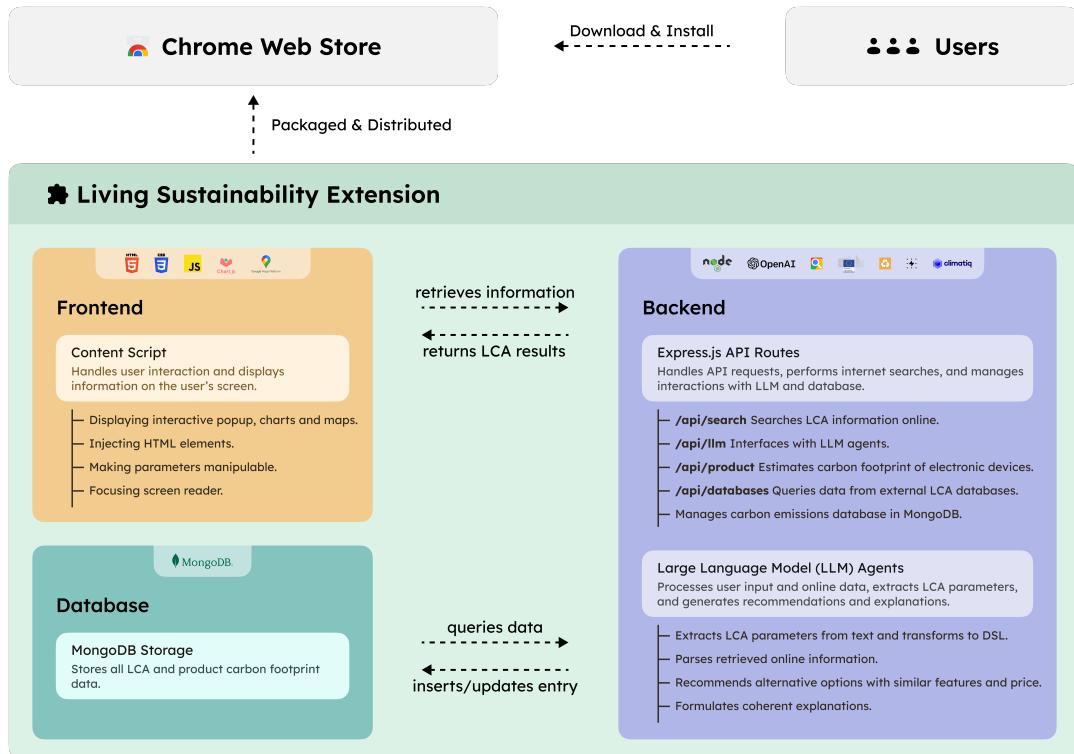


Fig. 11. Overview of Living Sustainability implementation architecture. Our system consists of a Chromium extension frontend that is packaged and can be installed directly by users, with a backend server that connects to a MongoDB Atlas database to store carbon emissions data to reduce runtime and computational overhead.

object, and 3) sending it to the backend. Following backend processing, the system dynamically injects HTML components to display carbon footprint data related to identified parameters. These components are encapsulated within a Shadow DOM to maintain isolation from the webpage's existing styles and scripts, preventing potential conflicts. The system also enhances the visual presentation of input text. This involves referencing the original HTML node of the highlighted text and replacing the node with a new HTML element with visual effects while copying the original styles (see example in [Figure 5c](#)).

Interactive Input. The system transforms static numeric values on a webpage into interactive input fields (see example in [Figure 10](#)). This transformation replaces original HTML nodes with custom HTML elements while maintaining visual consistency. An input listener tracks value changes and triggers corresponding updates in the extension. This real-time coupling between input fields and visualizations allows users to explore “what-if” scenarios and immediately observe their EI.

Visual Decision Support. For shopping and shipping platforms, the system implements EI tagging. This process involves systematically traversing the webpage DOM to identify options, calculating the EI of each option, and dynamically injecting green indicators. For example, for FedEx, the tool analyzes shipping alternatives and visually highlights the options by injecting “most eco-friendly” icons (see example in [Figure 7](#)). These visual cues help users without having to compare carbon emissions across multiple options manually.

Screen Reader Support. To ensure our tool serves Blind or Low Vision (BLV) users, we implemented screen reader integration through DOM manipulation. We use JavaScript’s *focus* and *tabindex* functions to prioritize the extension in screen readers. We make all the elements in the extension focusable and in the right order. When the system injects new elements or updates existing content, We again use *focus* and *ARIA live regions* to direct screen readers to these changes.

5.3.3 Parameter Visualization. Our system implements a suite of interactive visualization components to communicate EI-related information. For raw material analysis, we utilize Chart.js to present hierarchical breakdowns of material compositions and their associated carbon emissions, with visualizations updating dynamically as users modify parameters through the interactive input fields.

For transportation, we integrate Google Maps API to create interactive geo visualizations of routes. To circumvent Chrome’s Content Security Policy restrictions on loading external resources and dynamically evaluated scripts like Google Maps API, we host an external HTML webpage on our server that connects with our backend. The extension calls an endpoint on our backend to trigger the hosted HTML webpage to generate a geo map. Finally, the extension uses *<iframe>* property to fetch the element from the external HTML page and display them on the extension UI.

6 EVALUATION

To evaluate Living Sustainability, we conducted 1) a quantitative technical analysis of the system’s accuracy in transforming unstructured text to LCA data abstraction and case studies comparing our results against traditional full LCA, and 2) a qualitative user study to assess the system’s usability and effectiveness in real-world scenarios. Below, we detail the methodology and findings of our two-pronged evaluation approach.

6.1 Technical Evaluation

Our system automates the three key steps an LCA expert would use to perform this analysis. The first is identifying relevant information about the product or process and sorting them into life cycle stages. The next step is extracting the relevant information such as the exact materials used, their amounts and units, and other processes into our DA to create a complete representation of the life cycle inventory needed to compute EI. The

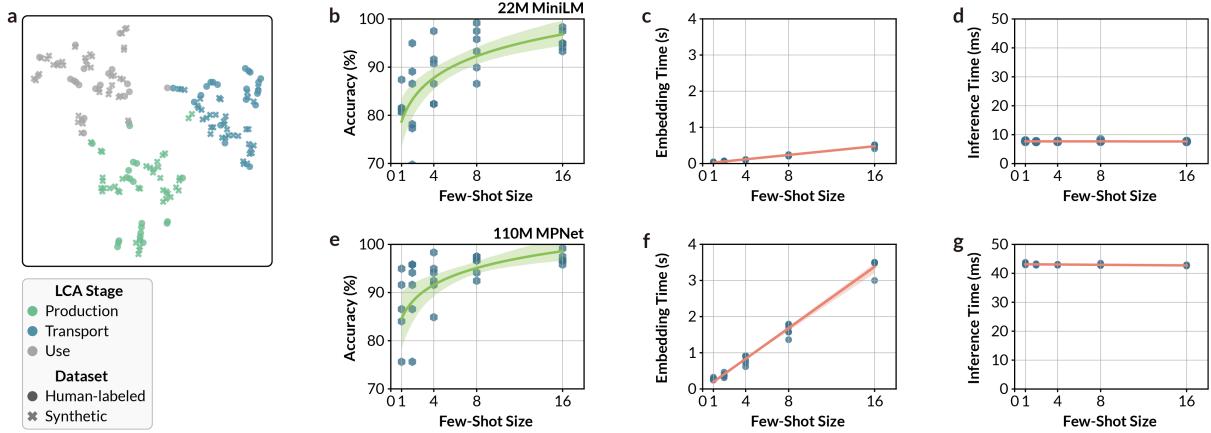


Fig. 12. (a) t-SNE plot of sentence examples in MiniLM embeddings, showing the separation of clusters corresponding to each LCA stage. The “Production” stage includes both raw material extraction and manufacturing. (b-g) Accuracy and average runtime of 22M MiniLM and 110M MPNet as the size of few-shot examples per LCA stage increases ($N = 6$, shaded region = $\pm \sigma$).

final step is matching elements of the DA to emission factors and using those to compute EI. We evaluate our system with benchmarks for the first two steps followed by detailed end-to-end case studies to evaluate accuracy in real-world scenarios.

6.1.1 Semantic LCA Stage Classification. A critical component of our system is the ability to accurately identify text relevant to LCA modelings and classify it into appropriate life cycle stages. We evaluated several pre-trained Sentence-BERT model architectures and sizes: 6-layer and 12-layer MiniLM [84] with 22M and 33M parameters, respectively, and an MPNet [71] with 110M parameters.

We created two distinct datasets. For testing, we manually collected 20 example sentences for each LCA stage from various online sources, ensuring a diverse representation of real-world textual expressions (see examples in Appendix Table 6). For training, we used LLM (gpt-4o-mini-2024-07-18) to generate synthetic data, including both descriptive explanations of each LCA stage and 20 reference examples per stage. This follows established practice of leveraging LLM distillation to train smaller models [41]. We used these synthetic examples for few-shot learning [14], while reserving the manually collected dataset for testing to better evaluate the generalizability.

Our evaluation focused on three key metrics: classification accuracy, embedding computation time, and inference latency. To understand how model performance scales with few-shot sizes, we randomly sampled from our synthetic dataset to create datasets of 1, 2, 4, 8, and 16 examples per LCA stage. As shown in Figure 12b,e, classification accuracy improves with more few-shot examples, reaching convergence at approximately 16 examples per stage. The embedding computation time scales linearly $O(n)$ with dataset size, remaining under 0.6 s for MiniLM 22M with 16 examples per stage, compared to ~4 s for MPNet (Figure 12c,f). We note that these embeddings of reference examples can be precomputed and loaded only once during system initialization, so this computation doesn't affect user experience or response latency. The inference time remains constant regardless of size (Figure 12d,g) because Sentence-BERT compares the embedding of the test example to the precomputed embeddings of each LCA stage category, so it depends only on the fixed number of LCA stages. The final evaluation results, using 20 random synthetic reference examples per stage and the manually collected test set (20 examples per stage), are shown in Table 2. The results show that our system can extract this critical information needed for LCA with high accuracies of over 97% with sub-second latency.

Model	Description Only		Few-shot Only		Description + Few-shot	
	Acc.	Emb. Time (s)	Acc.	Emb. Time (s)	Acc.	Emb. Time (s)
MiniLM 22M	83.19	0.051	97.48	0.643	97.48	0.704
MiniLM 33M	88.24	0.099	95.80	1.273	95.80	1.417
MPNet 110M	90.76	0.325	95.80	4.383	97.48	4.849

Table 2. Performance of MiniLM and MPNet sentence transformers under three evaluation setups: LCA stage description only, 20-shot reference examples of each stage only, and hybrid. Metrics include model size, accuracy (Acc.), and embedding time (Emb. Time).

6.1.2 End-to-End Data Abstraction Transformation. The primary contribution of Living Sustainability lies in bridging the gap between unstructured natural language descriptions and LCA modeling for users without domain expertise. We evaluated the accuracy of this transformation using our end-to-end automated pipeline (§4).

To create ground truth, we constructed a dataset using the 60 natural language examples (20 examples per stage) from our initial manual data collection in §6.1.1. Two researchers with LCA experiences independently labeled the dataset with the correct data abstraction representations required for LCA modeling, followed by cross-validation. We also developed an automated scoring program that iteratively checks each entry, using a rubric where any error in material extraction, process identification, or any missing entry would result in the entire example being marked as incorrect.

Our system achieved 98.3% accuracy in end-to-end transformation. The single observed error occurred in a transportation example (Example 1 in Appendix Table 6 Transport), where the system incorrectly applied a freight transport emission factor instead of the appropriate gasoline car driving. We also note that even in the few cases where the Sentence-BERT misclassified the LCA stage, the subsequent agent could still construct the correct DA by leveraging contextual parameters that were incompatible with the misclassified stage, this echoes with our DA principle of "validation by construction" (discussed earlier in §4.1).

6.1.3 Case Study. We conduct case studies to validate Living Sustainability’s automated analysis against full LCA for representative examples described in §5.2. To benchmark, we compared the EI results from Living Sustainability against a manual full LCA conducted using openLCA² with ecoinvent v3.10 database [26] and CML v4.8 LCIA method.

Package Shipping (from §5.2.2). The functional unit is defined as “the transportation of a 10.12 lb package from Seattle, WA (98195) to Atlanta, GA (30112)”. This scenario allowed us to validate Living Sustainability for a combined transport mode profile. In the full LCA, ground transport is modeled using the “transport, freight, lorry, >16 metric ton, diesel” LCI, and air is modeled using “transport, freight, aircraft, medium haul”. The shipment involved 2635 miles of ground travel; or 2182 miles of air travel, with an additional 29 miles of ground handling at origin and destination. For air transport, our tool uses IATA’s RP1678 methodology [46].

As shown in Table 3, both transportation modes showed only 10% discrepancy between the full LCA and Living Sustainability. The manual LCA required approximately 2 expert-hours, primarily to manually identify vehicle specifications (e.g., truck size, fuel type) and air freight profiles (e.g., belly freight vs. dedicated freighter, haul distance category) of FedEx package shipping, as well as to source accurate transport routes and distances.

Cloud Computing (from §5.2.3). The functional unit is defined as “the use of an Azure D2s v3 virtual machine in the Japan East region for 30 days”. In the full LCA, total emissions are decomposed into operational emissions

²www.openlca.org

By Ground	Vehicle	Airplane	Total (kg CO2e)	By Air	Vehicle	Airplane	Total
Full LCA	2.16	N/A	2.16	Full LCA	0.025	25.94	25.965
Ours	1.93	N/A	1.93	Ours	0.03	23.23	23.26
Diff (%)	-10.65%	N/A	-10.65%	Diff (%)	20%	-10.45%	-10.42%
Baselines							
o4-mini	1.30–2.41	N/A	1.30–2.41	o4-mini	0.99–1.95	8.86–19.50	9.85–21.45
GPT-4.1	1.24–3.00	N/A	1.24–3.00	GPT-4.1	0.11–1.00	9.60–9.70	9.71–10.70

Table 3. Comparison of package shipping impacts in CO2e between Living Sustainability and traditional full LCA for ground and air transportation. Baseline values from leading LLMs, including OpenAI o4-mini (o4-mini-2025-04-16) and GPT-4.1 (gpt-4.1-2025-04-14), are reported as min–max ranges over independent runs (N=5).

(electricity consumption during runtime) and embodied emissions (hardware manufacturing impacts amortized over expected equipment lifetime). Azure Japan East is located in the Saitama Prefecture, Tokyo. Operational emissions are estimated using a weighted electricity carbon intensity (around 70% conventional energy and 30% renewable energy), based on Tokyo Metropolitan’s energy mix [27, 45, 78]. Embodied emissions are modeled from a component-level LCA of Azure hardware, adjusting the solid-state drive (SSD) emissions based on ACT data [37], and amortized over a 5-year expected lifespan. The detailed methodology for modeling embodied emissions for cloud is provided in Appendix D.

As shown in Table 4, the final results show only a 6.86% difference. The manual LCA process required >3 expert-hours, given the manual effort needed to investigate and align cloud instance specifications to physical hardware LCAs, and to determine accurate local carbon intensities from government energy reports.

Upon further investigation, discrepancies observed in these case studies are largely attributed to differences in emission factor sources. For example, emission factor databases such as Ecoinvent often have multiple options even for a specific process (truck shipping) as well as a variety of customizable inputs, and updates between database versions. Correct or standardized emission factor mapping is a common challenge in LCA and an area of active research [7, 89].

We note that the manual full LCAs required approximately 3 expert-hours to complete, which is considered relatively quick. This is due both to the simplicity of the inventory, and because we did not have to consult external sources to find proprietary data. This process can take months for products such as consumer electronics devices [28, 89], whereas our tool can deliver comparable results in real-time.

	Operational	Embodied	Total (kg CO2e)
Full LCA	2.97	1.40	4.37
Ours	3.10	0.97	4.07
Diff (%)	4.38%	-30.71%	-6.86%
Baselines			
o4-mini	2.2–30	0.3–1	2.5–30.9
GPT-4.1	2.5–27	0.8–2	3.3–29

Table 4. Comparison of cloud computing impacts in CO2e between Living Sustainability and traditional full LCA.

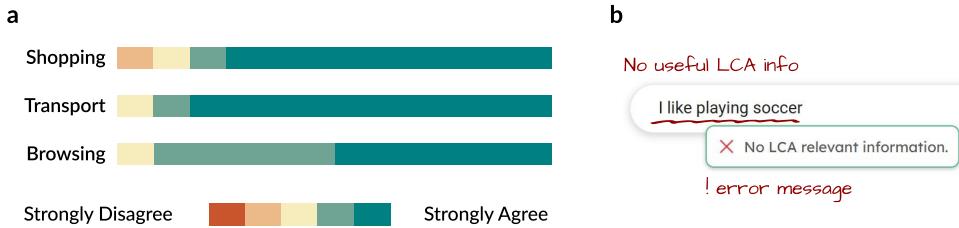


Fig. 13. (a) Participant feedback on the different augmented scenarios of Living Sustainability for usefulness. (b) System failure case when processing input with no LCA-relevant information.

Baselines. To further contextualize our results, we evaluated the state-of-the-art LLMs, the latest GPT (GPT-4.1) and reasoning (o4-mini) models, as baselines. They are provided with identical system prompts and inputs, and are equipped with web search tools. However, unlike Living Sustainability, these LLMs lack a DA layer and infer EI directly from the natural language input. We found that both LLMs failed to produce robust or reproducible EI estimates. Across five independent runs, their outputs varied one order of magnitude, with discrepancies from the full LCA often exceeding 100%. These results highlight the limitations of relying solely on black-box LLM to hallucinate EI, and underscore the importance of integrating domain-specific DA layers for better constraint.

6.2 User Evaluation

We evaluated our tool with users in the loop through a preliminary user study with 12 participants, followed by open-ended discussions. The study aims to measure the system’s usability and gather preferences and feedback from potential users across multiple use cases, providing insights into how the general public might benefit from such automated LCA tools. The study method is detailed in Appendix E.1.

6.2.1 Usability Results. This section presents the outcomes of our user evaluation. Our system achieved an overall System Usability Scale (SUS) score of 82.71 (SD = 11.92). When rating different scenarios on a 5-point Likert scale, participants particularly found intuitive and useful (Figure 13a) of the shopping scenario (Avg = 4.50 and 4.50), shipping and other transportation-related (Avg = 4.58 and 4.75), and text highlighting during general browsing (Avg = 4.58 and 4.42). Overall, participants appreciate all of the features and scenarios.

We found that only 75% of our participants ever considered environmental impact in their daily lives. This finding is noteworthy given that our participant pool consisted of individuals who voluntarily signed up for a sustainability-related study, suggesting that adoption rates among the general public might be even lower. This insight highlights a crucial challenge: technical solutions alone may not be sufficient to drive sustainability behavioral change.

The participant who provided the lowest SUS score offered valuable insights into the tool’s contextual utility, explained, “I don’t see myself using it when I buy a laptop or a phone, because those purchases are too big and have too much impact on my daily life for me to make decisions based on sustainability.” However, the same participant emphasized the tool’s value for transportation decisions, noting that “shipping and travel analysis is a useful example because it may affect my decision, such as store pickup versus home delivery.”³

One participant proposed integrating a chatbot interface, suggesting users could “click on the floating icon and open a chatbot that can enable searching for anything’s carbon footprint.” Others expressed interest in more interactive information panels, also with chatbot integration, as the current static text format limits QA engagement with the underlying methodology and data sources.

³Participant quotes have been slightly modified for concision, grammar, and anonymity.

	Raw Materials	Manufacturing	Transport	Use	Total (g CO2e)
Full LCA	22.71	28.62	571.78	21.59	644.7
Ours	19.99	28.68	582.51	24.00	655.18
Diff (%)	-11.98%	0.21%	1.88%	11.16%	1.63%
Baselines					
o4-mini	137–3095	680–1040	3–89	100–118	1090–3879
GPT-4.1	103–320	24–266	2.4–7	30–50	225.4–591

Table 5. Comparison of the environmental impact in CO2e of recyclable PVA PCB substrates [18] between Living Sustainability and traditional full LCA across four LCA stages: raw materials, manufacturing, transport, and use.

The presentation of CO2e equivalencies emerged as a complex design challenge. While one LCA startup founder cited it as their favorite feature, another participant found the three CO2e equivalencies we provided insufficient for building intuition. The participant suggested including “a quote of the estimation of people’s daily CO2e” as a reference point, though this raises important questions about the feasibility of such personal benchmarks due to diverse lifestyles and habits.

6.2.2 Other Use Cases and System Limitations. Throughout the open exploration, participants discovered several new applications that demonstrated the tool’s versatility.

In one compelling case, one participant utilized the tool to retrospectively assess the EI of their sustainable materials research paper, for which they had not conducted an LCA.

We also identified several technical limitations through participant interactions. When one participant attempted to analyze a full academic paper of >7,000 words, the response time of the agent system exceeded 1 min. In addition, when users input no LCA-relevant information, the system returned NULL and an error message (Figure 13b).

6.2.3 More Case Study. The use case of conducting LCA for research papers discovered during the user study’s open exploration provided another opportunity for an in-depth case study validation.

The functional unit is “a 37g recyclable PCB substrate with PVA and glycerin” [18]. To ensure comparability, we used identical system boundaries (encompassing raw materials, manufacturing, transportation, and use phases). A notable challenge in the manual LCA was the absence of PVA in the ecoinvent v3.10 database. Following standard LCA practice, the LCA practitioner selected polyvinyl chloride (PVC) as the closest approximation based on expert domain knowledge. This substitution exemplifies a common challenge in traditional LCA: practitioners often need to justify appropriate proxies when exact matches are unavailable in standard databases [89], which contributes significantly to the time-consuming nature of traditional LCA, and also highlights why it remains largely inaccessible to non-experts. For the use phase, we standardized energy consumption to 1 kWh, using “electricity, low voltage, renewable energy” in openLCA and selecting Seattle (which uses 100% hydroelectricity) as the location in our tool.

Living Sustainability successfully computed the total CF in CO2e for every LCA stage, and demonstrated remarkable accuracy when compared to the full LCA. As shown in Table 5, the overall discrepancy is only 1.63% in total carbon footprint. This small discrepancy is attributed to two key factors upon investigation: the different emission factors used for materials (PVC in the full LCA as aforementioned versus PVA in ours), and in the electricity modeling, where the full LCA used average global grid mix data while ours can be precise to a specific region (e.g., Seattle).

Our tool again computed comparable results in real-time in a single “highlight”.

7 CONCLUSION AND DISCUSSION

In this paper, we introduced Living Sustainability, a novel AI agent-powered system to communicate sustainability information by automatically augmenting webpages and text documents with interactive, in-context visualizations. Our system automatically transforms unstructured natural language into LCA data abstractions using a set of NLP techniques, then conducts a simplified LCA to create real-time EI insights. Through our technical evaluations and user study, we demonstrated that our system can generate accurate DA representations for LCA and quantitative EI results comparable to traditional full LCA, while reducing analysis time from hours of expert time to one click. Our work seeks to bridge a critical gap between the public interest in sustainability and accessibility of EI measurements. Just like measuring power or water consumption, our goal is to enable users with zero LCA expertise to gain meaningful insights into how their everyday choices relate to abstract environmental metrics like carbon footprint. To the best of our knowledge, Living Sustainability is the first system capable of conducting LCA from unstructured language input and creating real-time EI communication with reactive runtime and explorable explanations.

Below, we situate our results within related work and discuss multiple key avenues for future work.

7.1 Advancing Complex LCA

While our current system effectively supports various LCA stages and shows promising accuracy for common, well-understood scenarios, extending its capabilities to handle more complex or ambiguous cases presents opportunities for future work.

Prior SHCI research highlights how synthetic data in eco-feedback prototyping leads to misleading insights [13]. Therefore, we designed our current system through "transparent accounting and accountability" [12] in our data abstraction and deliberately limit the scope to scenarios where high fidelity and confidence can be achieved based solely on the information explicitly provided by the user, rather than providing estimates with high uncertainty.

For instance, "Ultimaker Cura estimates that printing this shape will use 34g of PLA filament" (Example 1 in Appendix [Table 6 Production](#)), our system computes EI for the raw material (PLA) only. While 3D printing undeniably consumes electricity, we do not estimate energy usage due to the lack of specifics (e.g., printer model, print duration).

Similarly, inputs like "I like playing soccer", the system currently returns an error message due to the absence of actionable LCA-relevant content. While a possible future direction could involve multi-step agent reasoning (e.g., inferring likely intent, such as estimating emissions from a soccer ball or field use), our current version limits such inferences to avoid overstepping user intent and to maintain semantic clarity.

The system also has limitations when handling 1) long technical datasheets; 2) intricate nested fabrication processes that require multi-step interpretation and sequencing; 3) niche electronic products with proprietary components for which no public teardown data or product reports exist. These limitations were also observed during open exploration in the user study (as discussed in [§6.2](#)). Given that our system backend is built on a modular agent framework, we anticipate that many current limitations, particularly regarding context length and multi-step reasoning, will be progressively addressed as the base models are upgraded.

To address these limitations, we identify two promising areas for further advancement.

Graph-based Data Representations. Our current hierarchical data abstraction, while efficient for frontend parameter manipulation, limits the system's ability to represent complex LCA, such as nested processes. Redesigning a node graph data representation would allow for recursive flow where materials and processes have interdependencies, and multi-path simulations for alternative production routes.

Automated LCA Verification. A natural next step is the automated analysis of full LCA reports published by companies. By parsing the document, we could extract the assumptions, methodologies, and impact factors

used for the LCA report. This would allow for systematic comparison of corporate sustainability claims and support regulatory compliance verification. Our system could also integrate domain-specific heuristics to provide intervention [53] when identifying potential inconsistencies between materials and processes, or deviations from specific industry standards.

7.2 Deploying at Scale: From Individual Behavior to Collective Insights

Existing research suggests that active interventions can drive behavioral change in various domains, such as personal health [2] and digital well-being [62, 87]. While the study between human behavior and active sustainability feedback is limited, we believe that such active personalized feedback can be extended to promote similar positive changes in sustainability behavior.

Uncertainty Transparency and Trust-building. LCA inherently involves estimations and uncertainties. Currently, our system provides explanations of data sources and assumptions made, but future work could expose modeling pathways and highlight where uncertainty arises in EI estimates to ensure transparency, which will improve trust and persuasiveness. We advocate that both the general public and policymakers should learn to make informed decisions about sustainability under uncertainty rather than waiting for perfect information [68]. Future work could also allow users to manually intervene in real time to adjust parameters or refine assumptions to reduce uncertainty in their specific scenarios. We can then analyze these user actions to refine our LCA algorithms through contrastive or reinforcement learning, continuously improving the tool’s accuracy.

Collaborative LCA. To scale beyond individual users, we propose making Living Sustainability a collaborative platform where users can: 1) comment on, refine, and validate system-generated LCAs; 2) suggest alternative data sources and highlight missing factors; 3) flag potential misinformation and erroneous estimates. This approach aligns with prior work on peer-driven knowledge sharing and could enable crowdsourced EI assessments with higher accuracy and coverage. A longitudinal deployment could study how self-learning and community contributions improve system reliability over time.

7.3 Reflecting on Operational Impact

While Living Sustainability aims to promote sustainability, it is important to recognize that the system itself consumes energy and computational resources that contribute to emissions, even when our implementation has integrated several optimizations, including 1) GPT-4o-mini for agent backbone over larger LLMs for its efficiency; 2) custom lightweight NLP algorithms to minimize LLM usage; 3) regular expressions and rule-based methods to reduce token sizes during parsing. Every GPT query consumes 0.34 Wh of electricity [3] and generates approximately 4.32 g CO₂e [67], meaning that a high-frequency, unoptimized system could paradoxically increase emissions while promoting sustainable behaviors. Future work should explore fine-tuned, domain-specific smaller LMs as our task does not require a large output window, and parameter-efficient fine-tuning, such as LoRA [42].

In this work, we take the first step to make accurate EI information accessible across diverse everyday scenarios. While domain-specific environmental calculators exist [23, 34, 54, 57, 76], they remain limited in scope. We acknowledge that such specialized implementations might achieve greater efficiency for narrow use cases. Future work could explore hybrid approaches that leverage pre-computed values from domain-specific tools for common scenarios while maintaining generalizability.

Community-Supported Database and Estimation. We also mitigate the computational cost of autonomous retrieval by maintaining a dynamic database of previously calculated footprints and retrieved emission factor data. This dramatically reduces computational overhead by minimizing redundant queries to LLM agents to retrieve information from online sources. As discussed above, a community-driven approach could further improve efficiency. If Living Sustainability is widely adopted, crowdsourced contributed data could be incorporated into a

shared database. This could enable the use of lightweight nearest-neighbor to estimate the emission factor of new materials or components based on similarity.

Acknowledgments

We thank Electricity Maps for providing electricity carbon intensity data, and all anonymous participants for their contribution to the human studies of this work. We thank Shumin Yu for feedback on the figures, and Jeffrey Heer for discussion. This research was supported by gift funds from Amazon, and NSF awards IIS 2212049, CNS 2310515, CNS 2338736, CNS 2401177.

References

- [1] Omnicl Solutions AB. 2023. The 2030 Calculator. <https://www.2030calculator.com/>
- [2] Abhishek Aggarwal, Cheuk Chi Tam, Dezhi Wu, Xiaoming Li, and Shan Qiao. 2023. Artificial Intelligence-Based Chatbots for Promoting Health Behavioral Changes: Systematic Review. *Journal of Medical Internet Research* 25, 1 (Feb. 2023), e40789. [doi:10.2196/40789](https://doi.org/10.2196/40789)
- [3] Sam Altman. 2025. The Gentle Singularity. <https://blog.samaltman.com/the-gentle-singularity>
- [4] Anders S. G. Andrae and Otto Andersen. 2010. Life cycle assessments of consumer electronics – are they consistent? *The International Journal of Life Cycle Assessment* 15, 8 (Sept. 2010), 827–836. [doi:10.1007/s11367-010-0206-1](https://doi.org/10.1007/s11367-010-0206-1)
- [5] Anders S. G. Andrae and Mikko Samuli Vaija. 2014. To Which Degree Does Sector Specific Standardization Make Life Cycle Assessments Comparable?—The Case of Global Warming Potential of Smartphones. *Challenges* 5, 2 (Dec. 2014), 409–429. [doi:10.3390/challe5020409](https://doi.org/10.3390/challe5020409)
- [6] Autodesk. 2024. Manufacturing Sustainability Insights. <https://apps.autodesk.com/FUSION/en/Detail/Index?id=5902568957837856470>
- [7] Bharathan Balaji, Venkata Sai Gargeya Vunnava, Nina Domingo, Shikhar Gupta, Harsh Gupta, Geoffrey Guest, and Aravind Srinivasan. 2023. Flamingo: Environmental Impact Factor Matching for Life Cycle Assessment with Zero-shot Machine Learning. *ACM J. Comput. Sustain. Soc.* 1, 2 (Dec. 2023), 11:1–11:23. [doi:10.1145/3616385](https://doi.org/10.1145/3616385)
- [8] Bharathan Balaji, Venkata Sai Gargeya Vunnava, Geoffrey Guest, and Jared Kramer. 2023. CaML: Carbon Footprinting of Household Products with Zero-Shot Semantic Text Similarity. In *Proceedings of the ACM Web Conference 2023*. ACM, Austin TX USA, 4004–4014. [doi:10.1145/3543507.3583882](https://doi.org/10.1145/3543507.3583882)
- [9] Tereza Bicalho, Ildo Sauer, Alexandre Rambaud, and Yulia Altukhova. 2017. LCA data quality: A management science perspective. *Journal of Cleaner Production* 156 (2017), 888–898.
- [10] Christina Bremer. 2024. *Exploring the Role of Digital Interventions in Realising Energy Savings*. phd. Lancaster University. <https://eprints.lancs.ac.uk/id/eprint/219129/>
- [11] Christina Bremer, Harshit Gujral, Michelle Lin, Lily Hinkers, Christoph Becker, and Vlad C. Coroamă. 2023. How Viable are Energy Savings in Smart Homes? A Call to Embrace Rebound Effects in Sustainable HCI. *ACM J. Comput. Sustain. Soc.* 1, 1 (Sept. 2023), 4:1–4:24. [doi:10.1145/3608115](https://doi.org/10.1145/3608115)
- [12] Christina Bremer, Bran Knowles, and Adrian Friday. 2022. 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 (CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–11. [doi:10.1145/3491102.3517609](https://doi.org/10.1145/3491102.3517609)
- [13] Christina Bremer, Christian Remy, and Adrian Friday. 2024. Fake Data Leads to Fake Insights: The Challenges of Prototyping Energy Dashboards. *10th Workshop on Computing within Limits (LIMITS 2024)* (2024). <https://computingwithinlimits.org/2024/papers/limits2024-bremer-energy-dashboards.pdf>
- [14] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., 1877–1901. <https://papers.nips.cc/paper/2020/hash/1457c0d6bfc4967418fb8ac142f64a-Abstract.html>
- [15] Dan Cascaval, Rastislav Bodik, and Adriana Schulz. 2023. A Lineage-Based Referencing DSL for Computer-Aided Design. *Proc. ACM Program. Lang.* 7, PLDI (June 2023), 109:76–109:99. [doi:10.1145/3591223](https://doi.org/10.1145/3591223)
- [16] Interuniversity Microelectronics Centre. 2024. imec.netzero. <https://netzero.imec-int.com/>
- [17] Tejaswini Chatty, Bryton L. Moeller, Ioana A. Pantelimon, Catherine D. Parnell, Tahsin M. Khan, Lise Laurin, Jeremy Faludi, and Elizabeth L. Murnane. 2024. EcoSketch: Promoting Sustainable Design through Iterative Environmental Assessment during Early-Stage Product Development. *ACM J. Comput. Sustain. Soc.* 2, 2 (June 2024), 22:1–22:29. [doi:10.1145/3648436](https://doi.org/10.1145/3648436)

- [18] Tingyu Cheng, Zhihan Zhang, Han Huang, Yingting Gao, Wei Sun, Gregory D. Abowd, HyunJoo Oh, and Josiah Hester. 2024. Recy-
tronics: Designing Fully Recyclable Electronics With Varied Form Factors. <https://arxiv.org/abs/2406.09611v1>
- [19] Neil Chulpongsatorn, Mille Skovhus Lunding, Nishan Soni, and Ryo Suzuki. 2023. Augmented Math: Authoring AR-Based Explorable Explanations by Augmenting Static Math Textbooks. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3586183.3606827
- [20] Climatiq. 2025. Climatiq - API for Carbon Footprint Calculations. <https://www.climatiq.io/>
- [21] European Commission. 2005. European Platform on LCA. <https://eplca.jrc.ec.europa.eu/>
- [22] Matthew Conlen and Jeffrey Heer. 2018. Idyll: A Markup Language for Authoring and Publishing Interactive Articles on the Web. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology (UIST '18)*. Association for Computing Machinery, New York, NY, USA, 977–989. doi:10.1145/3242587.3242600
- [23] Microsoft Corporation. 2024. Microsoft Sustainability Calculator. https://appsource.microsoft.com/en-us/product/power-bi/coin-sustainability.emissions_impact_dashboard
- [24] Distill Editorial Team. 2021. Distill Hiatus. *Distill* 6, 7 (July 2021), e31. doi:10.23915/distill.00031
- [25] Didem Dizdaroglu. 2017. The Role of Indicator-Based Sustainability Assessment in Policy and the Decision-Making Process: A Review and Outlook. *Sustainability* 9, 6 (June 2017), 1018. doi:10.3390/su9061018 Number: 6 Publisher: Multidisciplinary Digital Publishing Institute.
- [26] ecoinvent. 2024. ecoinvent - Data with purpose. <https://ecoinvent.org/>
- [27] Ember. 2025. Japan remained the fourth-largest solar generating country globally in 2024, but lags far behind on wind. <https://ember-energy.org/countries-and-regions/japan>
- [28] Zachary Englhardt, Felix Hähnlein, Yuxuan Mei, Tong Lin, Connor Masahiro Sun, Zhihan Zhang, Shwetak Patel, Adriana Schulz, and Vikram Iyer. 2025. Incorporating Sustainability in Electronics Design: Obstacles and Opportunities. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3706598.3713299
- [29] Pierre Friedlingstein, Michael O'Sullivan, Matthew W. Jones, Robbie M. Andrew, Judith Hauck, Peter Landschützer, Corinne Le Quéré, Hongmei Li, Ingrid T. Luijkx, Are Olsen, Glen P. Peters, Wouter Peters, Julia Pongratz, Clemens Schwingshackerl, Stephen Sitch, Josep G. Canadell, Philippe Ciais, Robert B. Jackson, Simone R. Alin, Almut Arneth, Vivek Arora, Nicholas R. Bates, Meike Becker, Nicolas Bellouin, Carla F. Berghoff, Henry C. Bittig, Laurent Bopp, Patricia Cadule, Katie Campbell, Matthew A. Chamberlain, Naveen Chandra, Frédéric Chevallier, Louise P. Chini, Thomas Colligan, Jeanne Decayeux, Laique Djetchouang, Xinyu Dou, Carolina Duran Rojas, Kazutaka Enyo, Wiley Evans, Amanda Fay, Richard A. Feely, Daniel J. Ford, Adrianna Foster, Thomas Gasser, Marion Gehlen, Thanos Gkrizalis, Giacomo Grassi, Luke Gregor, Nicolas Gruber, Özgür Gürses, Ian Harris, Matthew Hefner, Jens Heinke, George C. Hurtt, Yosuke Iida, Tatiana Ilyina, Andrew R. Jacobson, Atul Jain, Tereza Jarníková, Annika Jersild, Fei Jiang, Zhe Jin, Etsushi Kato, Ralph F. Keeling, Kees Klein Goldewijk, Jürgen Knauer, Jan Ivar Korsbakken, Siv K. Lauvset, Nathalie Lefèvre, Zhu Liu, Junjie Liu, Lei Ma, Shamil Maksyutov, Gregg Marland, Nicolas Mayot, Patrick McGuire, Nicolas Metzl, Natalie M. Monacci, Eric J. Morgan, Shin-Ichiro Nakaoaka, Craig Neill, Yosuke Niwa, Tobias Nützel, Lea Olivier, Tsuneo Ono, Paul I. Palmer, Denis Pierrot, Zhangcai Qin, Laure Resplandy, Alizée Roobaert, Thais M. Rosan, Christian Rödenbeck, Jörg Schwinger, T. Luke Smallman, Stephen Smith, Reinel Sospedra-Alfonso, Tobias Steinhoff, Qing Sun, Adrienne J. Sutton, Roland Séférian, Shintaro Takao, Hiroaki Tatebe, Hanqin Tian, Bronte Tilbrook, Olivier Torres, Etienne Tourigny, Hiroyuki Tsujino, Francesco Tubiello, Guido van der Werf, Rik Wanninkhof, Xuhui Wang, Dongxu Yang, Xiaojuan Yang, Zhen Yu, Wenping Yuan, Xu Yue, Sönke Zaehle, Ning Zeng, and Jiye Zeng. 2024. Global Carbon Budget 2024. *Earth System Science Data Discussions* (Nov. 2024), 1–133. doi:10.5194/essd-2024-519 Publisher: Copernicus GmbH.
- [30] Jon Froehlich. 2009. Promoting Energy Efficient Behaviors in the Home through Feedback: The Role of Human-Computer Interaction. *HCIC 2009* (2009). https://makeabilitylab.cs.washington.edu/media/publications/Promoting_Energy_Efficient_Behaviors_in_the_Home_through_Feedback_The_Role_of_Human-Computer_Interaction_mnfyndD.pdf
- [31] Jon Froehlich, Tawanna Dillahunt, Predrag Klasnja, Jennifer Mankoff, Sunny Consolvo, Beverly Harrison, and James A. Landay. 2009. UbiGreen: investigating a mobile tool for tracking and supporting green transportation habits. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. Association for Computing Machinery, New York, NY, USA, 1043–1052. doi:10.1145/1518701.1518861
- [32] Jon Froehlich, Leah Findlater, and James Landay. 2010. The design of eco-feedback technology. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. Association for Computing Machinery, New York, NY, USA, 1999–2008. doi:10.1145/1753326.1753629
- [33] Jon Froehlich, Leah Findlater, Marilyn Ostergren, Solai Ramanathan, Josh Peterson, Inness Wragg, Eric Larson, Fabia Fu, Mazhengmin Bai, Shwetak Patel, and James A. Landay. 2012. The design and evaluation of prototype eco-feedback displays for fixture-level water usage data. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. Association for Computing Machinery, New York, NY, USA, 2367–2376. doi:10.1145/2207676.2208397
- [34] Google. 2021. Google Maps Eco-Friendly Routing. <https://www.gstatic.com/gumdrop/sustainability/google-maps-eco-friendly-routing.pdf>

[35] Greenhouse Gas Protocol. 2011. *Product Life Cycle Accounting and Reporting Standard*. World Resources Institute and World Business Council for Sustainable Development. https://ghgprotocol.org/sites/default/files/standards/Product-Life-Cycle-Accounting-Reporting-Standard_041613.pdf

[36] Aditya Gunturu, Yi Wen, Nandi Zhang, Jarin Thundathil, Rubaiat Habib Kazi, and Ryo Suzuki. 2024. Augmented Physics: Creating Interactive and Embedded Physics Simulations from Static Textbook Diagrams. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*. Association for Computing Machinery, New York, NY, USA, 1–12. [doi:10.1145/3654777.3676392](https://doi.org/10.1145/3654777.3676392)

[37] Udit Gupta, Mariam Elgamal, Gage Hills, Gu-Yeon Wei, Hsien-Hsin S. Lee, David Brooks, and Carole-Jean Wu. 2022. ACT: designing sustainable computer systems with an architectural carbon modeling tool. In *Proceedings of the 49th Annual International Symposium on Computer Architecture (ISCA '22)*. Association for Computing Machinery, New York, NY, USA, 784–799. [doi:10.1145/3470496.3527408](https://doi.org/10.1145/3470496.3527408)

[38] Andrew Head, Kyle Lo, Dongyeop Kang, Raymond Fok, Sam Skjonsberg, Daniel S. Weld, and Marti A. Hearst. 2021. Augmenting Scientific Papers with Just-in-Time, Position-Sensitive Definitions of Terms and Symbols. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–18. [doi:10.1145/3411764.3445648](https://doi.org/10.1145/3411764.3445648)

[39] Jeffrey Heer, Matthew Conlen, Vishal Devireddy, Tu Nguyen, and Joshua Horowitz. 2023. Living Papers: A Language Toolkit for Augmented Scholarly Communication. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. Association for Computing Machinery, New York, NY, USA, 1–13. [doi:10.1145/3586183.3606791](https://doi.org/10.1145/3586183.3606791)

[40] Chris Hendrickson, Arpad Horvath, Satish Joshi, and Lester Lave. 1998. Economic Input–Output Models for Environmental Life-Cycle Assessment. *Environmental Science & Technology* 32, 7 (April 1998), 184A–191A. [doi:10.1021/es983471i](https://doi.org/10.1021/es983471i) Publisher: American Chemical Society.

[41] Cheng-Yu Hsieh, Chun-Liang Li, Chih-kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alex Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. Distilling Step-by-Step! Outperforming Larger Language Models with Less Training Data and Smaller Model Sizes. In *Findings of the Association for Computational Linguistics: ACL 2023*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 8003–8017. [doi:10.18653/v1/2023.findings-acl.507](https://doi.org/10.18653/v1/2023.findings-acl.507)

[42] Edward J. Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=nZeVKeefYF9>

[43] Idemant. 2024. Idemant App. <https://idemantapp.com/>

[44] Thoughtworks Inc. 2025. Cloud Carbon Footprint - An open source tool to measure and analyze cloud carbon emissions. <https://www.cloudcarbonfootprint.org/>

[45] Institute for Sustainable Energy Policies. 2024. 2023 Share of Electricity from Renewable Energy Resources in Japan. <https://www.isep.or.jp/en/1529/>

[46] International Air Transport Association. 2022. Recommended Practice 1678 for Cargo CO₂ Emissions Measurement Methodology. <https://www.iata.org/en/programs/cargo/sustainability/carbon-footprint/>

[47] J.T. Houghton, G.J. Jenkins, and J.J. Ephraums. 1990. *Climate Change: The IPCC Scientific Assessment*. Cambridge University Press. https://www.ipcc.ch/site/assets/uploads/2018/03/ipcc_far_wg_I_full_report.pdf

[48] George Kamiya and Steven Moore. 2025. *Rethinking Mobile Phones: The business case for circularity*. Technical Report. GSMA. https://www.gsma.com/solutions-and-impact/connectivity-for-good/external-affairs/wp-content/uploads/2025/04/Rethinking-Mobile-Phones_Web.pdf

[49] Dae Hyun Kim, Enamul Hoque, Juho Kim, and Maneesh Agrawala. 2018. Facilitating Document Reading by Linking Text and Tables. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology (UIST '18)*. Association for Computing Machinery, New York, NY, USA, 423–434. [doi:10.1145/3242587.3242617](https://doi.org/10.1145/3242587.3242617)

[50] Jesper Kjeldskov, Mikael B. Skov, Jeni Paay, and Rahuvaran Pathmanathan. 2012. Using mobile phones to support sustainability: a field study of residential electricity consumption. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. Association for Computing Machinery, New York, NY, USA, 2347–2356. [doi:10.1145/2207676.2208395](https://doi.org/10.1145/2207676.2208395)

[51] Marion Koelle, Madalina Nicolae, Aditya Shekhar Nittala, Marc Teyssier, and Jürgen Steinle. 2022. Prototyping Soft Devices with Interactive Bioplastics. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. ACM, Bend OR USA, 1–16. [doi:10.1145/3526113.3545623](https://doi.org/10.1145/3526113.3545623)

[52] Kun-Mo Lee and Atsushi Inaba. 2004. *Life Cycle Assessment Best Practices of ISO 14040 Series*. Technical Report. Asia-Pacific Economic Cooperation.

[53] Chu Li, Zhihan Zhang, Michael Saugstad, Esteban Safranchik, Chaitanya Sheerif Kulkarni, Xiaoyu Huang, Shwetak Patel, Vikram Iyer, Tim Althoff, and Jon E. Froehlich. 2024. LabelAIId: Just-in-time AI Interventions for Improving Human Labeling Quality and Domain Knowledge in Crowdsourcing Systems. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–21. [doi:10.1145/3613904.3642089](https://doi.org/10.1145/3613904.3642089)

[54] Google LLC. 2024. Travel Impact Model 1.9.1. <https://github.com/google/travel-impact-model>

[55] Jasmine Lu, Beza Desta, K. D. Wu, Romain Nith, Joyce E Passananti, and Pedro Lopes. 2023. ecoEDA: Recycling E-waste During Electronics Design. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3586183.3606745

[56] Makersite. 2024. Makersite. <https://makersite.io/>

[57] Electricity Maps. 2025. Electricity Maps Carbon Intensity Data. <https://www.electricitymaps.com>

[58] Damien Masson, Sylvain Malacria, Géry Casiez, and Daniel Vogel. 2023. Charagraph: Interactive Generation of Charts for Realtime Annotation of Data-Rich Paragraphs. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3544548.3581091

[59] Eleonora Mencarini, Valentina Nisi, Christina Bremer, Chiara Leonardi, Nuno Jardim Nunes, Jen Liu, and Robert Soden. 2024. Imagining Sustainable Futures: Expanding the Discussion on Sustainable HCI. *interactions* 31, 2 (Feb. 2024), 39–43. doi:10.1145/3643493

[60] Madalina Nicolae, Vivien Roussel, Marion Koelle, Samuel Huron, Jürgen Steimle, and Marc Teyssier. 2023. Biohybrid Devices: Prototyping Interactive Devices with Growable Materials. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3586183.3606774

[61] Elsa Olivetti and Randolph Kirchain. 2012. A Product Attribute to Impact Algorithm to Streamline IT Carbon Footprinting. In *Design for Innovative Value Towards a Sustainable Society*, Mitsutaka Matsumoto, Yasushi Umeda, Keijiro Masui, and Shinichi Fukushige (Eds.). Springer Netherlands, Dordrecht, 747–749. doi:10.1007/978-94-007-3010-6_151

[62] Adiba Orzikulova, Han Xiao, Zhipeng Li, Yukang Yan, Yuntao Wang, Yuanchun Shi, Marzyeh Ghassemi, Sung-Ju Lee, Anind K Dey, and Xuhai Xu. 2024. Time2Stop: Adaptive and Explainable Human-AI Loop for Smartphone Overuse Intervention. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–20. doi:10.1145/3613904.3642747

[63] Katherine Perkins, Wendy Adams, Michael Dubson, Noah Finkelstein, Sam Reid, Carl Wieman, and Ron LeMaster. 2006. PhET: Interactive Simulations for Teaching and Learning Physics. *The Physics Teacher* 44, 1 (Jan. 2006), 18–23. doi:10.1119/1.2150754

[64] United Nations Development Programme. 2024. People's Climate Vote 2024. https://peoplesclimate.vote/document/Peoples_Climate_Vote_Report_2024.pdf

[65] Davide Proserpio, Ali Goli, Tyler Mangini, Ken Lau, and Daniela Yu. 2024. The impact of sustainability programs on consumer purchase behavior: Evidence from Amazon. doi:10.2139/ssrn.5045830

[66] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, Hong Kong, China, 3982–3992. doi:10.18653/v1/D19-1410

[67] Matthias C. Rillig, Marlene Ågerstrand, Mohan Bi, Kenneth A. Gould, and Uli Sauerland. 2023. Risks and Benefits of Large Language Models for the Environment. *Environmental Science & Technology* 57, 9 (March 2023), 3464–3466. doi:10.1021/acs.est.3c01106 Publisher: American Chemical Society.

[68] Gabrielle Samuel, Federica Lucivero, Bran Knowles, and Katherine Wright. 2024. Carbon Accounting in the Digital Industry: The Need to Move towards Decision Making in Uncertainty. *Sustainability* 16, 5 (Jan. 2024), 2017. doi:10.3390/su16052017 Number: 5 Publisher: Multidisciplinary Digital Publishing Institute.

[69] M. Six Silberman, Lisa Nathan, Bran Knowles, Roy Bendor, Adrian Clear, Maria Håkansson, Tawanna Dillahunt, and Jennifer Mankoff. 2014. Next steps for sustainable HCI. *interactions* 21, 5 (Sept. 2014), 66–69. doi:10.1145/2651820

[70] Sluicebox. 2024. Sluicebox. <https://www.sluicebox.ai/>

[71] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. MPNet: masked and permuted pre-training for language understanding. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (NIPS '20)*. Curran Associates Inc., Red Hook, NY, USA, 16857–16867.

[72] Katherine W Song and Eric Paulos. 2021. Unmaking: Enabling and Celebrating the Creative Material of Failure, Destruction, Decay, and Deformation. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3411764.3445529

[73] Katherine W. Song, Samar Sabie, Steven J. Jackson, Kristina Lindström, Eric Paulos, Åsa Ståhl, and Ron Wakkary. 2025. Unmaking and HCI: Techniques, Technologies, Materials, and Philosophies beyond Making. *ACM Trans. Comput.-Hum. Interact.* 31, 6 (Jan. 2025), 76:1–76:6. doi:10.1145/3689047

[74] Sphera. 2024. Life Cycle Assessment Software and Data | Sphera (GaBi). <https://sphera.com/solutions/product-stewardship/life-cycle-assessment-software-and-data/>

[75] Kleis Technology. 2024. LCA as Code. <https://lca-as-code.com/book/>

[76] The Nature Conservancy. 2025. What is your carbon footprint? <https://www.nature.org/en-us/get-involved/how-to-help/carbon-footprint-calculator/>

[77] Vanessa Thomas, Christian Remy, Mike Hazas, and Oliver Bates. 2017. HCI and Environmental Public Policy: Opportunities for Engagement. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 6986–6992. doi:10.1145/3025453.3025579

- [78] Tokyo Electric Power Company. [n. d.]. *ESG Data 2024*. Technical Report. https://www4.tepco.co.jp/en/hd/about/esg/pdf/e_indicators_2024_eng.pdf
- [79] United Nations Environment Programme. 2024. Emissions Gap Report 2024: No more hot air ... please! With a massive gap between rhetoric and reality, countries draft new climate commitments. (Oct. 2024). <https://wedocs.unep.org/20.500.11822/46404> Publisher: United Nations Environment Programme.
- [80] Eldy S. Lazaro Vasquez and Katia Vega. 2019. Myco-accessories: sustainable wearables with biodegradable materials. In *Proceedings of the 2019 ACM International Symposium on Wearable Computers (ISWC '19)*. Association for Computing Machinery, New York, NY, USA, 306–311. [doi:10.1145/3341163.3346938](https://doi.org/10.1145/3341163.3346938)
- [81] Eldy S. Lazaro Vasquez, Hao-Chuan Wang, and Katia Vega. 2020. Introducing the Sustainable Prototyping Life Cycle for Digital Fabrication to Designers. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference (DIS '20)*. Association for Computing Machinery, New York, NY, USA, 1301–1312. [doi:10.1145/3357236.3395510](https://doi.org/10.1145/3357236.3395510)
- [82] Bret Victor. 2011. Explorable Explanations. <https://worrydream.com/ExplorableExplanations/>
- [83] Bret Victor. 2011. Kill Math. <https://worrydream.com/KillMath/>
- [84] Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. MINILM: deep self-attention distillation for task-agnostic compression of pre-trained transformers. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (NIPS '20)*. Curran Associates Inc., Red Hook, NY, USA, 5776–5788.
- [85] Chenming Wu, Haisen Zhao, Chandrakana Nandi, Jeffrey I. Lipton, Zachary Tatlock, and Adriana Schulz. 2019. Carpentry compiler. *ACM Trans. Graph.* 38, 6 (Nov. 2019), 195:1–195:14. [doi:10.1145/3355089.3356518](https://doi.org/10.1145/3355089.3356518)
- [86] Han Xiao. 2024. Regex for chunking by using all semantic cues. <https://gist.github.com/hanxiao/3f60354cf6dc5ac698bc9154163b4e6a>
- [87] Xuhai Xu, Tianyuan Zou, Han Xiao, Yanzhang Li, Ruolin Wang, Tianyi Yuan, Yuntao Wang, Yuanchun Shi, Jennifer Mankoff, and Anind K Dey. 2022. TypeOut: Leveraging Just-in-Time Self-Affirmation for Smartphone Overuse Reduction. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–17. [doi:10.1145/3491102.3517476](https://doi.org/10.1145/3491102.3517476)
- [88] Zhihan Zhang, Agni K. Biswal, Ankush Nandi, Kali Frost, Jake A. Smith, Bichlien H. Nguyen, Shwetak Patel, Aniruddh Vashisth, and Vikram Iyer. 2024. Recyclable vitrimer-based printed circuit boards for sustainable electronics. *Nature Sustainability* 7, 5 (May 2024), 616–627. [doi:10.1038/s41893-024-01333-7](https://doi.org/10.1038/s41893-024-01333-7) Publisher: Nature Publishing Group.
- [89] Zhihan Zhang, Felix Hählein, Yuxuan Mei, Zachary Englhardt, Shwetak Patel, Adriana Schulz, and Vikram Iyer. 2024. DeltaLCA: Comparative Life-Cycle Assessment for Electronics Design. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8, 1 (March 2024), 29:1–29:29. [doi:10.1145/3643561](https://doi.org/10.1145/3643561)
- [90] Zhihan Zhang, Mallory Parker, Kuotian Liao, Jerry Cao, Anandghan Waghmare, Joseph Breda, Chris Matsumura, Serena Eley, Eleftheria Roumeli, Shwetak Patel, and Vikram Iyer. 2024. Biodegradable Interactive Materials. [doi:10.48550/arXiv.2404.03130](https://doi.org/10.48550/arXiv.2404.03130) arXiv:2404.03130 [cs].

A Formative Study

A.1 Study Method

A.1.1 Participant Recruitment. We recruited 6 participants (4 women, 2 men) via industry partnerships and snowball sampling. Our participants represented various domains of expertise in sustainability, including 1 PhD student and 1 professor in sustainable HCI, 1 senior LCA expert and 1 UI/UX designer in industry, and 2 clean energy education specialists. Each participant completed a study session of approximately one hour and was compensated at a rate of \$20 per hour for their time.

A.1.2 Protocol. After obtaining verbal consent, we recorded each session. We began by providing participants with an overview of EI assessment tools and explaining our study objectives. The protocol consisted of three main phases: First, we explored participants' current practices through semi-structured interviews covering: 1) understanding of carbon footprint metrics and EI, 2) current methods for assessing sustainability in daily activities, 3) barriers and motivations for making sustainable choices, and 4) experience with existing environmental assessment tools

Second, we conducted a design elicitation study using two complementary prototypes: a functional web-based tool demonstrating core features, and Figma sketches exploring additional visualization and interaction concepts. Participants were first demonstrated with the working prototype to understand basic functionality, then explored expanded design possibilities through the sketches.

Finally, we gathered feedback on visualization preferences, trust-building elements, and integration requirements. Throughout the sessions, we encouraged participants to draw from their professional expertise while considering the needs of general users, with key findings and suggestions documented directly informing our system design.

B Data Abstraction

LCA Data Abstraction JSON Template

```

# Functional Unit: the reference quantity for all LCA results
"functional_unit": {
  "description": "", # Human-readable definition (e.g. "1 unit of product")
  "source": "" # Whether it was provided or inferred
},

"production": {
  # Materials defined via ratios or mixtures (e.g. 1:2 mix of A and B)
  "related_materials": [
    {
      "text_source": "<text from input>",
      "ratio": [
        {
          "name": "Material 1",
          "ratio_value": <value>, # Numerical ratio value
          "unit": "<unit>", # Unit of measurement (e.g. g)
          "ratio_value_source": "present", # Whether value was present or inferred
          "unit_source": "present",
          "carbon_emission_factor": "<value> g CO2e per <unit>",
          "carbon_emission_source": "<url>", # Source of the emission factor
          "index": 0
        },
        {
          "name": "Material 2",
          "ratio_value": <value>,
          "unit": "<unit>",
          "ratio_value_source": "present",
          "unit_source": "inferred",
          "carbon_emission_factor": "<value> g CO2e per <unit>",
          "carbon_emission_source": "<url>",
          "index": 1
        }
      ],
      # Subprocesses applied to raw materials
      "processes": [
        {
          "name": "Process 1", # Descriptive process name
          "inputs": ["Material 1", "Material 2"],
          "power": <value>,
          "power_unit": "W",
          "power_original": <value>,
          "power_original_unit": "W",
          "time": <value>, # Duration of process
          "process_type": "mixing"
        }
      ]
    }
  ]
}

```

```

    "time_unit": "s",
    "time_original": <value>,
    "time_original_unit": "hr",
    "power_source": "present",
    "time_source": "present",
    "energy": <value>,
    "energy_unit": "kWh", # Derived kWh
    "energy_source": "calculated",
    "outputs": [],
    "waste": [],
    "location": "<location>" # Whether it was provided or missing
    "carbon_emission_factor": "",
    "carbon_emission_source": "<url>",
    "index": 2
  }
],
},
],
}

# Materials directly specified by quantity
"independent_material": [
{
  "name": "Material 3",
  "amount": <value>,
  "unit": "<unit>",
  "amount_source": "present",
  "unit_source": "present",
  "carbon_emission_factor": "<value> g CO2e per <unit>",
  "carbon_emission_source": "<url>",
  "index": 3
}
],
}

# General processing (e.g. assembly)
"processes": [
{
  "name": "Process 2",
  "inputs": [],
  ...
  "index": 4
}
],
}

"transport": {
  # Transport routes for moving materials or products
  "segments": [

```

```
{
  "from_location": "<Location A>", # Origin
  "to_location": "<Location B>", # Destination
  "weight": {
    "value": <value>,
    "unit": "<unit>", # e.g., kg, lb
    "weight_source": "present"
  },
  "vehicle_type": "<vehicle type>", # e.g., refrigerator truck, ship
  "vehicle_type_source": "",
  "carbon_emission_factor": "", # Transport emission factor
  "carbon_emission_source": "<url>",
  "index": 5
},
{
  "from_location": "<Location B>",
  "to_location": "<Location C>",
  ...
  "index": 6
}
],
},
"use": {
  # How the product consumes energy in operation or use
  "operations": [
    {
      "name": "Operation 1",
      "power": <value>,
      "power_unit": "W",
      "power_original": <value>,
      "power_original_unit": "W",
      "time": <value>, # Duration of use
      "time_unit": "hr",
      "time_original": <value>,
      "time_original_unit": "day",
      "power_source": "present",
      "time_source": "present",
      "energy": <value>,
      "energy_unit": "kWh", # Derived kWh
      "energy_source": "calculated",
      "location": "<location>", # Whether it was provided or missing
      "carbon_emission_factor": "<value> g CO2e per kWh",
      "carbon_emission_source": "<url>",
      "index": 7
    }
  ]
}
```

```
  ],
},
# Plain-language summary for the general public in UI information panel
"explanation": "<generated summary>"
}
```

LCA Data Abstraction Example of Figure 5

```
{
  "functional_unit": {
    "description": "1 unit of GFRV composite",
    "source": "inferred"
  },
  "production": {
    "related_materials": [
      {
        "text_source": "stoichiometric ratio (1:1:5 mol% to the acid) of epoxy (EPON 828, Skygeek), adipic acid (Sigma Aldrich) and 1,5,7-triazabicyclo[4.4.0]dec-5-ene (TBD, Sigma Aldrich)",
        "ratio": [
          {
            "name": "adipic acid",
            "ratio_value": 146.14,
            "unit": "g/mol",
            "ratio_value_source": "inferred",
            "unit_source": "inferred",
            "carbon_emission_factor": "8.87 g CO2e per g",
            "carbon_emission_source": "https://ecoinvent.org/ecoinvent-v3-10/",
            "index": 0
          },
          {
            "name": "EPON 828",
            "ratio_value": 340,
            "unit": "g/mol",
            "ratio_value_source": "inferred",
            "unit_source": "inferred",
            "carbon_emission_factor": "5.97 g CO2e per g",
            "carbon_emission_source": "https://single-market-economy.ec.europa.eu/industry/sustainability/sustainable-production",
            "index": 1
          },
          {
            "name": "TBD",
            "ratio_value": 6.96,
            "unit": "g/mol",
            "ratio_value_source": "inferred",
            "unit_source": "inferred",
            "carbon_emission_factor": "3.51 g CO2e per g",
            "carbon_emission_source": "https://www.climatiq.io/data",
            "index": 2
          }
        ],
      }
    ]
  }
}
```

```

  "processes": [],
  },
  "independent_material": [
    {
      "name": "woven glass fibre sheets",
      "amount": 1,
      "unit": "g",
      "amount_source": "inferred",
      "unit_source": "inferred",
      "carbon_emission_factor": "2.60 g CO2e per g",
      "carbon_emission_source":
        "https://legacy.winnipeg.ca/finance/findata/matmgt/documents/2012/682-2012/682-2012_app"
      "index": 3
    },
    "processes": []
  },
  "transport": {
    "segments": []
  },
  "use": {
    "operations": []
  },
  "explanation": "The GFRV composite is produced by combining vitrimer polymer and
  woven glass fibre sheets. The vitrimer polymer is synthesized through a
  stoichiometric reaction involving EPON 828, adipic acid, and TBD."
}

```

LCA Data Abstraction Example of [Figure 9](#)

```
{
  "functional_unit": {
    "description": "Shipping a package from an unspecified location to Berkeley, CA, USA",
    "source": "inferred"
  },
  "production": {
    "related_materials": [],
    "independent_material": [],
    "processes": []
  },
  "transport": {
    "segments": [
      {
        "from_location": null, # Missing location!
        "to_location": "Berkeley, CA, USA",
        "weight": {
          "value": null, # Missing weight!
          "unit": null,
          "weight_source": null
        },
        ...
      }
    ]
  },
  "use": {
    "operations": []
  },
  "explanation": "The user wants to ship a package to Berkeley, the weight and origin were not specified."
}
```

LCA Data Abstraction Example of [Figure 10](#)

```
{
  "functional_unit": {
    "description": "Training a transformer model using an Nvidia RTX 4090 for 24
      hours",
    "source": "present"
  },
  "production": {
    "related_materials": [],
    "independent_material": [],
    "processes": []
  },
  "transport": {
    "segments": []
  },
  "use": {
    "operations": [
      {
        "name": "Transformer model training on Nvidia RTX 4090",
        "power": 450,
        "power_unit": "W",
        "power_original": 450,
        "power_original_unit": "W",
        "time": 24,
        "time_unit": "hr",
        "time_original": 24,
        "time_original_unit": "hr",
        "power_source": "inferred",
        "time_source": "present",
        "energy": 10.8,
        "energy_unit": "kWh",
        "energy_source": "calculated",
        "location": null, # Unspecified location
        "carbon_emission_factor": "480 g CO2e per kWh", # Use an global average
          grid emission factor
        "carbon_emission_source":
          "https://ember-energy.org/latest-insights/global-electricity-review-2024/global-electri
        "index": 7
      }
    ]
  },
  "explanation": "The training of a transformer model using an Nvidia RTX 4090 for
  24 hours was modeled. The RTX 4090 has an estimated power draw of 450 W. Over
  24 hours, this results in a total energy consumption of 10.8 kWh. Using an
  average global grid emission factor because the location is unspecified."
}
```

LCA Data Abstraction of Example 2 in Table A6 Production

```
{
  "functional_unit": {
    "description": "1 batch of bread",
    "source": "present"
  },

  "production": {
    "related_materials": [
      {
        "text_source": "644g bread flour, 414g milk, 86g eggs, 64g sugar, and 6g active dry yeast",
        "ratio": [
          {
            "name": "bread flour",
            "ratio_value": 644,
            "unit": "g",
            "ratio_value_source": "present",
            "unit_source": "present",
            "carbon_emission_factor": "0.52 g CO2e per g",
            "carbon_emission_source": "https://shrinkthatfootprint.com/processed-agricultural-products-carbon-footprint-data",
            "index": 0
          },
          {
            "name": "milk",
            "ratio_value": 414,
            "unit": "g",
            "ratio_value_source": "present",
            "unit_source": "present",
            "carbon_emission_factor": "1.64 g CO2e per g",
            "carbon_emission_source": "https://klimato.com/klimato-insights/how-to-optimize-recipes-in-klimato-an-expert-guide",
            "index": 1
          },
          {
            "name": "eggs",
            "ratio_value": 86,
            "unit": "g",
            "ratio_value_source": "present",
            "unit_source": "present",
            "carbon_emission_factor": "4.67 g CO2e per g",
            "carbon_emission_source": "https://ourworldindata.org/environmental-impacts-of-food",
            "index": 2
          }
        ]
      }
    ]
  }
}
```

```
{
  "name": "sugar",
  "ratio_value": 64,
  "unit": "g",
  "ratio_value_source": "present",
  "unit_source": "present",
  "carbon_emission_factor": "0.62 g CO2e per g",
  "carbon_emission_source":
    "https://shrinkthatfootprint.com/processed-agricultural-products-carbon-footprint-data",
  "index": 3
},
{
  "name": "active dry yeast",
  "ratio_value": 6,
  "unit": "g",
  "ratio_value_source": "present",
  "unit_source": "present",
  "carbon_emission_factor": "3.29 g CO2e per g",
  "carbon_emission_source":
    "https://shrinkthatfootprint.com/processed-agricultural-products-carbon-footprint-data",
  "index": 4
},
],
"independent_material": [],
"processes": []
},
"transport": {
  "segments": []
},
"use": {
  "operations": []
},
"explanation": ""
}
```

C Dataset

LCA Stage	# Example
Production	<ol style="list-style-type: none"> 1. Ultimaker Cura estimates that printing this shape will use 34g of PLA filament. 2. This bread recipe calls for 644g bread flour, 414g milk, 86g eggs, 64g sugar, and 6g active dry yeast. 3. Two 25-μm-thick sheets of copper foil (1 oz, McMaster-Carr) were first laser cut.
Transport	<ol style="list-style-type: none"> 1. My friend took an 8-day road trip to move from the East Coast to Seattle in her car with her belongings and her cat. 2. The headphones are transported from the final assembly in China to the distribution hub in the Netherlands mostly by ship. 3. Moving this machine of size 71" \times 39" \times 35" and weighing 750 lbs will require a U-Haul truck.
Use	<ol style="list-style-type: none"> 1. The shape that I'm going to 3D print will take about 4 hours on the Ultimaker S3. 2. Bake bread dough in the oven for 40 minutes at 400 F. 3. A washing machine consumes 1.5 kWh of electricity per cycle.

Table 6. Natural language description examples of each LCA stage: production, transportation, and use.

D Embodied Emissions for Cloud

The embodied emissions of the virtual machine instance are calculated based on the following amortization methodology, which allocates the embodied emissions over the design lifetime of the cloud infrastructure. The embodied emissions $E_{embodied}$ for a cloud instance are given by:

$$E_{embodied} = E_{total} \times \frac{t_{use}}{t_{lifetime}} \quad (5)$$

where:

- E_{total} is the total embodied emissions, the cumulative amount of emissions produced by all hardware within the server.
- t_{use} is the duration for which the hardware is reserved to provide the cloud service.
- $t_{lifetime}$ is the expected lifetime of the server hardware before refresh.

E User Evaluation Study

E.1 Study Method

E.1.1 Participant Recruitment. We recruited 12 participants via university mailing lists, industry partnerships, and snowball sampling. Based on self-reported demographics, we had 6 men, 5 women, and 1 preferred not to say. Our participants included 2 founders of AI for LCA companies. Each participant was compensated at a rate of \$20 per hour for their participation.

E.1.2 Protocol. Sessions were conducted both in-person and via Zoom to accommodate geographically distributed participants from the EU, US, and Asia. Throughout the session, we employed a think-aloud methodology, where participants were encouraged to verbalize their thoughts and any confusion while interacting with the tool.

Each session lasted approximately one hour. After obtaining verbal consent, we began with a demo walkthrough using the paper reading example (§5.1) to introduce participants to the system's basic functionality. Participants then independently interacted with four predefined use cases: online shopping, shipping, cloud server, and general web browsing. We deliberately provided minimal guidance during this phase to assess whether users could intuitively understand and navigate the tool's interface and functionality. Once participants were familiar with the system's capabilities, they transitioned to an open exploration phase where they could freely use the tool on any website or text of their choice for EI analysis, helping us to understand users' actual interests in sustainability information in their everyday decision-making and uncovering previously unconsidered use cases.

After the hands-on portion, participants completed a questionnaire on the system's usability, also intuitiveness and usefulness across different scenarios. It is noteworthy that we did not include comparative evaluations against existing LCA tools, as current solutions either lack the capability to assess EI across general daily activities or require significant expertise to operate which we do not expect participants to learn within the session period.

The session concluded with a semi-structured interview exploring three key areas: the tool's triggering scenarios, user expectations, and potential improvements. We discussed their experiences and desired additional features so the tool could better support sustainability awareness and decision-making. This final discussion was particularly valuable for understanding potential market applications as we prepared for a public launch on the Chrome web store.