

DOI: https://doi.org/10.48009/2_iis_2026_126

The cognitive footprint of large language models: User awareness and sustainable intentions

Michael Greco, *Middle Georgia State University*, mikegreco2@gmail.com

Abstract

The rapid adoption of large language models (LLMs) has raised concerns about environmental impact, yet limited research has examined how users understand and respond to these costs. This qualitative study explores the cognitive footprint of LLMs, defined as the psychological processes through which users become aware of, interpret, and respond to a technology's environmental impacts. Guided by the Theory of Planned Behavior (TPB) and Value-Belief-Norm (VBN) Theory, semi-structured interviews were conducted with fifteen adult participants representing diverse professional backgrounds and levels of LLM engagement. Thematic analysis identified five themes: environmental impacts remained abstract until quantified through relatable comparisons; emotional responses varied by technical background; participants weighed environmental costs against perceived utility; responsibility was assigned primarily to technology providers; and sustainable intentions depended on low behavioral friction. Findings suggest that the awareness gap identified in prior research is not only an information deficit but also a framing deficit. They also indicate that sustainable AI usage decisions may depend on interface-level transparency, visible choices, and users' perceived ability to act. The study proposes cognitive footprint as an exploratory concept for understanding user responses to the environmental impacts of LLMs.

Keywords: Cognitive footprint, large language models, environmental sustainability, user awareness, sustainable intentions.

Introduction

Since the public release of ChatGPT in late 2022, large language models (LLMs) have become widely integrated into professional and personal digital life. Major corporations, including Google, Microsoft, and X, have launched competing models, and public reports suggest that leading systems such as ChatGPT and Gemini have reached hundreds of millions of users (Paris, 2025; Sanz, 2025). This rapid adoption reflects the versatility of LLMs for tasks such as natural language processing, question answering, content generation, and complex problem solving (Cheung et al., 2025; Maliakel et al., 2025; Sánchez-Mompó et al., 2025).

However, the growth of LLM use has raised important questions about environmental sustainability. The development and operation of these systems require substantial energy and water during both training and inference (Jegham et al., 2025; Morrison et al., 2025). Estimates suggest that training large models can consume significant electricity and produce substantial carbon emissions, while data centers also require water for cooling specialized hardware (Jegham et al., 2025; Morrison et al., 2025). These impacts have

contributed to growing interest in Green IT and sustainable computing as researchers seek to address the resource demands of digital technologies (Ashiq et al., 2019; Esfahani et al., 2015).

Despite increasing documentation of the technical environmental footprint of LLMs, less is known about the cognitive footprint, or the extent to which users recognize, interpret, and internalize these environmental costs. Prior research has examined model efficiency and environmental impact, but the role of the end user remains less developed (Cheung et al., 2025; Doudkin et al., 2025; Dungo et al., 2025; Jiang et al., 2024; Liu & Yin, 2024). Limited evidence suggests that many users have incomplete awareness of the energy and water consumption associated with everyday LLM use (Dungo et al., 2025). This awareness gap presents a challenge for sustainable technology use because user behavior may contribute to the overall environmental impact of AI systems (Ashiq et al., 2019; Esfahani et al., 2015; Wu et al., 2025).

To examine this issue, the study draws on the Theory of Planned Behavior (TPB) and the Value-Belief-Norm Theory (VBN). These frameworks suggest that awareness alone may be insufficient to produce sustainable behavior; perceived behavioral control, personal values, responsibility attribution, and behavioral intention also shape whether users believe action is possible or worthwhile (Bosnjak et al., 2020; Stern et al., 1999). Currently, it remains unclear how users interpret environmental information about LLMs or how that awareness influences their attitudes, perceived ability, responsibility, and intentions regarding sustainable use.

The purpose of this qualitative study is to explore how users understand and interpret the environmental impacts of LLMs and how this awareness influences their attitudes, perceived behavioral control, personal values, and intentions to engage in sustainable usage behaviors. Specifically, the study examines how users respond to information about the energy consumption, carbon emissions, and water usage associated with LLMs, and whether this information shapes their willingness to consider practices such as reducing unnecessary queries or selecting lower-impact AI systems. To identify themes emerging from participants' interview data, the study was guided by the following research questions:

RQ1: How do users describe their baseline awareness and understanding of the environmental impacts of LLMs before being presented with factual information?

RQ2: How do users interpret and make sense of new information about the energy, carbon, and water footprint of LLMs when it is presented to them?

RQ3: How does the presentation of environmental impact data influence users' attitudes toward AI sustainability?

RQ4: After receiving environmental impact information, how do users perceive their ability and responsibility to engage in sustainable LLM usage behaviors?

RQ5: How do users articulate their motivations or hesitations regarding sustainable LLM usage after receiving environmental impact information?

Review of the Literature

The increasing use of large language models (LLMs) has raised significant concerns about environmental sustainability because the development, training, and operation of these models require substantial energy and water and may contribute to electronic waste (Kocak et al., 2025; Jegham et al., 2025; Morrison et al., 2025; Wang et al., 2025; Wu et al., 2025). This situation has fostered the growth of fields such as "Green IT" and Sustainable Human-Computer Interaction (HCI), which focus on designing and deploying systems with greater awareness of their environmental and social impacts (Ashiq et al., 2019; Doudkin et al., 2025; Esfahani et al., 2015).

While significant research has been conducted on the technical aspects of LLM efficiency, the role of the end user has not been adequately explored (Dungo et al., 2025; Cheung et al., 2025; Liu & Yin, 2024; Doudkin et al., 2025; Jiang et al., 2024). This section reviews the literature regarding the environmental footprint of LLMs, the debate surrounding their net impact, and the theoretical frameworks, specifically the Theory of Planned Behavior (TPB) and Value-Belief-Norm (VBN) Theory, that provide the lens for understanding user-driven sustainability (Ashiq et al., 2019; Doudkin et al., 2025; Esfahani et al., 2015).

The Environmental Footprint of Large Language Models

The environmental footprint of LLMs stems from substantial energy and water consumption and spans their lifecycle, from development and training to daily use by end users (Jegham et al., 2025; Morrison et al., 2025). Because environmental estimates vary substantially depending on model size, workload, hardware configuration, data-center location, cooling method, and local energy mix, the figures below should be interpreted as illustrative estimates rather than fixed values.

Energy Consumption in the AI Lifecycle

LLM energy demands span both training and inference, or daily use (Jegham et al., 2025; Morrison et al., 2025; Pimenow et al., 2024). Training requires significant computational resources, with one estimate indicating that training an LLM consumed approximately 1,287 MWh of electricity and emitted over 550 tons of CO₂ equivalent (Morrison et al., 2025). Inference may account for much of lifecycle energy consumption as models respond to large volumes of user queries, and inference energy varies by model size, task type, hardware configuration, and system design (Jegham et al., 2025; Maliakel et al., 2025; Sánchez-Mompó et al., 2025). Water consumption is also important because data centers require water for cooling, and prior estimates describe both substantial water use during model development and possible water costs during individual interactions (Jegham et al., 2025; Jiang et al., 2024; Morrison et al., 2025).

The Debate on Net Environmental Impact: The Sustainability Paradox

Although energy and water consumption raise important sustainability concerns, the overall environmental impact of LLMs is complex because the technology may function as both a contributor to environmental harm and a tool for efficiency gains. This duality is often described as the sustainability paradox (Kocak et al., 2025). Some research suggests that LLMs may reduce environmental costs in certain task-specific comparisons; however, these potential benefits remain in tension with the growing resource demands of large-scale AI development and deployment (Kocak et al., 2025; Morrison et al., 2025; Ren et al., 2024). This uncertainty underscores the need for transparent environmental reporting and greater attention to how users understand the sustainability implications of LLM use.

Sustainable Human-Computer Interaction and Green IT

Addressing the environmental footprint of LLMs requires a focus on the fields of “Green IT/IS” and human-computer interaction (HCI). The discussion surrounding technology and sustainability can be divided into two main approaches. The first is Green IT, which focuses on mitigating the negative environmental impacts of technology across its entire lifecycle, including production, use, and disposal (Ashiq et al., 2019; Esfahani et al., 2015). In the second approach, Green IS is viewed as part of the solution (Esfahani et al., 2015), involving the use of information systems to improve environmental sustainability, for example, using logistics software to optimize delivery routes (Ashiq et al., 2019; Esfahani et al., 2015). The study of LLM user awareness and its environmental implications is situated at the intersection of these fields in the domain of sustainable HCI. A key principle of Sustainable HCI is that user awareness is essential for promoting pro-environmental behavior (Ashiq et al., 2019; Esfahani et al., 2015; Si et al., 2022; Wei et al., 2025). A large-scale empirical study involving 1,005 residents confirmed this connection, revealing that environmental awareness has a significant positive impact on an individual's pro-

environmental actions (Si et al., 2022). However, simply having access to information is not sufficient; sustainable action also depends on attitudes, knowledge, willingness to act, perceived control, and responsibility attribution (Ashiq et al., 2019). Therefore, behavioral intention theories are useful for examining how users interpret environmental information and decide whether sustainable LLM use feels possible.

Theoretical Frameworks for Sustainable Intention Formation

The Theory of Planned Behavior (TPB) and Value-Belief-Norm (VBN) Theory provide models for examining how individual values, beliefs, perceived control, responsibility, and social contexts influence pro-environmental intentions and actions (Stern et al., 1999).

The Theory of Planned Behavior (TPB)

The Theory of Planned Behavior (TPB) posits that behavioral intentions are shaped by attitudes, subjective norms, and perceived behavioral control (PBC) (Bosnjak et al., 2020). In sustainability contexts, PBC is especially relevant because individuals may be less likely to form sustainable intentions when they do not believe they have meaningful control over the behavior or its outcomes (Bosnjak et al., 2020; Wei et al., 2025; Zhao et al., 2025). Prior sustainability research has identified PBC as an important predictor of pro-environmental intention, making it useful for examining whether LLM users believe they can meaningfully reduce the environmental impact of their own AI use (Wei et al., 2025; Zhao et al., 2025).

The Value-Belief-Norm (VBN) Theory

Complementing this, the Value-Belief-Norm (VBN) Theory illustrates how personal values (biospheric, altruistic, or egoistic) shape environmental beliefs and personal norms that drive pro-environmental behavior (Stern et al., 1999; Kumar et al., 2025; Al Mamun et al., 2022). It proposes a causal chain in which personal values shape a person's ecological worldview, which in turn increases their awareness of the consequences of their actions and their sense of responsibility (Stern et al., 1999; Kumar et al., 2025; Shang et al., 2023).

Prior studies have applied VBN theory to explain pro-environmental intentions and energy-conservation behaviors across consumer and workplace contexts (Al Mamun et al., 2022; Kumar et al., 2025). These findings suggest that users' underlying value orientations may shape how they interpret environmental consequences, assign responsibility, and develop personal norms related to sustainable technology use (Stern et al., 1999; Al Mamun et al., 2022; Kumar et al., 2025).

Research Gap: The Cognitive Footprint

While the environmental implications of LLMs are established, empirical research linking this awareness to user meaning-making is limited (Ashiq et al., 2019; Bosnjak et al., 2020; Dungo et al., 2025; Jegham et al., 2025; Morrison et al., 2025; Stern et al., 1999; Wu et al., 2025). Research has prioritized model optimization over user perception, resulting in an “awareness gap” in which users remain largely unaware of the resources required for generative AI (Dungo et al., 2025). A 2025 survey of college students revealed that, while they were generally aware of the environmental impact of generative AI, their understanding was limited; energy consumption was better recognized, whereas awareness of water usage was notably less comprehensive (Dungo et al., 2025). Consequently, the behavioral factors and meaning-making processes that shape end users' engagement with these technologies have been overlooked. Few studies have examined in depth how users understand the environmental impacts of LLMs or how this awareness shapes their attitudes, values, perceived responsibility, and intentions regarding sustainable use practices. This lack of awareness presents a significant barrier to

promoting sustainable technology use and advancing Green IT initiatives (Ashiq et al., 2019; Esfahani et al., 2015).

Methodology

Research Design

This study utilized a qualitative descriptive research design. This approach was appropriate because it allowed for an in-depth exploration of participants' perspectives, reactions, and meaning-making processes regarding LLM environmental impacts (Dungo et al., 2025; Kumar et al., 2025). The design aligned with the interpretive nature of the guiding theoretical frameworks, the Theory of Planned Behavior and Value-Belief-Norm Theory, which emphasize user beliefs and subjective interpretations as precursors to behavioral intentions (Al Mamun et al., 2022; Bosnjak et al., 2020). By employing semi-structured interviews, the study captured rich, descriptive data on the awareness gap and the behavioral factors influencing the use of sustainable technology (Ashiq et al., 2019; Esfahani et al., 2015).

Participants and Sampling

The target population for this study was adults aged 18 years or older who were familiar with large language models (LLMs) such as ChatGPT, Gemini, Copilot, or Claude. A purposive sampling strategy was used to recruit 15 participants with direct experience relevant to the research questions. This sample size was considered sufficient for a qualitative descriptive study because prior research suggests that saturation in semi-structured interview studies with relatively homogeneous participants often occurs after approximately 12 interviews (Guest et al., 2006). In this study, thematic saturation was observed during the later stages of data collection. By approximately the twelfth interview, participant responses continued to reinforce existing categories, including abstract environmental awareness, surprise regarding water usage, responsibility attribution to providers, and conditional willingness to reduce low-value LLM use. The final interviews confirmed these patterns rather than producing new thematic categories, indicating that sufficient depth had been reached for the purposes of this qualitative descriptive study.

The sampling strategy sought to include participants with varying levels of LLM use and diverse professional perspectives, while maintaining a shared baseline of prior familiarity with LLMs. Recruitment was conducted through direct outreach within the researcher's professional and social networks, including LinkedIn, Instagram, Reddit, and professional contacts. Individuals meeting the eligibility criteria were contacted privately with a brief description of the study and an invitation to participate. Participation was voluntary. To minimize potential power differentials, individuals under the researcher's direct supervision were not recruited.

Participant characteristics were collected during the opening portion of each interview. Participants were asked to describe their age range, work or school background, general comfort with technology, frequency of LLM use, and common LLM use cases. The final sample included adults with varying professional backgrounds, levels of comfort with technology, and patterns of LLM engagement. Participants ranged in age from their 30s to 60s, with work or educational backgrounds spanning information technology, healthcare, business, communications, and environmental science. Participants described their LLM use as ranging from occasional personal use to frequent professional use, with common tasks including information seeking, writing support, brainstorming, summarization, learning assistance, and work-related productivity.

For analytic comparison, participants were categorized as technical or non-technical users. Technical users were participants whose education, occupation, or regular work involved information technology,

cybersecurity, software development, data systems, or related technical decision-making. Non-technical users were participants who used LLMs personally or professionally but did not work in a technical computing role. The final sample included 8 technical users and 7 non-technical users.

Instrumentation

Data were collected using a researcher-developed, semi-structured interview protocol. The instrument was deductively developed from the study's guiding theoretical frameworks, the Theory of Planned Behavior (TPB) (Ajzen, 1991) and Value-Belief-Norm (VBN) theory (Stern et al., 1999). Interview questions were intentionally aligned with core TPB constructs, including attitudes toward the behavior, subjective norms, and perceived behavioral control, as well as VBN constructs such as values, ecological beliefs, awareness of consequences, and personal norms.

The protocol was structured to explore four domains: (a) participants' baseline awareness of LLM environmental impacts, (b) interpretation and meaning-making of standardized factual information, (c) attitudinal and value-based reactions, and (d) stated behavioral intentions related to sustainable technology use. Questions were sequenced first to elicit unprompted perceptions and beliefs, followed by reflective prompts after exposure to a standardized informational summary. This structure enabled examination of potential shifts in attitudes, perceived responsibility, and behavioral intentions consistent with TPB and VBN theoretical pathways.

The semi-structured format ensured consistency across interviews while allowing flexibility for probing and clarification. Follow-up questions were used to explore participants' reasoning processes, perceived social influences, and perceived constraints or facilitators of behavior in greater depth.

The interview included a scripted Standardized Information Prompt summarizing empirically documented findings on LLM energy consumption, water usage, carbon emissions, and the debate over AI's net environmental impact. The prompt was read verbatim to each participant before the reflective questions:

I would like to share a quick summary of the environmental resources required to build and run these models:

First, Training: Teaching a single large model like GPT-4 to 'learn' requires significant computing power. For example, research shows that training a single large model can consume over 1,000 megawatt-hours of electricity, equivalent to the annual usage of hundreds of homes.

Second, Water: Data centers consume water for cooling. Research estimates that training a specific large model consumed approximately 2.7 million liters of freshwater—roughly the amount one person would use over 24 years. In addition to electricity and water, training a single large language model has been estimated to emit hundreds of tons of CO₂-equivalent emissions, comparable to the lifetime emissions of dozens of cars.

Third, Daily Use: Every time you ask a question, the system uses a small amount of energy and water. Some benchmarks suggest that under certain conditions, a single interaction can have a water footprint of up to 500mL.

However, it is a debate. Some researchers argue that AI is a 'net positive' because it helps us work faster and solve complex climate problems efficiently. Do you have any questions about this information before we continue?

The use of a standardized stimulus ensured that all participants received the same information prior to reflective questioning, thereby enhancing procedural consistency across interviews. The development and structured use of this prompt were consistent with systematic approaches to semi-structured interview guide design that emphasize the use of prior knowledge and transparent methodological justification (Kallio et

al., 2016). Participants were then asked to reflect on their understanding and potential behavioral intentions immediately following the prompt.

Procedures

Upon receiving approval from the Institutional Review Board (IRB), interviews were scheduled via secure video conferencing platforms. Before the interview, participants provided informed consent, acknowledging the study's purpose, confidentiality measures, and their agreement to be audio-recorded. The interview sessions lasted between 20 and 30 minutes. During the session, the researcher followed the approved protocol, administering the informational prompt and asking the associated research questions. All interviews were recorded and subsequently transcribed verbatim for analysis. To ensure privacy, all data were de-identified by removing names and other personal identifiers during transcription.

Data Analysis

The transcribed interview data were analyzed using thematic analysis as outlined by Braun and Clarke (2006). Analysis followed the six-phase process of familiarization, initial coding, theme development, theme review, defining and naming themes, and reporting. The researcher first read each transcript in full to become familiar with participants' responses and to note early patterns related to awareness, interpretation, responsibility, perceived control, and behavioral intention.

Coding was conducted manually by the researcher using a structured Microsoft Excel codebook. Initial codes were developed through a hybrid coding strategy. Deductive codes were informed by the study's theoretical frameworks, including TPB constructs such as attitudes, subjective norms, perceived behavioral control, and behavioral intention, as well as VBN constructs such as values, awareness of consequences, responsibility attribution, and personal norms. Inductive codes were also developed from participants' own language and recurring patterns in the data, including abstract environmental awareness, surprise regarding water usage, cost-benefit reasoning, responsibility assigned to providers, and willingness to reduce low-value or frivolous LLM use.

The codebook was iteratively refined throughout the analysis. Codes were reviewed for overlap, combined when conceptually similar, and separated when participant responses reflected distinct meanings. Candidate themes were developed by grouping related codes into broader patterns of shared meaning across the dataset. Themes were then reviewed against the full set of transcripts to ensure they were grounded in participant responses and addressed the study's research questions. Discrepant cases were retained during analysis rather than excluded, particularly when technical users expressed concern but had low willingness to change their behavior, or when non-technical users showed surprise without clear behavioral intention. Direct quotations were mapped to corresponding codes and themes to support analytic credibility and to preserve a clear connection between the data and the findings.

Researcher Positionality

The researcher has a professional background in information technology and familiarity with LLMs and sustainability concerns. This prior knowledge may have influenced the interpretation of participant responses. To mitigate potential bias, the researcher used a semi-structured interview protocol, a standardized informational prompt, systematic coding procedures, and documentation of analytic decisions throughout the research process.

Trustworthiness

To support the rigor of this qualitative study, trustworthiness was addressed through credibility, dependability, confirmability, and transferability. Credibility was supported by a standardized informational prompt, ensuring that all participants received the same environmental impact information

before responding to reflective questions. Interviews were audio-recorded and transcribed verbatim, and the semi-structured interview format allowed the researcher to ask follow-up questions for clarification and elaboration.

Dependability was supported through a systematic coding process and the use of an audit trail. Coding decisions, code definitions, revisions, exemplar quotations, and theme development notes were documented in a Microsoft Excel codebook. This process allowed the researcher to track how initial codes were developed, revised, combined, and organized into final themes. The use of Braun and Clarke's (2006) six-phase thematic analysis also provided a consistent analytic structure.

Confirmability was addressed by maintaining a clear link between participant responses and the study findings. Direct quotations were used to demonstrate how themes were grounded in the data rather than solely in the researcher's interpretation. Because the researcher served as the sole coder, reflexive attention was given to potential bias throughout the analysis. The researcher's professional background in information technology and familiarity with LLMs could have shaped interpretation; therefore, the study relied on a standardized interview protocol, consistent coding procedures, documentation of analytic decisions, and attention to discrepant cases to reduce the risk of selective interpretation.

Transferability was supported by providing descriptive information about the sample, including participants' technical or non-technical orientation, professional backgrounds, age ranges, LLM usage patterns, and common use cases. Although the findings are not intended for statistical generalization, these contextual details allow readers to assess whether the findings may be applicable to similar populations or digitally mediated sustainability contexts.

Results

This analysis is based on data gathered from fifteen adult participants with diverse professional backgrounds and varying levels of engagement with LLMs, ranging from occasional personal use to daily professional reliance. The Methodology section provides a detailed description of participant demographics and sampling procedures. This section presents the findings as patterns of shared meaning across the dataset, supported by illustrative participant quotations (Braun & Clarke, 2006). Five themes emerged from the analysis: environmental impact remains abstract until it is quantified; uneven distribution of surprise and concern; sustainability paradox and cognitive tension; responsibility assigned upward; and conditional and context-dependent intentions. These themes are presented below and aligned with the study's five research questions.

Environmental Impact Remains Abstract Until Quantified

This theme addresses RQ1 by showing that participants' baseline awareness of LLM environmental impacts was generally abstract, incomplete, or limited prior to receiving the informational prompt. While it was generally assumed that technology consumes energy, the specific magnitude of this consumption and the existence of a water footprint were largely unrecognized. As Participant 10 (P10) acknowledged, "I'm sure there's some sort of data storage area where they have to keep all the data... I know it's happening in the background. I just don't think about it" (P10, non-technical user). The lack of awareness surrounding water usage was particularly striking, with P12 noting, "Definitely the water usage, that's kind of crazy... I had some idea there was some electricity there, but the water usage was a surprise" (P12, non-technical user).

The presentation of the Standardized Information Prompt, particularly through comparisons to household energy use or water consumption in liters, served as a catalyst for understanding. P3 explained, "So the way that you've broken it down makes it relatable. You know that use one person's usage over 24 years, that's an astounding amount, but I understand that" (P3, non-technical user).

Uneven Distribution of Surprise and Concern

Addressing Research Questions 2 and 3, this theme highlights participants' differing reactions to the environmental data. Participants classified as non-technical users described more immediate, affective responses, often because they had previously perceived LLMs as immaterial. P7's response illustrates how the water comparison reframed LLM use as physically wasteful: "Wow... That's [a] staggering amount... the water usage really stood out to me... I mean, 24 years... That blows my mind... that just is an alarming amount of waste" (P7, non-technical user). Participants with technical backgrounds, however, viewed the findings as confirmatory. P8 stated, "You didn't bring any new information to me... It confirms my concern" (P8, technical user).

Although participants expressed concern, many did not describe clear or immediate actions they could take in response to the information. As P2 summarized, "I don't know how I feel... anytime you do use it, you're... contributing to those problems... because I have no idea how to use it any differently than what I already do" (P2, non-technical user). No participants dismissed the environmental impact information as irrelevant; however, they varied in the intensity of concern and perceived ability to act.

The Sustainability Paradox and Cognitive Tension

This theme addresses RQ3 and RQ5 by showing that participants weighed environmental concerns against the perceived usefulness of LLMs, especially when distinguishing between high-value and low-value uses. High-value tasks, such as medical research and complex problem-solving, were seen as acceptable trade-offs, while low-value uses led to moral discomfort. P9 emphasized, "If you're looking for... a scientific improvement, or, let's say, for cancer research... to help out mankind in order to save lives. That's where the benefit comes" (P9, technical user). P12 drew a line around trivial uses: "I use it for some really silly stuff sometimes, and I don't think that's necessarily worth it... [I will] maybe not ask it silly questions outside of work" (P12, non-technical user). This cognitive tension remained largely unresolved, with several participants framing LLM use as an unavoidable feature of modern digital life. As P2 noted, "It is... definitely concerning but I don't see it going away... I doubt it will go anywhere" (P2).

Responsibility Assigned Upwards

This theme addresses RQ4 by showing that participants generally assigned responsibility for environmental mitigation to AI developers, platform providers, and data center operators rather than primarily to individual users. P6 captured this view: "I think the largest party who should be bearing some responsibility for these things [it is] the providers... the Googles and the ChatGPT... because they are enabling all of this... because of the scale of use" (P6, technical user). Participants also described uncertainty about what sustainable LLM use would look like in practice. Several participants noted that they lacked clear information about the environmental impact of specific tools, prompts, or usage choices. P15 stated, "I would have no idea how to check that off top [of] my head. I would assume that I would either [Google] it or ask the LLM what that looks like. Thus, I'm creating the problem" (P15, technical user). This uncertainty was more pronounced in professional settings where tool choice was determined by organizational policy. P5 explained, "From a work standpoint, we have to use Co-Pilot, so we can't change" (P5, technical user).

Conditional and Context-Dependent Intentions

This theme addresses RQ5 by showing that participants' willingness to adopt sustainable LLM behaviors depended on whether those behaviors were convenient, visible, and easy to integrate into existing use patterns. While absolute commitments to cease LLM use were rare, participants articulated conditional intentions to adopt sustainable behaviors, sensitive to the friction involved. A common intention was to reduce frivolous queries. P11 noted, "I'll try to limit the inconsequential... questions... because it's just as easy to open up, you know, Chrome on my phone and type it into the search bar" (P11, non-technical user).

However, when sustainable actions required compromises in speed or productivity, intention diminished. P10 expressed, “I’m still going to use it as much as I do, just because it does help in everyday activities... Because I’m [going to] use it anyway, if it helps me be more productive” (P10, non-technical user). Several participants described greater willingness to change their behavior if environmental information or lower-impact options were made visible within the tools themselves. P9 mentioned, “If a company can show that... with this particular query prompt, you utilize X amount of water... that would influence [if] I want to be this deep, or how to minimize my query so it will be more efficient” (P9, technical user).

Discussion

The five themes identified in this study suggest that users are not indifferent to the environmental impacts of LLMs, but the pathway from awareness to sustainable behavioral intention is constrained by cognitive distance, limited perceived behavioral control, and diffused responsibility. This section situates the findings within the existing literature and discusses their theoretical and practical implications. Participants’ abstract baseline awareness is consistent with Dungo et al. (2025), who found that users may recognize generative AI’s environmental footprint in general terms while lacking specific understanding, particularly regarding water consumption. The present findings add that awareness was shaped not only by the availability of information but also by how that information was framed. Relatable comparisons, such as household energy or long-term water use, helped participants move from abstract recognition to more concrete understanding. This suggests that the awareness gap identified in prior Green IT and sustainability research may be partly a framing deficit rather than only an information deficit (Ashiq et al., 2019; Dungo et al., 2025; Esfahani et al., 2015; Si et al., 2022). Participants’ reactions also varied by technical background. Non-technical users more often expressed surprise or concern, while technical users tended to view the information as confirmatory. However, concern did not consistently translate into clear behavioral intention. This pattern aligns with TPB research suggesting that attitudes alone are insufficient for intention formation when users lack perceived control or feasible action pathways (Bosnjak et al., 2020; Wei et al., 2025; Zhao et al., 2025). The findings also show that participants weighed environmental costs against perceived utility. This individual-level tension mirrors the broader sustainability paradox described in prior LLM research, where AI may create environmental burdens while also supporting efficiency and problem-solving (Kocak et al., 2025; Ren et al., 2024). In this study, participants were more willing to justify LLM use for high-value tasks, such as professional work, research, or complex problem-solving, while expressing more discomfort with casual or low-value uses. Responsibility attribution emerged as a central barrier to sustainable intention formation. Participants generally assigned responsibility to AI providers, platform developers, and data center operators rather than to individual users. From a TPB perspective, this suggests that perceived behavioral control may be limited in opaque digital systems where users lack feedback, alternatives, or clear metrics for evaluating environmental cost (Bosnjak et al., 2020). From a VBN perspective, responsibility diffusion may weaken the pathway from awareness of consequences to personal norm activation, even when environmental concern is present (Stern et al., 1999). Together, these findings support the usefulness of cognitive footprint as an exploratory concept for examining how users become aware of, interpret, justify, and respond to the environmental impacts of AI systems.

Theoretical Implications

These findings suggest that system opacity may be an important consideration when TPB is applied to AI sustainability contexts. In opaque AI systems, perceived behavioral control may depend less on users' individual skills and more on the availability of transparent information and actionable choices. Responsibility diffusion may similarly complicate the VBN pathway from awareness of consequences to personal norm activation. This study, therefore, proposes cognitive footprint as an exploratory concept describing the psychological processes—including awareness, interpretation, emotional response, cost-benefit reasoning, responsibility attribution, and perceived behavioral control—through which users understand and respond to the environmental impacts of technology. The five themes offer preliminary dimensions that future research may examine across different populations and technological contexts.

Practical Implications

Findings suggest that sustainable AI adoption requires both greater transparency and reduced behavioral friction. The consistent finding that impacts remain abstract until quantified through relatable comparisons indicates that the format of transparency matters: providers should translate environmental data into personally meaningful terms rather than raw technical units. Embedding sustainability features directly into AI interfaces—such as real-time environmental indicators, lower-impact model options, and configurable sustainability preferences—may enhance PBC more effectively than external educational campaigns, aligning with the Green IS approach described by Esfahani et al. (2015). Organizations can support sustainable practices by establishing efficient guidelines for AI use and incorporating environmental criteria into procurement decisions. At the industry level, standardized environmental reporting for AI systems, comparable to energy efficiency labels on appliances, could reduce the cognitive abstraction identified in this study.

Limitations

This study involved fifteen participants and was not designed for statistical generalization. All data are based on self-reported attitudes and stated intentions rather than observed behaviors. The study design introduced environmental impact information during the interview and immediately elicited reactions, which may have produced demand characteristics that encouraged socially desirable responses; the long-term durability of attitudinal shifts was not assessed. Additionally, the researcher served as the sole instrument for data collection and analysis. While a structured codebook and systematic procedures enhanced rigor, the absence of a secondary coder introduces the possibility of interpretive bias. The researcher's familiarity with the environmental impacts of AI may have influenced probing decisions and pattern recognition during coding. Finally, recruitment through professional networks may have limited the demographic diversity of participants.

Future Research

Future research should examine whether cognitive footprint can be operationalized and validated as a measurable construct across larger and more diverse samples. Quantitative studies could test whether the dimensions identified in this study, including awareness, interpretation, emotional response, cost-benefit reasoning, responsibility attribution, and perceived behavioral control, form a reliable scale for assessing user responses to the environmental impacts of AI systems. Longitudinal research is also needed to determine whether increased awareness leads to sustained behavioral change over time. Experimental studies testing interface-level interventions, such as real-time environmental indicators, lower-impact model options, or sustainability prompts, could provide stronger evidence regarding how design features influence user decision-making. Future studies in organizational settings could also clarify how workplace policies, tool availability, and institutional constraints shape sustainable AI usage. Finally, comparative research across different AI technologies could determine whether the patterns identified here are specific to LLMs or reflect broader dynamics of user-level sustainability cognition in opaque digital systems.

Conclusion

This study examined how users cognitively process the environmental impacts of large language models. Although quantified information increased awareness and often elicited concern, sustainable behavioral intentions were constrained by low perceived behavioral control and diffused responsibility. Users did not dismiss environmental impacts; rather, they struggled to identify meaningful avenues for action within opaque technological systems. The sustainability challenge of generative AI is therefore not solely technical but cognitive. Addressing environmental impact requires not only infrastructure-level optimization but also system designs that enhance transparency, reduce cognitive distance, and increase perceived agency.

Acknowledgements

The author used generative AI tools to support manuscript editing, organization, and revision. The author reviewed, verified, and approved all final content.

References

- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179–211. [https://doi.org/10.1016/0749-5978\(91\)90020-T](https://doi.org/10.1016/0749-5978(91)90020-T)
- Al Mamun, A., Hayat, N., Mohiuddin, M., Salameh, A. A., Ali, M. H., & Zainol, N. R. (2022). Modelling the significance of value-belief-norm theory in predicting workplace energy conservation behaviour. *Frontiers in Energy Research*, 10. <https://doi.org/10.3389/fenrg.2022.940595>
- Ashiq, S., Chrysoulas, C., & Banissi, E. (2019). Conceptualising green awareness as moderator in technology acceptance model for green IS/IT. In *2019 International Conference on Innovative Computing (ICIC)* (pp. 1–6). IEEE. <https://doi.org/10.1109/ICIC48496.2019.8966710>
- Bosnjak, M., Ajzen, I., & Schmidt, P. (2020). The theory of planned behavior: Selected recent advances and applications. *Europe's Journal of Psychology*, 16(3), 352–356. <https://doi.org/10.5964/ejop.v16i3.3107>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Cheung, K. S., Kaul, M., Jahangirova, G., Mousavi, M. R., & Zie, E. (2025). *Comparative analysis of carbon footprint in manual vs. LLM-assisted code development*. ACM. <https://doi.org/10.1145/3711919.3728678>
- Doudkin, A., Pataranutaporn, P., & Maes, P. (2025). *AI persuading AI vs AI persuading Humans: LLMs' Differential Effectiveness in Promoting Pro-Environmental Behavior* (No. arXiv:2503.02067). arXiv. <https://doi.org/10.48550/arXiv.2503.02067>
- Dungo, C. A. B., Beltran, Z. L. E., Declaro, B. C., Dela-Cruz, J. J. C., & Viray, R. U. (2025). Students' level of awareness on the environmental implications of generative AI. *Journal of Education in Science, Environment and Health*, 11(2), Article 2. <https://doi.org/10.55549/jeseh.777>

- Esfahani, M., Nilashi, M., Rahman, A., Ghapanchi, A., & Zakaria, N. (2015). Psychological factors influencing the managers' intention to adopt green IS: A review-based comprehensive framework and ranking the factors. *International Journal of Strategic Decision Sciences*, 6(2), 28–56. <https://doi.org/10.4018/ijds.2015040103>
- Guest, G., Bunce, A., & Johnson, L. (2006). How many interviews are enough? An experiment with data saturation and variability. *Field Methods*, 18(1), 59–82. <https://doi.org/10.1177/1525822X05279903>
- Jegham, N., Abdelatti, M., Elmoubarki, L., & Hendawi, A. (2025). How hungry is AI? Benchmarking energy, water, and carbon footprint of LLM inference [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.2505.09598>
- Jiang, P., Sonne, C., Li, W., You, F., & You, S. (2024). Preventing the immense increase in the life-cycle energy and carbon footprints of LLM-powered intelligent chatbots. *Engineering*, 40, 202–210. <https://doi.org/10.1016/j.eng.2024.04.002>
- Kallio, H., Pietilä, A. M., Johnson, M., & Kangasniemi, M. (2016). Systematic methodological review: Developing a framework for a qualitative semi-structured interview guide. *Journal of Advanced Nursing*, 72(12), 2954–2965. <https://doi.org/10.1111/jan.13031>
- Kocak, B., Ponsiglione, A., Romeo, V., Ugga, L., Huisman, M., & Cuocolo, R. (2025). Radiology AI and sustainability paradox: Environmental, economic, and social dimensions. *Insights into Imaging*, 16(1), 1–18. <https://doi.org/10.1186/s13244-025-01962-2>
- Kumar, V., Kaur, A., Cao, D., Sindhvani, R., Mathiyazhagan, K., & Lin, B. (2025). Exploring sustainable energy consumption practices: An extended environmental value-belief-norm framework using SEM analysis. *Energy Economics*, 145, 108408. <https://doi.org/10.1016/j.eneco.2025.108408>
- Liu, V., & Yin, Y. (2024). Green AI: Exploring carbon footprints, mitigation strategies, and trade-offs in large language model training. *Discover Artificial Intelligence*, 4(1), 1–12. <https://doi.org/10.1007/s44163-024-00149-w>
- Maliakel, P. J., Ilager, S., & Brandic, I. (2025). *Investigating energy efficiency and performance trade-offs in LLM inference across tasks and DVFS settings* (Version 1) [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.2501.08219>
- Morrison, J., Na, C., Fernandez, J., Dettmers, T., Strubell, E., & Dodge, J. (2025). *Holistically evaluating the environmental impact of creating language models* (Version 1) [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.2503.05804>
- Paris, M. (2025, April 12). ChatGPT hits 1 billion users, OpenAI CEO says, doubled in weeks. *Forbes*. <https://www.forbes.com/sites/martineparis/2025/04/12/chatgpt-hits-1-billion-users-openai-ceo-says-doubled-in-weeks/>
- Pimenow, S., Pimenowa, O., & Prus, P. (2024). Challenges of artificial intelligence development in the context of energy consumption and impact on climate change. *Energies*, 17(23), 5965. <https://doi.org/10.3390/en17235965>

- Ren, S., Tomlinson, B., Black, R. W., & Torrance, A. W. (2024). Reconciling the contrasting narratives on the environmental impact of large language models. *Scientific Reports*, 14(1), 1–8. <https://doi.org/10.1038/s41598-024-76682-6>
- Sánchez-Mompó, A., Mavromatis, I., Li, P., Katsaros, K., & Khan, A. (2025). Green MLOps to green GenOps: An empirical study of energy consumption in discriminative and generative AI operations. *Information*, 16(4), 281. <https://doi.org/10.3390/info16040281>
- Sanz, R. (2025, April 12). Google's Gemini AI app hits 400M users, 'doubled in weeks' says CEO Sundar Pichai. *Yahoo Finance*. <https://finance.yahoo.com/news/googles-gemini-ai-app-400m-174500173.html>
- Shang, D., Wu, W., & Schroeder, D. (2023). Exploring determinants of the green smart technology product adoption from a sustainability adapted value-belief-norm perspective. *Journal of Retailing and Consumer Services*, 70, 103169. <https://doi.org/10.1016/j.jretconser.2022.103169>
- Si, W., Jiang, C., & Meng, L. (2022). The relationship between environmental awareness, habitat quality, and community residents' pro-environmental behavior—Mediated effects model analysis based on social capital. *International Journal of Environmental Research and Public Health*, 19(20), 13253. <https://doi.org/10.3390/ijerph192013253>
- Stern, P. C., Dietz, T., Abel, T., Guagnano, G. A., & Kalof, L. (1999). A value-belief-norm theory of support for social movements: The case of environmentalism. *Human Ecology Review*, 6(2), 81–97.
- Wang, H., Hua, W., Peng, J., & Hu, M. (2025). Public sentiment analysis of data center energy consumption using social media data and large language models. *Energy and Buildings*, 341, 115802. <https://doi.org/10.1016/j.enbuild.2025.115802>
- Wei, W., Prasetyo, Y. T., Belmonte, Z. J. A., Cahigas, M. M. L., Nadlifatin, R., & Gumasing, Ma. J. J. (2025). Applying the technology acceptance model – Theory of planned behavior (TAM-TPB) model to study the acceptance of building information modeling (BIM) in green building in China. *Acta Psychologica*, 254, 104790. <https://doi.org/10.1016/j.actpsy.2025.104790>
- Wu, Y., Hua, I., & Ding, Y. (2025). *Unveiling environmental impacts of large language model serving: A functional unit view* (Version 1) [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.2502.11256>
- Zhao, L., Chen, T., & Huang, M. (2025). Predicting consumers' green purchasing intentions using the theory of planned behavior: Evidence from China. *Sustainability*, 17(4), 1894. <https://doi.org/10.3390/su17041894>