# THE COGNITIVE FOOTPRINT OF LARGE LANGUAGE MODELS: USER AWARENESS AND SUSTAINABLE INTENTIONS

by

MICHAEL ANTHONY GRECO

B.S., University of Colorado Global, 2016

M.I.T.M., University of Colorado Global, 2022

A Research Paper Submitted to the School of Computing Faculty of

Middle Georgia State University in

Partial Fulfillment for the Requirements for the Degree

DOCTOR OF SCIENCE IN INFORMATION TECHNOLOGY

MACON, GEORGIA

2026

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

**Michael Greco,** *Middle Georgia State University,* Michael.Greco@mga.edu

**Abstract**

The rapid adoption of Large Language Models (LLMs) has raised concerns about their environmental impacts, 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 internalize and respond to a technology's environmental impacts, including awareness, interpretation, emotional response, cost-benefit reasoning, responsibility attribution, and perceived behavioral control (PBC). 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 core themes: (a) environmental impacts remained abstract until quantified through relatable comparisons, (b) emotional responses varied by technical background, (c) participants engaged in cost-benefit reasoning reflecting an individual-level sustainability paradox, (d) responsibility was externalized toward technology providers, and (e) sustainable intentions were conditional on low behavioral friction. Findings indicate that the awareness gap identified in prior research is not solely an information deficit but a framing deficit, and that PBC in opaque digital systems is structurally suppressed by a lack of system transparency and user-facing choice architecture. This study extends TPB and VBN into digitally mediated sustainability contexts, introduces the cognitive footprint as a measurable construct, and highlights the role of interface design in bridging the gap between environmental concern and sustainable AI usage.

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

## Introduction

In the five years since the release of OpenAI's ChatGPT-3, Large Language Models (LLMs) have become widely integrated into both professional and personal aspects of digital life. Major corporations, including Google, Microsoft, and X, have launched competing models to capture this market. As of 2025, OpenAI's ChatGPT boasts 500 million weekly active users, while Google's Gemini has 400 million monthly active users (Paris, 2025; Sanz, 2025). This surge in popularity is primarily due to their versatility in handling a wide range of tasks including 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 rapid adoption of this technology has raised important questions regarding the environmental implications of these powerful language models (Jiang et al., 2024; Kocak et al., 2025; Pimenow et al., 2024). As society and the global economy become increasingly dependent on technologies such as AI and LLMs, the environmental footprint of these technologies has become a significant concern (Jiang et al., 2024; Maliakel et al., 2025; Pimenow et al., 2024). Concepts such as "Green IT" and "Sustainable Computing" have become burgeoning areas of research as scholars seek to address the escalating energy and resource demands of the digital age (Ashiq et al., 2019; Esfahani et al., 2015).

The environmental impact of LLMs stems from energy-intensive computational processes and water usage during both training and inference phases (Jegham et al., 2025; Morrison et al., 2025). A 2022 study analyzing the carbon intensity of AI found that training an LLM to just 13% completion required 103.5 MWh of energy (Dodge et al., 2022). More recent estimates indicate that training a full LLM can consume approximately 1,287 megawatt-hours (MWh) of electricity and emit over 550 tons of $CO_2$ equivalents

(Morrison et al., 2025). To put this in perspective, the energy used to train a single model is roughly equivalent to the annual energy consumption of hundreds of homes (Morrison et al., 2025).

Beyond energy, these models have a substantial water footprint. Data centers require significant amounts of water to cool the specialized hardware used for running large LLMs (Jegham et al., 2025; Kocak et al., 2025). Estimates suggest that the water consumed during the training phase of several LLMs amounted to 2.769 million liters, roughly the same amount of water one person in the United States would use over 24 years (Morrison et al., 2025; U.S. Environmental Protection Agency, n.d.). Furthermore, the inference phase, in which the model generates responses to user queries, accounts for the majority of energy consumption across the model's lifecycle (Jegham et al., 2025). With over 700 million queries submitted daily, the operational energy usage is equivalent to the annual consumption of about 35,000 United States homes (Jegham et al., 2025).

Despite the documentation of this technical carbon footprint, a critical gap remains regarding the "cognitive footprint," or the extent to which these environmental costs are recognized and internalized by users. While the technical aspects of LLM efficiency have been explored, the role of the end user has not been adequately addressed (Cheung et al., 2025; Doudkin et al., 2025; Dungo et al., 2025; Jiang et al., 2024; Liu & Yin, 2024). Limited evidence suggests that most users are unaware of the energy and water consumption associated with their everyday interactions with LLMs (Dungo et al., 2025). This "awareness gap" poses a significant barrier to promoting sustainable technology use and advancing Green IT initiatives, as collective user behavior strongly influences the technology's overall environmental footprint (Ashiq et al., 2019; Esfahani et al., 2015; Wu et al., 2025).

To understand how users might adopt more sustainable behaviors, this study draws on behavioral science theories, specifically the Theory of Planned Behavior (TPB) and the Value-Belief-Norm (VBN) Theory. These frameworks suggest that awareness alone is insufficient; factors such as PBC and value-driven norms determine whether awareness translates into sustainable intentions and actions (Bosnjak et al., 2020; Stern et al., 1999). Currently, it remains unclear how awareness of these environmental impacts affects users' intentions to adopt more sustainable practices when using LLMs.

This qualitative study explores how users understand and interpret the environmental impacts of LLMs and how this awareness influences their attitudes, PBC, personal values, and intentions to engage in sustainable usage behaviors. Specifically, the research aims to understand how users comprehend information regarding the energy consumption, carbon emissions, and water usage associated with LLMs. The study also seeks to examine how their responses to this information affect their willingness to adopt more sustainable practices, such as reducing unnecessary queries or selecting energy-efficient 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 their environmental sustainability (Kocak et al., 2025; Jegham et al., 2025; Morrison et al., 2025; Wang et al., 2025; Wu et al., 2025). The development, training, and operation of these models require substantial energy and water and generate considerable electronic waste (Kocak et al., 2025; Jegham et al., 2025; Morrison 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 chapter 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

While LLMs represent a new era in computing, their widespread adoption has been accompanied by a significant environmental cost (Kocak et al., 2025; Jegham et al., 2025; Morrison et al., 2025; Wang et al., 2025; Wu et al., 2025). 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 (Jegham et al., 2025; Morrison et al., 2025).

### Energy Consumption in the AI Lifecycle

The energy demands of LLMs can be categorized into two phases: training and inference (Jegham et al., 2025; Pimenow et al., 2024).

#### *Training Phase*

The training phase of a large language model (LLM) is a lengthy and highly energy-intensive process that demands significant computational resources (Pimenow et al., 2024). During this phase, vast datasets are processed, resulting in a significant increase in CO2 emissions (Pimenow et al., 2024). According to a study by Morrison et al. (2025), the training of an LLM was estimated to consume approximately 1,287 MWh of electricity and emit over 550 tons of CO2 equivalent. A report from Meta's AI research team noted that 20% of an LLM's energy consumption occurs during the training phase (Sánchez-Mompó et al., 2025). Furthermore, to gain a comprehensive understanding of the environmental impact of LLMs, it is crucial to consider the "embodied carbon" associated with the manufacturing processes involved in creating the necessary hardware, which can account for 20% to 40% of the model's total emissions over its lifetime (Morrison et al., 2025).

#### *Inference (Usage) Phase*

The inference phase, where an LLM generates responses to user queries, accounts for most of the model's energy consumption throughout its lifecycle (Jegham et al., 2025). Each query processed by a model like

OpenAI's GPT-4o consumes approximately 0.43 watt-hours (Jegham et al., 2025). With over 700 million queries submitted daily, the total energy usage during this phase is equivalent to the annual energy consumption of about 35,000 U.S. homes (Jegham et al., 2025). Notably, the ongoing operation of a large language model (LLM) can account for up to 90% of its total energy consumption over its lifetime (Jegham et al., 2025).

However, the energy consumption of LLMs during the inference phase is highly malleable (Sánchez-Mompó et al., 2025). Larger and more powerful models, such as Mistral-7B and Falcon-7B, consume four to six times as much energy as smaller models like GPT-2 (Maliakel et al., 2025). Additionally, system-level settings affect energy efficiency; one study reported that increasing a GPU's clock speed reduced energy consumption by 30%-50%, primarily due to faster query completion times (Maliakel et al., 2025). This variability suggests numerous opportunities for reduction, including choices made by end users themselves (Sánchez-Mompó et al., 2025).

### Carbon Emissions and Greenhouse Gases

The vast energy consumption of large language models (LLMs) directly contributes to significant carbon emissions, as determined by multiplying total energy consumption by the local power grid's carbon intensity (Wu et al., 2025). A 2025 study on software development found that while LLM-assisted code generation was faster, it resulted in a carbon footprint 19-43 times higher than manual coding across all scenarios (Cheung et al., 2025). Another study focusing on the training phase estimated that training a single LLM with 213 million parameters can produce around 300 metric tons of CO2 equivalents, which is roughly equal to the emissions from 125 round-trip flights between Beijing and New York City (Jiang et al., 2024).

### Water Consumption

Beyond energy use, water consumption is a significant yet often overlooked aspect of LLMs' environmental impact (Jegham et al., 2025). Data centers that host these models require substantial water to cool their servers (Jegham et al., 2025; Jiang et al., 2024). During the training phase of GPT-3 alone, OpenAI is estimated to have used 5.4 million liters of water (Jiang et al., 2024).

Another study estimates that the total water footprint of developing a series of LLMs could reach 2.769 million liters, equivalent to the average amount of water consumed by one person in the United States over 24.5 years (Morrison et al., 2025). During the inference phase, the water footprint can reach up to 500 mL per query, excluding water used for energy generation (Jegham et al., 2025). As global demand for AI continues to rise, this concern is becoming increasingly urgent (Kocak et al., 2025). One study predicts that by 2027, global demand for AI could result in annual water consumption ranging from 4.2 to 6.6 billion cubic meters (Kocak et al., 2025).

### The Debate on Net Environmental Impact: The Sustainability Paradox

Although the data on energy and water consumption offer a troubling insight, the overall environmental impact of large language models (LLMs) is more complex and nuanced, as the technology is both a contributor to environmental issues and a potential solution (Kocak et al., 2025). This duality is often described as the 'sustainability paradox' (Kocak et al., 2025).

Several studies support the net positive narrative. For instance, a study conducted by Ren et al. (2024) examined the energy required to write a 500-word essay. The findings revealed that a human consumes 0.85

kWh for this task, while the Llama-3-70B LLM model consumes only 0.02 kWh. This indicates an efficiency ratio of 44:1 in favor of the LLM (Ren et al., 2024). When lighter-weight LLMs are considered, this ratio can improve even further, reaching as high as 3,500 to 1 (Ren et al., 2024).

However, these optimistic findings stand in direct tension with the immense and escalating resource consumption required to develop and operate these same models (Kocak et al., 2025; Ren et al., 2024). The trend toward building ever-larger LLMs threatens to offset any such efficiency gains, creating significant uncertainty around the technology's ultimate environmental outcome (Kocak et al., 2025; Ren et al., 2024). This uncertainty underscores the pressing need for further research, transparent reporting of carbon footprints, and greater user awareness to guide the technology toward sustainability.

**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; it also requires a combination of a person's attitude, knowledge, and willingness to act (Ashiq et al., 2019).

**Theoretical Frameworks for Sustainable Intention Formation**

To understand the factors that drive sustainable intentions and behaviors, it is important to reference established psychological theories. Frameworks like the Theory of Planned Behavior (TPB) and the Value-Belief-Norm (VBN) Theory provide models for examining how individual values, beliefs, and social contexts influence pro-environmental actions (Stern et al., 1999).

**The Theory of Planned Behavior (TPB)**

The Theory of Planned Behavior (TPB) posits that attitudes, subjective norms, and PBC influence intentions (Bosnjak et al., 2020). In sustainability contexts, PBC is often a strong predictor of action (Bosnjak et al., 2020; Wei et al., 2025; Zhao et al., 2025). For instance, a study involving 354 consumers found that both attitude toward green consumption and PBC predicted intention to engage in green consumption, with PBC having the strongest effect (Zhao et al., 2025). Similarly, in a study of 588 professionals on the adoption of Building Information Modeling (BIM) in green building projects, PBC was identified as having the greatest influence on its use (Wei et al., 2025). This suggests that users' intention to act is closely related to their perceived ability to make a difference (Wei 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).

Empirical studies have validated this model. For instance, a study involving Indian consumers found that biospheric and altruistic values positively predicted the development of an ecological worldview, whereas egoistic values negatively affected it (Kumar et al., 2025). Similarly, research on employees in China demonstrated that the VBN theory can effectively predict workplace energy conservation efforts, in which biospheric values emerged as a significant driver of pro-environmental beliefs (Al Mamun et al., 2022). This suggests that a user's underlying value orientation can significantly influence their response to learning about technology's environmental impact (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. No in-depth study has examined how users understand the environmental impacts of LLMs or how this awareness shapes their attitudes, values, 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 (18 years or older) who were familiar with Large Language Models such as ChatGPT, Gemini, Copilot, or Claude. A purposive sampling strategy was used to recruit 10–15 participants, as it enabled the intentional selection of individuals with direct experience relevant to the research questions. This sample size was considered sufficient to achieve thematic saturation in a qualitative descriptive study, as prior research suggests that saturation in semi-structured interview studies with relatively homogeneous participants often occurs after approximately 12 interviews (Guest et al., 2006). Thematic saturation was observed during the later stages of data collection, as no new codes or themes emerged after approximately the twelfth interview. Subsequent interviews reinforced existing patterns rather than introducing novel insights, indicating that sufficient depth had been achieved. The sampling strategy sought to ensure diverse usage levels and perspectives while maintaining a shared baseline of LLM familiarity.

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 (i.e., adults with prior LLM experience) were contacted privately with a brief description of the study and an invitation to participate. Participation was entirely voluntary. To minimize potential power differentials, individuals under the researcher's direct supervision were not recruited.

## 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, 2000). 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.

Crucially, the interview included a scripted Standardized Information Prompt summarizing empirically documented findings on LLM energy consumption and water usage. The inclusion of a standardized stimulus ensured that all participants were exposed to 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. The semi-structured format ensured consistency while allowing the researcher to use follow-up probes to explore specific responses in depth.

**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 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.

All transcripts were reviewed manually by the researcher. Coding was conducted by hand, with codes recorded and organized in a structured Microsoft Excel codebook. The Excel file was used to systematically document code names, operational definitions, exemplar quotations, and associated transcript references. This facilitated transparent tracking of code development and revisions throughout the analytic process.

A hybrid coding strategy was employed. Deductive codes were developed a priori from the study's theoretical frameworks, including constructs from the Theory of Planned Behavior and Value-Belief-Norm theory. Simultaneously, inductive codes emerging directly from the data were incorporated to capture unanticipated patterns and meanings. The codebook was iteratively refined as analysis progressed, with analytic decisions documented to maintain consistency and auditability. Direct quotations were retained and mapped to corresponding codes and themes to support analytic credibility.

**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, a structured interview protocol, a standardized informational prompt, and systematic coding procedures were employed to maintain consistency and transparency throughout the research process.

**Trustworthiness**

To ensure the rigor of this qualitative study, trustworthiness was established using the criteria of credibility, dependability, confirmability, and transferability.

Credibility was maintained by employing a standardized informational prompt, ensuring all participants received identical stimuli before reflective questioning. The interviews were recorded and transcribed word-for-word. The semi-structured protocol permitted probing to clarify participant responses. Additionally, aligning the interview protocol with well-established theoretical frameworks, such as the Theory of Planned Behavior and the Value-Belief-Norm Theory, enhanced the interpretive validity.

Dependability was enhanced through a structured coding process and an audit trail. All coding decisions were documented in a Microsoft Excel codebook, including code definitions, revisions, and associated excerpts. The iterative refinement of codes and themes followed a consistent analytic process aligned with Braun and Clarke's (2006) six-phase thematic analysis.

Confirmability was supported by maintaining a clear link between the data and the findings. Direct participant quotations were systematically included to illustrate themes, ensuring that interpretations were grounded in the data rather than in the researcher's assumptions. The structured codebook and documentation of analytic decisions further enhanced transparency.

Transferability was addressed by including participants with diverse professional backgrounds and varying levels of familiarity with large language models. Although the findings are not intended for statistical generalization, the detailed descriptions of participant perspectives and contexts enable readers to assess their applicability to similar populations and settings.

## 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 chapter provides a detailed description of participant demographics and sampling procedures. This chapter 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 and are presented in the sections that follow. These themes are aligned with and address the study's five research questions.

**Table 1**

*Summary of Themes, Descriptions, and Exemplar Quotes*

| Theme | Description | Exemplar Quote |
|---|---|---|
| Environmental Impact Remains Abstract Until Quantified | Users lacked a concrete understanding until information was framed using relatable comparisons | "That's an astounding amount… but I understand that." |
| Uneven Distribution of Surprise and Concern | Non-technical users showed stronger emotional reactions, while technical users viewed information as confirmatory | "That blows my mind… that just is an alarming amount of waste." |
| Sustainability Paradox and Cognitive Tension | Users weighed environmental costs against utility, creating internal conflict | "If it helps mankind… that's where the benefit comes." |
| Responsibility Assigned Upwards | Responsibility for sustainability was attributed to providers rather than users | "The providers… should be bearing responsibility." |
| Conditional and Context-Dependent Intentions | Users expressed willingness to act only when behavioral friction was low | "I'll try to limit inconsequential questions…" |

### Environmental Impact Remains Abstract Until Quantified

This theme addresses Research Question 1, which aimed to describe users' baseline awareness of environmental impacts. Participants consistently described their initial understanding of environmental costs as limited, abstract, or nonexistent prior to receiving the informational prompt. While a general assumption existed 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 with lower technical literacy 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).

Crucially, these reactions revealed a disconnect between concern and agency. Even when participants experienced unease, that unease did not consistently translate into a clear sense of feasible action. 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 reflects the complex negotiation participants undertook as they weighed the utility of LLMs against their environmental costs. 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

Addressing Research Question 4, participants predominantly placed responsibility for environmental mitigation on AI developers and data center operators rather than viewing it as a personal obligation. P6 captured this: "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). Perceived behavioral control (PBC) was also particularly low. Participants identified barriers related to a lack of transparency and choice architecture. 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 limitation was even more pronounced in professional settings, where tool choice is determined by 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

The final theme addresses Research Question 5. 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).

Notably, participants expressed willingness to act if the technology could bridge the gap through Green IS features. 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). These insights suggest that design interventions making sustainable choices clear and low-cost could transform existing concerns into practical, repeatable behaviors.

## Discussion

The five themes identified in this study reveal a consistent pattern: users are not indifferent to the environmental impacts of LLMs, but the pathway from awareness to sustainable behavioral intention is constrained by cognitive distance, low perceived behavioral control, and diffused responsibility. This chapter situates these findings within the existing literature, discusses theoretical and practical implications, and identifies limitations and directions for future research.

This study contributes to Information Systems research by extending the Theory of Planned Behavior into opaque, digitally mediated environments, refining the Value-Belief-Norm framework through the identification of responsibility diffusion, and introducing the cognitive footprint as a novel construct for understanding how users cognitively process the environmental impacts of AI systems.

The finding that participants possessed only abstract awareness of environmental impacts prior to quantification is consistent with Dungo et al. (2025), who reported that college students demonstrated general awareness of generative AI's footprint but lacked specific understanding, particularly regarding water consumption. This study deepens this finding by revealing the mechanism through which awareness shifts: participants did not simply lack information but lacked relatable frames of reference. It was the use of household-level comparisons, such as equating water consumption to that of one person over 24 years, that catalyzed comprehension. This suggests that the awareness gap in prior research (Ashiq et al., 2019; Dungo et al., 2025; Esfahani et al., 2015) is not solely an information deficit but a framing deficit, extending Si et al.'s (2022) finding that environmental awareness predicts pro-environmental behavior by demonstrating that awareness operates along a continuum from abstract recognition to concrete understanding.

The divergent emotional responses, with non-technical users expressing surprise and technical users treating data as confirmatory, add nuance to TPB applications in sustainability research, which have generally treated attitude formation as a uniform process (Wei et al., 2025; Zhao et al., 2025). The present findings suggest that prior technical knowledge moderates both the intensity and character of attitudinal responses, providing qualitative evidence for the moderating role of green awareness proposed by Ashiq et al. (2019). However, even strong emotional responses did not consistently translate into behavioral intention, aligning with Bosnjak et al.'s (2020) observation that attitudes are necessary but insufficient for intention formation under TPB.

Participants' cost-benefit reasoning mirrors the sustainability paradox described at the systems level by Kocak et al. (2025) and Ren et al. (2024), but this study demonstrates that the same tension operates at the individual cognitive level. Unlike more straightforward environmental behaviors such as household energy conservation (Kumar et al., 2025) or workplace energy behavior (Al Mamun et al., 2022), LLM use involves a perceived tradeoff in which the environmentally costly behavior is also the productive, socially valued, and often professionally mandated behavior. This duality complicates the VBN causal chain from awareness to personal norm activation (Stern et al., 1999) because the same behavior that triggers environmental concern also triggers justification based on perceived utility.

Perhaps the most consequential finding is the consistent externalization of responsibility. Under TPB, PBC is recognized as a strong, often the strongest, predictor of behavioral intention (Bosnjak et al., 2020; Wei et al., 2025; Zhao et al., 2025). The present study confirms PBC's centrality but reveals that in the context of LLM use, PBC is structurally suppressed by the opacity of the technological system itself. Participants could not identify what a sustainable LLM behavior would look like, not because they lacked motivation, but because the system provides no feedback mechanisms, no visible alternatives, and no metrics for evaluating environmental cost. This represents a qualitatively different barrier than those in prior TPB applications, where PBC is typically constrained by skill deficits or resource limitations (Bosnjak et al., 2020). From a VBN perspective, the externalization of responsibility interrupts the pathway between awareness of consequences and personal norm activation. Stern et al. (1999) posited that awareness leads to the ascription of personal responsibility, which in turn activates norms to act. Here, participants attributed responsibility externally, suggesting that in digitally mediated contexts, locus of responsibility may function as a moderating variable that interrupts the VBN causal chain even when awareness and concern are present.

**Theoretical Implications**

This study expands the TPB into contexts of digitally mediated sustainability, which often feature unclear infrastructure. It demonstrates that PBC relies more on system transparency and the available choices than on individual skills or resources. Additionally, it refines the VBN framework by indicating that the diffusion of responsibility can disrupt the connection between awareness and the activation of personal norms in large-scale technological systems. Furthermore, this research introduces the concept of the cognitive footprint, which is defined as the set of psychological processes—including awareness, interpretation, emotional responses, cost-benefit reasoning, responsibility attribution, and perceived behavioral control—through which users understand and react to the environmental impacts of a technology. The five themes discussed provide a comprehensive view of this construct, offering a framework that future researchers can apply across various 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 employ longitudinal designs to assess whether increased awareness leads to sustained behavioral changes over time. Experimental studies testing specific interface-level interventions, such as real-time carbon dashboards or low-impact modes, would provide causal evidence for the design recommendations suggested here. Quantitative research should further develop the cognitive footprint construct, potentially through a Cognitive Footprint Awareness Scale incorporating the dimensions identified in this study. Studies situated in organizational contexts could clarify how workplace constraints and institutional policies shape sustainable AI usage. Finally, comparative research across different AI technologies would determine whether the patterns identified here are specific to LLMs or reflect broader dynamics of user-level sustainability cognition in opaque 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.

## References


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

Dodge, J., Prewitt, T., Combes, R. T. D., Odmark, E., Schwartz, R., Strubell, E., Luccioni, A. S., Smith, N. A., DeCario, N., & Buchanan, W. (2022). Measuring the carbon intensity of AI in cloud instances. *In 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)* (pp. 2026-2038). Association for Computing Machinery. https://doi.org/10.1145/3531146.3533234

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/ijsds.2015040103

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

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

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., Sindhwani, 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.

U.S. Environmental Protection Agency. (n.d.). *Statistics and facts*. Retrieved July 13, 2025, from https://www.epa.gov/watersense/statistics-and-facts

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