



Public perceptions of AI energy consumption: Awareness, responsibility attribution, and strategies

Michael Klesel^{1,2} · Uwe Messer³

Received: 24 September 2025 / Accepted: 18 May 2026
© The Author(s) 2026

Abstract

Generative artificial intelligence (GenAI) is widely used in the gig economy and electronic markets. However, its progress incurs substantial environmental costs due to its high and constantly increasing energy consumption. As inference—the process by which GenAI generates output—accounts for a large part of this demand, individual usage patterns are becoming a key factor in understanding the overall energy footprint and a crucial prerequisite for developing effective educational strategies and policies. Drawing from the construal-level theory of psychological distance, this study explores individuals' perceptions related to GenAI energy consumption. Based on a nationally representative sample of 1080 participants from the USA, we provide evidence that individuals are aware of GenAI-related energy consumption. Moreover, we show that, compared to household appliances, there is a diffusion of responsibility for GenAI's energy consumption, with more emphasis placed on policy-makers and developers than on users. We also explore how individuals intend to implement pro-environmental strategies and shed light on different avenues suggested by individuals to lower energy consumption when using GenAI. We conclude with implications for theory, practice, and policymaking.

Keywords (Generative) Artificial intelligence · Energy consumption · Green AI · Public perception · Psychological distance

JEL Classification O33

Introduction

Generative artificial intelligence (GenAI), particularly multimodal systems like ChatGPT, is widely used in the gig economy and is rapidly proliferating across diverse sectors (Shin et al., 2023; Humlum & Vestergaard, 2025), offering remarkable performance that can enhance human productivity (Noy & Zhang, 2023; Brynjolfsson et al., 2025) and even surpass human capabilities in specific domains (Bran et al., 2024; Trinh et al., 2024). However, this advancement incurs

substantial environmental costs, primarily through increased energy consumption and associated anthropogenic CO₂ emissions¹ (International Energy Agency, 2024; Bourzac, 2024; Ananthaswamy, 2023). For instance, training a large language model such as BLOOM (176-billion parameters) is estimated to generate up to 50.5 tonnes of CO₂ (Luccioni et al., 2023). Although training is energy-intensive, the inference stage, when the model is deployed for real-world applications, requires as much energy, if not more. The Electric Power Research Institute estimates that inference accounts for approximately 60% of AI's total energy consumption (Aljbour et al., 2024, p. 15). Even simple tasks like question-answer queries consume measurable energy (e.g., 0.001 Wh), with more complex operations like video generation requiring substantially more (0.477 Wh) (Luccioni et al., 2024a).

Responsible Editor: Lars Andraschko

✉ Michael Klesel
michael.klesel@fra-uas.de
Uwe Messer
uwe.messer@unibw.de

¹ Frankfurt University of Applied Sciences, Nibelungenplatz 1, 60318 Frankfurt, Deutschland

² hessian.AI, Landwehrstr. 50A, 64293 Darmstadt, Germany

³ Universität der Bundeswehr München, Werner-Heisenberg-Weg 39, 85577 Neubiberg, Germany

¹ Carbon emissions are a function of total energy demand and the carbon intensity of the underlying energy infrastructure (Luccioni & Hernandez-Garcia, 2023; Strubell et al., 2020). For the general public, electricity consumption provides a more tangible and actionable metric than carbon emissions, which is why this study focuses on energy consumption.

Large language models (LLMs), primarily built on the transformer architecture (Vaswani et al., 2017), have historically been advanced by scaling training compute, model parameters, and dataset sizes (*training-time scaling*) (Hoffmann et al., 2022). With trends towards diminishing returns and escalating resource demands indicating a saturation point for purely training-focused scaling, the research landscape is increasingly shifting towards optimizing inference-time performance. Techniques such as Chain-of-Thought prompting (Wei et al., 2022), for example, aim to improve how LLMs handle complex requests by generating explicit reasoning steps (*inference-time scaling*). While such approaches can further enhance output quality, they cause a substantial increase in computational costs for the inference stage (Brown et al., 2025).

The resource consumption of GenAI,² presents a challenge that can be approached from technical and socio-behavioral perspectives. From a technical perspective, the focus is on achieving efficiency gains designed to stagnate or even shrink AI's environmental footprint (Patterson et al., 2022). This implies an emphasis on hardware and data center optimization (Mankoff et al., 2007). However, the sufficiency of purely technical solutions in ensuring AI sustainability is debated. Increased efficiency can lead to rebound effects, where cost reduction or ease of use results in greater overall consumption and a net increase in resource demands (Luccioni et al., 2025; Varoquaux et al., 2025). Consequently, socio-technical research can serve as a valuable complementary strand to existing research in technical domains, thereby reducing GenAI's energy consumption on a large scale, particularly in light of the global adoption of GenAI by end users (Bick et al., 2026).

As inference increasingly drives AI's energy consumption, individual usage patterns will collectively exert a greater influence on global energy demand. User decisions, such as how often, how extensively, and for what tasks they employ GenAI (e.g., prompt complexity, duration, model choice), significantly impact global energy demand. For instance, using an LLM for an Internet search can consume considerably more energy than a conventional Google search (Aljbour et al., 2024). Similarly, selecting task-specific models over large, general-purpose ones can yield substantial energy savings (Luccioni et al., 2024b; Klesel & Messer, 2026). Every token processed and generated consumes energy (Jegham et al., 2025). This includes common politeness phrases, such

as “thank you” and “please,” which incur a measurable and cumulatively substantial energy cost.

Although extensive research has explored the nexus between individual attitudes, behaviors, and energy consumption in domains such as household appliances and transportation (Attari et al., 2010; Attari, 2014; Dietz, 2014; Andor et al., 2020; Allcott, 2011; Larrick & Soll, 2008), there is a lack of understanding of the role of the user in AI-related energy consumption. Despite a consensus in academia and industry on the pressing need to reduce AI's CO₂ emissions (Chen et al., 2025; Luccioni et al., 2024a), public perception of this issue remains largely unexamined. Key questions regarding public awareness, responsibility attribution, and knowledge of mitigation strategies lack empirical answers, hindering the efficacy of policy interventions aimed at fostering *green AI* (Schwartz et al., 2020).

This study makes the case that understanding public perception is essential to reveal knowledge gaps and ultimately to target those gaps from an end-user perspective and is guided by the following research question (RQ):

RQ: To what extent (a) is the public aware about GenAI's energy consumption, (b) does the public feel responsible for GenAI's energy consumption, and (c) does the public come up with strategies to reduce GenAI's energy consumption?

To theorize on the perceptions of GenAI-related energy demands, we use construal-level theory, the concept of psychological distance, and diffusion of responsibility (Bandura et al., 1996; Latane & Darley, 1968; Trope et al., 2007; Trope & Liberman, 2010) as a theoretical perspective that guides our conceptual model and which are used to interpret our results. Using data from a nationally representative sample of the US population ($N = 1080$), we provide empirical insights and discuss their implications for advancing both academic understanding and practical solutions.

This study makes several contributions. First, we contribute to an emerging discourse on the intersection of GenAI and sustainability research. Specifically, we emphasize the importance of a socio-technical perspective and reveal empirical insights about individuals' perceptions in this area. Using construal-level theory and the concept of psychological distance (Trope et al., 2007; Trope & Liberman, 2010) as a theoretical lens, we introduce these theories as meaningful lenses for better understanding differences in public perception of information technology. Second, we provide new insights into the diffusion of responsibility in the context of GenAI showing to what extent individuals distribute responsibility and contrast our results with responsibility attribution in the domestic context. This knowledge adds to current research on digital responsibility and puts individual perception into a better context. Third, we add to the behavioral IS

² Note that GenAI is a subclass of artificial intelligence (see e.g., Banh & Strobel, 2023). Although other AI categories, such as machine learning (e.g., decision trees) and deep learning (e.g., neural networks), also consume energy, this study focuses on GenAI due to its widespread adoption by end users. Thus, when referring to AI, we primarily consider generative AI.

research that focuses on sustainability by providing insights into individual perception in the context of energy consumption and GenAI. Due to the massive use of GenAI, our results inform a broad range of behavioral IS research that goes beyond perceptions to behavior related with GenAI.

Related work

Energy consumption in IS literature

Digital sustainability is an important field in Information Systems research, with three distinct outcomes: (1) *environmental sustainability*, (2) *social sustainability*, and (3) *economic sustainability* (Kotlarsky et al., 2023). Here, we focus on environmental sustainability, which encompasses two related concepts, namely green IT and green IS. *Green IT* is commonly defined as the practice of creating or using environmentally sustainable technology. In contrast, *Green IS* refers to the use of technology to achieve environmental objectives (for a more comprehensive discussion on green IS and green IT, see, for instance, Dedrick, 2010; Watson et al., 2010; vom Brocke et al., 2013). With the rise of Artificial Intelligence, the current literature also employs the notion of *Green AI* (Schwartz et al., 2020).

Within the field of green IS, green IT, and green AI, the role of energy consumption has been identified as one grand societal challenge (Ketter et al., 2016). For this reason, there is an increasing number of studies that seek to identify and evaluate new avenues to reduce energy consumption. Reviewing previous literature on energy consumption and energy-related use behavior, we find studies focusing on behavior, technology, or those that take a management perspective. With regard to use behavior, several studies have investigated the use of smart meters (Wendt et al., 2024; Dalén & Krämer, 2017; Sim et al., 2023) and web portals (Loock et al., 2013). In terms of technology, studies have been conducted on high-performance computing systems (Guo et al., 2023), enterprise systems (Hedwig et al., 2009), data cen-

ters (Fridgen et al., 2021), and carbon management systems (Zampou et al., 2022). Finally, there are also studies focusing on managerial issues, such as policy making (Watson et al., 2022) or the optimization of business process management (Hermann et al., 2024). Table 1 provides an overview of example studies on energy consumption.

Existing studies on energy consumption attest to the importance of this stream in IS research, which contributes to the overarching stream of sustainable IS. Various studies have already contributed to a better understanding of individual behavior, especially with regard to smart meters (Wendt et al., 2024; Dalén & Krämer, 2017; Sim et al., 2023), and have shown viable ways to enhance systems, such as high-performance computing systems (Guo et al., 2023), to reduce energy consumption. Besides behavioral and technological research, there are also studies that provide deeper insights into policy making (Watson et al., 2022) and the optimization of business process management (Hermann et al., 2024). At the same time, this review reveals that energy consumption related to AI has not been investigated thus far. Consequently, the IS literature lacks an understanding of the relationship between AI and energy consumption. Since organizations are increasingly challenged to reduce energy consumption, this knowledge gap is critical.

Energy consumption and AI

The impressive advancements of modern AI models come at a significant environmental and resource cost (Strubell et al., 2020; Luccioni et al., 2024b; Jegham et al., 2025; Li et al., 2025; Dauner & Socher, 2025). Initial research into AI's environmental footprint focused on the energy required to train deep learning models for natural language processing. These studies revealed that the carbon emissions from training a single model could be comparable to the entire lifecycle emissions of several cars (Strubell et al., 2020). While this initial focus on build-time (training) costs has been well-documented, subsequent work has broadened the scope to

Table 1 Example studies on energy consumption in IS literature

Focus	Context	References
Behavior	Smart meters	(Wendt et al., 2024; Dalén & Krämer, 2017; Sim et al., 2023)
	Web portals	(Loock et al., 2013)
Technology	High-performance computing systems	(Guo et al., 2023)
	Enterprise systems	(Hedwig et al., 2009)
	Data centers	(Fridgen et al., 2021)
	Carbon management systems	(Zampou et al., 2022)
Management	Policy	(Watson et al., 2022)
	Business processes	(Hermann et al., 2024)

include comprehensive lifecycle assessments of foundation models, which account for everything from manufacturing hardware to the energy used in training (Luccioni et al., 2023). In addition to its energy consumption, research has shown that AI also has a significant water footprint. This is primarily due to the freshwater required to cool the data centers used to train models like GPT-3 (Li et al., 2025; Jegham et al., 2025).

Historically, the focus of GenAI energy consumption has been on model training because the infrastructure required to train a state-of-the-art LLM is significant (Luccioni et al., 2023). However, current advancements, such as Chain-of-Thought prompting (Wei et al., 2022), have yielded massive increases in energy demand in the operating phase (i.e., for the inference). In fact, current research suggests that every single query consumes a small amount of energy, which scales up massively with millions of requests per day. For this reason, the Electric Power Research Institute estimates that inference accounts for approximately 60% of AI's total energy consumption (Aljbour et al., 2024, p. 15).

While there are growing numbers of studies that provide insight into the overall energy consumption of AI models (Luccioni et al., 2024b; Jegham et al., 2025), there is a significant gap in scientific research that investigates AI use from a socio-technical perspective. We argue that this gap is critical for the following two reasons. First, current research shows that generative AI is being adopted much more quickly than previous technologies, such as the personal computer (Bick et al., 2026). Therefore, the investigation of individual use behavior can have far-reaching implications. Second, there is a massive amount of energy demand related to AI and use behavior accounts for approximately 60%. Understanding AI use from a socio-technical perspective can thus be the foundation for meaningful and appropriate interventions.

This study is designed as a first step toward better understanding the relationship between AI use and energy consumption from an individual perspective. This approach is grounded in previous literature, which has investigated public perception in sustainability-related domains, including public perception of energy consumption (Attari et al., 2010), water consumption (Attari, 2014), climate change (Hansen et al., 2012; Mildemberger et al., 2025), environmental concerns (Pearson et al., 2018), and greenhouse gas emissions from food (Camilleri et al., 2019).

A key insight from these studies is that people have difficulty processing abstract, long-term, and less tangible threats, of which climate change is one example. This also results in biased estimates of demand, such as the energy costs of household devices (Attari et al., 2010). A common finding in most studies is that people tend to underestimate the relevance and impact of high levels of environmental concern among minority and low-income communities (Pearson et al., 2018). Dieckmann et al. (2021) argue that tangibility is a

crucial factor in shaping public perception. Building on this, we argue that the intangible nature of AI's energy consumption also leads to misaligned perceptions of its true scale and impact.

Theoretical foundation and hypotheses

Construal level theory and psychological distance

We conceptualize this study based on *psychological distance*, which is part of Construal-Level Theory (CLT; Trope et al., 2007; Trope & Liberman, 2010; Liberman & Trope, 2014, 2008). It provides a theoretical stance to explain the sustainable behavior of individuals (for an overview, we refer to Keller et al., 2022, and the references therein) and posits that the farther away an object, event, or person is from direct experience, the more abstractly and broadly we think about it. In contrast, the closer an object is, the more concrete and detailed our thoughts become (Trope et al., 2007; Trope & Liberman, 2010; Liberman & Trope, 2014). Psychological distance is defined as “*the extent of divergence from direct experience of me, here and now along the dimensions of time, space, social perspective, or hypotheticality*” (Liberman & Trope, 2014, p. 365). Temporal distance occurs when there is a time lag between events. For example, psychological distance will be greater if an event (e.g., the introduction of a new sustainability policy) occurred in the past (e.g., a decade ago) or will occur in the future than if it is happening now. Spatial distance relates to geographical locations. For example, a data center based in the United States feels psychologically more distant to a perceiver in the European Union (EU) than one situated within the EU itself. Social distance relates to differences between individuals. For example, psychological distance tends to be greater between individuals from different groups (e.g., managers and employees) than between those within the same group. Finally, there is the concept of hypotheticalness, which refers to the likelihood that an event will happen. For example, an event that is very unlikely, such as winning a lottery, comes with a greater psychological distance. The four distances converge into a subjective psychological distance. Psychological distance is not an aggregate of these four forces but rather a property triggered by information on all four correlated forces (Fiedler et al., 2012). Overall, psychological distance influences the way people mentally represent the world around them (Trope et al., 2007; Trope & Liberman, 2010; Liberman & Trope, 2014), which has implications for perception and behavior.

Psychological distance of GenAI's energy consumption

We posit that the energy consumption associated with GenAI is characterized by a high level of psychological distance.


This distance primarily arises from the *structural properties* of GenAI service provision, especially the inherent spatial separation between the point of GenAI application (e.g., a user's local device) and the physical location of its computational infrastructure (i.e., remote data centers) (Kim et al., 2025). The industry's nomenclature, such as "cloud service," reinforces the perception of remoteness, distinguishing GenAI from technologies where energy consumption and operation are co-located. For example, interacting with ChatGPT involves using a local user interface, while the actual computing is performed on geographically distributed servers, which may be located outside the user's country of residence (Tang, 2025; Kim et al., 2025). In contrast to other cloud-based services (e.g., Office software running in the cloud), many services targeting end-user were previously located on the users' devices. GenAI tools, including ChatGPT, have primarily been introduced as cloud-based services. This is in sharp contrast to other energy-consuming devices within users' perception, such as conventional household appliances like dishwashers or computers, which consume energy at their point of use, creating a direct perceptual link for users. This makes it seem as though the energy consumption of these appliances takes place "in a near place," while the energy use of GenAI takes place "in a distant place" (Fiedler et al., 2015). The *spatial distance* in GenAI systems is compounded by a perceptual distance stemming from a lack of transparency. Not only are the energy demands of cloud-based services hard to estimate because the reporting of such services is aggregated globally (Mytton, 2020), but the true energy demands of training and operating GenAI models are especially opaque to the end-user (Luccioni et al., 2024a). Combining these arguments, we conclude that psychological distance is a pivotal lens for examining perceptions of energy consumption when contrasting household devices with generative AI applications. However, we acknowledge that the structural realities of GenAI service provision, particularly cloud-based operation, operational opacity, and missing feedback (as noted by Luccioni et al., 2024a), can directly act upon outcomes such as awareness. Therefore, these factors are able to contribute to psychological distance and also act as obstacles, preventing users from understanding the energy consumption of GenAI.

Figure 1 illustrates our proposed research model. *Awareness of energy consumption* involves understanding which interactions with AI are more energy-intensive than others (Attari et al., 2010). Responsibility in the context of information systems involves taking ownership of the harmful and beneficial consequences that arise from technological interactions (Recker et al., 2025), such as energy consumption. In this research, we define responsibility attribution as the subjective assignment of responsibility for AI's energy consumption among the involved parties (Bleher & Braun, 2022; Latane & Darley, 1968). Pro-environmental behavior is broadly defined as "behavior that harms the environment as little as possible, or even benefits the environment" (Steg & Vlek, 2009, p. 309). Based on this definition, we define *pro-environmental strategies* as behavior that intends to reduce the energy footprint when using AI and AI-empowered tools (Attari et al., 2010). Note that this definition does not validate the objective impact of user-proposed strategies on energy reduction.

Awareness of GenAI's energy demand

According to CLT, psychological distance systematically alters how individuals mentally represent objects and events (Soderberg et al., 2015). Greater distance fosters higher-level construals, which are abstract mental representations that emphasize central, goal-relevant features. Conversely, shorter distances promote lower-level construals, which are more concrete and include specific, secondary details. As geographic distance increases, so does the level of abstraction, with remote events described in increasingly generalized terms rather than concrete details (Sneffjella & Kuperman, 2015). This shift in mental representation directly impacts judgment and behavior (Kogut et al., 2018). As argued above, the use of GenAI could come with a high degree of psychological distance. Following CLT, this distance should induce a high-level, abstract construal of GenAI's energy consumption. Consequently, users are predicted to focus on the primary, goal-relevant aspects of their interaction (e.g., "Can ChatGPT complete this task for me?") while peripheral, concrete, secondary details (e.g., the energy required to power the computation) become less salient.

Fig. 1 Research model



Because of the diminished focus on secondary attributes, we propose that users will be less aware of the energy consumption of GenAI. This has been shown for other cloud-based services, where users have low levels of awareness of the carbon emissions associated with cloud storage (Gujral et al., 2025) and often assume such services carry no environmental cost (Siy et al., 2025). We propose household devices as a baseline of low psychological distance. Unlike AI services, which are distal, cloud-based, and opaque, household devices are physically proximal and tangible, and their energy use is often associated with immediate feedback. Thus, we posit the following:

H_1 : Individuals are less aware of AI-related energy consumption compared to the energy consumption of household devices.

Responsibility diffusion

Psychological distance can influence perceived responsibility in an IS context (Tong et al., 2025). According to CLT, events that are psychologically distant are construed in more abstract terms, which can diminish an individual's sense of responsibility and urgency. We extend this principle and apply it in the context of AI's energy consumption. The abstract construal of AI energy use not only makes the impact feel more remote, but also facilitates the diffusion of responsibility by making the individual's role in the network of AI infrastructure less concrete. When responsibility is unclear or shared with (too) many others, this can lead to the feeling that no one is responsible (Bandura et al., 1996). This diffusion of responsibility is a key explanation for the bystander effect, where the presence of others hinders action because the focal person's sense of responsibility is lessened (Latane & Darley, 1968; Fischer et al., 2011). Similarly, in the context of GenAI, responsibility diffusion has been documented with AI-based clinical decision support systems due to the expanding number of possible moral agents involved (Bleher & Braun, 2022). The environmental impact of GenAI use is clearly a many-hands problem, and the impact of a single user's actions is not straightforward to quantify, yet the cumulative actions of all users have a huge effect (Luccioni et al., 2024b, 2025). Thus, we argue that responsibility is diffused across a vast network of actors, including developers, users, and policy-makers. This contrasts with typical household devices, where responsibility for energy consumption is more clearly assigned to the person operating the device. Consequently, there is a lower attribution of responsibility to the user for AI's energy footprint than there is for typical household devices. This leads us to propose our second hypothesis:

H_2 : Individuals perceive less responsibility for AI-related energy consumption compared to the energy consumption of household devices.

Pro-environmental strategies

Besides the impact on responsibility, psychological distance also affects how we evaluate actions and alternatives. When psychological distance increases, desirability (the "why" aspect) of an action becomes more important than feasibility (the "how" aspect; Todorov et al., 2007). With a higher distance, people are also able to find more pros (potential benefits) than cons (potential negative aspects) to an action (Eyal et al., 2004). Increasing psychological distance can mitigate the feeling of difficulty that is elicited by complex tasks (Thomas & Tsai, 2011). As argued above, energy consumption related to GenAI can be psychologically more distant, which is why we assume that people will have a high-level, abstract mindset. An abstract mindset makes it harder to answer "how" questions and lets one instead focus on the "why." In the context of GenAI energy consumption, and in line with our theoretical lens, we expect respondents to focus on seemingly easy-to-implement solutions that lack feasibility upon closer inspection, such as reducing their use of GenAI, abstaining from it entirely, or finding alternatives. Conversely, we expect respondents to be less likely to propose concrete solutions that emphasize the "how," such as prompt optimization, model selection, and modality choice. In addition to the cognitive mechanisms identified by CLT, other explanations are possible. In particular, the relative novelty of GenAI suggests that users' strategic repertoires for mitigation are likely still in the early stages of development. Consequently, we expect that individuals are more likely to state simplistic and optimistic solutions (e.g., abstain or reduce use) to mitigate AI-related energy consumption compared to concrete solutions such as prompt optimization or model selection as they are less knowledgeable and less informed about GenAI compared to ordinary household devices. Since this aspect is exploratory in nature, we pursue to investigate participants' responses also in an exploratory manner, guided by our two propositions (P):

P_1 : Individuals will be more likely to propose solutions that focus on limiting use, such as abstaining from or reducing GenAI usage.

P_2 : Individuals will be less likely to suggest concrete solutions that focus on optimizing interactions, such as prompt optimization, model selection, and modality choice.

Methodology

Mixed-method design

To address the outlined hypotheses and propositions, we adopt a mixed-methods approach (Venkatesh et al., 2013), combining quantitative and qualitative data to comprehen-

sively understand public perceptions of GenAI energy consumption. This study uses an online survey with a simplified, US Census-based sample where participants completed two ordinal ranking tasks—one for AI tasks and one for household devices—to capture awareness of energy consumption, as well as a modified dollar metric scale to elicit responsibility attributions for the energy consumption across five actors. For the qualitative analysis, participants were asked to provide open-ended responses about their personal energy-reduction strategies. These responses were coded using a theory-derived conceptual framework and an LLM ensemble approach to code the responses (see Appendix A).

Data collection

Data was collected using an online survey through the Prolific online platform, utilizing simplified US Census selection criteria that consider sex, age, and ethnicity. Because this study aims to investigate generic attitudes across a broad socioeconomic spectrum, rather than shared experiences within a specific organizational context, we chose a crowdsourcing panel for data collection (see Jia et al., 2017). We obtained answers from 1098 individuals and paid £1.24 for a median participation time of nine minutes. At the end of the survey, we asked participants if they had answered seriously and were not distracted. We dropped 18 observations from participants who indicated that they had not answered seriously or had been distracted. All the data used here, along with the corresponding analysis scripts, will be deposited in the Open Science Framework.³

Demographics

The final sample comprised 1080 participants ($\text{Mean}_{age} = 45.94$, $\text{Median}_{age} = 46$, $\text{SD} = 16.13$), whose ages ranged from 18 to 89 years. The gender breakdown was as follows: 550 (50.9%) identified as female, 513 (47.5%) as male, 12 (1.10%) as other, and 5 (0.50%) preferred not to disclose their gender. Regarding income, 0.6% of participants reported having no income. Among those reporting an annual income, 24.9% earned between \$20,000 and \$49,999, 23.0% between \$50,000 and \$79,999, 13.7% between \$80,000 and \$109,999, 11.5% between \$110,000 and \$139,999, 10.2% between \$140,000 and \$169,999, and 7.6% reported incomes above \$170,000. Participants also provided their political orientation: 10.6% identified as very liberal, 18.6% as liberal, 9.7% as slightly liberal, 22.3% as moderate, 10.6% as slightly conservative, 18.7% as conservative, and 9.5% as very conservative. The sample exhibited a broad range of educational backgrounds. Specifically, 1.0% reported having received only some schooling without earning a diploma or degree,

11.3% possessed a high school diploma or GED, 20.1% had attended some college, and 37.5% had earned a college degree. Furthermore, 3.6% had undertaken some graduate studies, and 26.5% had completed a graduate degree. An overview is provided in Table 2.

Tasks

Accurately estimating energy consumption can be difficult for respondents, as they often misjudge absolute units such as watts or kilowatts (Andor et al., 2020; Attari et al., 2010). While reference points or the specific choice of unit (e.g., watt versus kilowatt) can anchor judgments (Frederick et al., 2011), the absence of any reference can make judging energy use more difficult (Attari et al., 2011). We therefore

Table 2 Sample description

	n	Percentage
Gender		
Female	550	50.90%
Male	513	47.50%
Other	12	1.10%
Prefer not to answer	5	0.50%
Age		
18–24	122	11.30%
25–34	195	18.10%
35–44	193	17.90%
45–54	167	15.50%
55–64	257	23.80%
65–74	114	10.60%
Over 75	32	3.00%
Income		
Did not have an income	7	0.60%
< \$20,000	92	8.50%
\$20,000–\$49,999	269	24.90%
\$50,000–\$79,999	248	23.00%
\$80,000–\$109,999	148	13.70%
\$110,000–\$139,999	124	11.50%
\$140,000–\$169,999	110	10.20%
> \$170,000	82	7.60%
Political orientation		
Very liberal	114	10.60%
Liberal	201	18.60%
Slightly liberal	105	9.70%
Moderate	241	22.30%
Slightly conservative	114	10.60%
Conservative	202	18.70%
Very conservative	103	9.50%

³ <https://osf.io/8z4qj>

rely on an ordinal task-ranking approach (Attari, 2014). Participants performed two ranking exercises. First, they ordered five common AI tasks (e.g., image generation and question answering), inspired by previous literature (Luccioni et al., 2024a). Participants were invited to sort these tasks based on their energy demand. Second, the participants were asked to rank five common household tasks (e.g., room heating and computer use) from the most to the least energy-intensive in order to enable a systematic comparison. The sorting tasks measure respondents' ability to rank tasks within domains. To do so correctly, they need to understand which task is more energy-intensive than another (Attari et al., 2010), which reflects a form of relative energy awareness. Since users typically compare options rather than absolute consumption figures (e.g., 2 Wh), this operationalization aligns with how energy-related awareness manifests in practice. The order of the tasks was randomized in both domains, and there was no default order. Subjects had to drag and drop the tasks into a blank field. To simplify the comparison, all tasks were set to 8 h. We use five different AI tasks: **image generation** ("Using AI (e.g., DALL-E 3) to generate images for 8 h"), **image captioning** ("Using AI for image captioning (e.g., Claude) for 8 h (≈ 80 prompts). Example: Take an image and generate an appropriate caption describing the image."), **automatic speech recognition** ("Using AI (e.g., Whisper) to run automatic speech recognition for 8 h (≈ 3 h speech). Example: Transcribe a voice message to text."), **question answering** ("Using an AI (e.g., ChatGPT) to answer questions for 8 h (≈ 80 prompts). Example: Ask ChatGPT for a healthy recipe."), and **text classification** ("Using AI (e.g., ChatGPT) to classify text for 8 h (≈ 80 prompts). Example: Identify a document as an invoice."). Based on the estimated energy consumption reported in Luccioni et al. (2024a), we consider the following

vector as the true order (solution vector): Image generation $\Rightarrow$ Image captioning $\Rightarrow$ Automatic speech recognition $\Rightarrow$ Question answering $\Rightarrow$ Text classification. An overview is provided in Table 3.

Similar to the first task, we asked participants to sort five tasks related to domestic energy consumption. These include **Heating a small room** (16 ft $\times$ 10 ft $\times$ 6.5 ft) with a space heater for 8 h, **Streaming Netflix movies** on a large television for 8 h, **Using a slow cooker** to prepare a meal for 8 h, **Using Office applications** on a desktop computer for 8 h, and **Using a cable modem** for 8 h. Based on the average energy consumption values from the US Department of Energy, we consider the following vector as the true order (solution vector): Space heater $\Rightarrow$ Television $\Rightarrow$ Slow cooker $\Rightarrow$ Computer - desktop PC $\Rightarrow$ Router/DSL/cable modem. An overview is provided in Table 4.

Measurements

To measure responsibility, we asked participants to identify which actors should be responsible for energy consumption with regard to AI and, to contrast our findings, also with regard to household devices (hereafter: 'Home'). The question to assess responsibility for AI was: *Who should be taking care of the energy consumption of generative AI?* and the question to measure responsibility for household devices was: *Who should be taking care of the energy consumption of household appliances?*. We used a modified version of the dollar metric scale (Louviere & Islam, 2008), in which participants were given 100 points to allocate across five actors: (1) *Policy-makers*, (2) *Producers*, (3) *Users*, (4) *Engineers*, and (5) *No one* to indicate how much responsibility they believed each actor should take for the energy consumption of

Table 3 Energy consumption task with AI

Task	Text	Actual energy consumption ¹
Image generation	Using AI (e.g., DALL-E 3) to generate images for 8 h (≈ 80 prompts). <i>Example: Generate a picture of a dog.</i>	Average energy consumption = 477 Wh $\approx$ 3.8 kWh for the given scenario
Image captioning	Using AI for image captioning (e.g., Claude) for 8 h (≈ 80 prompts). <i>Example: Take an image and generate an appropriate caption describing the image.</i>	Average energy consumption = 109 Wh $\approx$ 0.8 kWh for the given scenario
Automatic speech recognition	Using AI (e.g., Whisper) to run automatic speech recognition for 8 h (≈ 3 h speech). <i>Example: Transcribe a voice message to text.</i>	Average energy consumption = 25 Wh $\approx$ 0.2 kWh for the given scenario
Question answering	Using an AI (e.g., ChatGPT) to answer questions for 8 h (≈ 80 prompts). <i>Example: Ask ChatGPT for a healthy recipe.</i>	Average energy consumption = 1 Wh $\approx$ 0.08 kWh for the given scenario
Text classification	Using AI (e.g., ChatGPT) to classify text for 8 h (≈ 80 prompts). <i>Example: Identify a document as an invoice.</i>	Average energy consumption = 0.5 Wh $\approx$ 0.04 kWh for the given scenario

Source: The average energy consumption values are based on Luccioni et al. (2024a)

Table 4 Energy consumption task with household devices

Task	Text	Actual energy consumption ¹
Space heater	Heating a small room (16 ft × 10 ft × 6.5 ft) with a space heater for 8 h	Wattage: 1320, utility rate: US average, hours used per day: 8, days used per year: 1. This results in 10.6 kWh
Television, SD/HD TV, > 40"	Streaming Netflix movies on a large television for 8 h	Wattage: 234, utility rate: US average, hours used per day: 8, days used per year: 1. This results in 1.9 kWh
Slow cooker	Using a slow cooker to prepare a meal for 8 h	Wattage: 200, utility rate: US average, hours used per day: 8, days used per year: 1. This results in 1.6 kWh
Computer - desktop PC	Using office applications on a desktop computer for 8 h	Wattage: 75, utility rate: US average, hours used per day: 8, days used per year: 1. This results in 0.6 kWh
Router/DSL/cable modem	Using a cable modem for 8 h	Wattage: 6, utility rate: US average, hours used per day: 8, days used per year: 1. This results in 0.048 kWh

Source: The actual energy consumption for home devices was estimated using the energy calculator provided by the US Department of Energy <https://www.energy.gov/energysaver/estimating-appliance-and-home-electronic-energy-use>, accessed October 24, 2025

either AI or common household devices. We provided examples of each entity for participants, i.e., policy-makers (e.g., the government), vendors of systems (e.g., OpenAI/General Electric), users (e.g., citizens), and developers (e.g., OpenAI employees/General Electric employees).

To identify promising strategies to reduce the energy consumption of AI during use, we relied on a qualitative response using an open-ended text field. We asked respondents: *What do you think is the most effective thing you could personally do to save energy when using generative AI?* which we adapted from Attari et al. (2010).

Fig. 2 A conceptual framework illustrating individual strategies to reduce AI energy consumption




Table 5 Strategies for energy reduction when using AI

Strategy	Example (quote)	Rationale	References
(1) Select provider	Select a provider that is more energy efficient than alternatives. (<i>I make use of AI from companies that support renewable energy [...]</i> , participant #634)	Data centers comprise servers, storage, cooling, and power that vary in energy consumption	(Lacoste et al., 2019; Hugging Face, 2024)
(2) Select model	Use a smaller, tailored model instead of general-purpose models and/or smaller models instead of larger models. For example, choose the “mini” model (e.g., GPT-4o mini) rather than the “reasoning” model (e.g., GPT-o3). (<i>Utilizing task-specific AI models [...]</i> , participant #978)	Large models (e.g., multi-purpose model with a high number of parameters) are more energy-intensive than small models (e.g., fewer parameters and fine-tuned)	(Luccioni et al., 2024b; Maliakel et al., 2025)
(3) Select modality	Prefer text whenever possible and avoid generation of visual content. For example, avoid illustration and ask for a textual description instead. (<i>Use it more for text rather than images.</i> , participant #669)	Multimedia data, such as images, require significantly more energy compared to text (e.g., text classification or text captioning)	(Luccioni et al., 2024b, a)
(4) Optimize prompt	Use specific and concise prompts, limit requests for lengthy explanations. (<i>Carefully select prompts to reduce revisions.</i> , participant #908)	Response length triggered by ambiguous prompts, long prompts, or commands to explain or justify increases energy demands	(Adamska et al., 2025)
(5) Reduce usage	Limit the amount and/or intensity of use. (<i>Limit my use of generative AI in general.</i> , participant #662)	Fewer interactions reduce cumulative energy consumption	(Chen et al., 2025)
(6) Use alternatives	Use alternatives to AI when suitable. For example use Google search instead of asking ChatGPT. (<i>To stop using it as often and just google it instead.</i> , participant #946)	Traditional technologies are known to consume significantly less energy compared to the energy needs of LLMs	(Aljbour et al., 2024)
(7) Abstain AI	Do not use AI at all. (<i>Just don't use it.</i> , participant #224)		—

Finally, we asked participants the following question: *Do you feel that you paid attention, avoided distractions, and took the survey seriously? Please answer honestly. Your answer will not affect your compensation.* This question has been previously used as a measure to identify low-quality responses (Stanley et al., 2020).

Coding of strategies

Previous literature has conceptualized the use of information technology, acknowledging the information system, the user, and the task (Burton-Jones & Straub, 2006). We build upon this stream of literature and use a conceptual framework to identify individual strategies that can be used to reduce AI energy consumption (see Fig. 2). We consider the user as a self-determined actor who can make decisions, including choosing between the use of AI or a non-AI alternative (e.g., using the World Wide Web or abstaining from technology) and the amount of use. If choosing in favor of AI, a user

can select a specific provider and a specific LLM. Moreover, when interacting with an LLM—and depending on the task—the modality can be chosen (e.g., text, audio, or video) and the prompt can be varied. Based on this conceptualization, we deduce seven strategies: (1) select provider, (2) select model, (3) select modality, (4) optimize prompt, (5) reduce usage, (6) use alternatives, and (7) abstain from AI.

Each strategy corresponds with rationales that justify energy reduction. For example, selecting a provider is relevant since different data centers vary in their energy consumption and are commonly evaluated using metrics such as Gigafllops/Watt,⁴ which can be linked to existing literature (Lacoste et al., 2019; Hugging Face, 2024). Therefore, choosing a more efficient model can be considered an appropriate strategy for energy reduction when using AI. Table 5 provides an overview of each strategy, including an exam-

⁴ See, for instance, the Green 500 list of High Performing Clusters <https://www.top500.org/lists/green500/>

ple quote from the data, the rationale in terms of how this strategy can foster energy reduction, and links to literature related to the strategy.

Those seven strategies are used as a foundation for our coding procedure, which is further detailed in Appendix A. Besides those strategies, individual responses also include answers that are in some way relevant to reducing energy consumption but do not fit into either one of the seven strategies. These are coded as “other.” Moreover, answers indicating that the respondent is not aware of any strategy are coded as “not sure/don’t know.” Finally, we also coded answers that are irrelevant to this research and coded them as “irrelevant.”

Results

Awareness of task-specific energy consumption

Most of the participants were able to solve the sorting tasks with only a few errors, and a positive correlation between the empirical data and the actual solution vector was found in both domains (i.e., $\rho_{AI} = 0.37$, $\rho_{Home} = 0.46$). Figure 3 shows the distribution of the number of swaps (right-skewed) needed to bring the solution vectors into the correct vector order. On average, there were 3.06 swaps needed for household tasks, and 3.46 swaps for the AI tasks, respectively. A t -test on the Fisher-transformed correlations ($z = \frac{1}{2}[\ln(1+r) - \ln(1-r)]$) between the solution vectors and the observed vectors revealed a significant difference of -0.391 ($t(2158) = -2.09$, $p = .037$).

These results suggest that participants can rank tasks according to their energy intensity relatively accurately. While they were marginally more accurate at ordering household tasks than AI tasks, they were generally aware that




Fig. 3 Comparison of the sorting-task performance of $N=1080$ participants with 5 sorting items. (A) Bar plot showing the number of swaps required to transform the response to the correct vector, ranging from 0 (perfect order) to 10 (maximum swaps). The maximum number of swaps is calculated as $\frac{n(n-1)}{2}$, where $n = 5$ items, resulting in 10 maximum swaps. (B) Inset showing the results of an additional inference test using a Fisher-transformed t -test

specific generative AI tasks (e.g., image generation) were more energy-intensive than other AI tasks (e.g., question answering). This indicates a basic awareness of energy consumption in relation to different tasks.

Perceived responsibility for AI energy consumption

“Policy Makers” received significantly more responsibility for AI ($M = 18.92$, $SD = 16.89$) than Home ($M = 13.34$, $SD = 14.77$; $t(2158) = 8.17$, $p < .001$). The effect size, as measured by Cohen’s d , was $d = 0.35$, indicating a small effect. No significant difference was observed for “Vendors” between AI ($M = 25.76$, $SD = 20.04$) and Home ($M = 27.15$, $SD = 19.90$; $t(2158) = -1.62$, $p = .106$). The effect size, as measured by Cohen’s d , was $d = -0.07$, indicating a very small effect. In the Home category ($M = 38.82$, $SD = 29.52$), “users” received significantly more responsibility than in the AI category ($M = 22.18$, $SD = 21.04$; $t(2158) = -15.09$, $p < .001$). The effect size, as measured by Cohen’s d , was $d = -0.65$, indicating a small effect. “Developers” received significantly more responsibility in the AI category ($M = 28.69$, $SD = 20.00$) compared to the Home category ($M = 17.72$, $SD = 15.50$; $t(2158) = 14.24$, $p < .001$). The effect size, as measured by Cohen’s d , was $d = 0.61$, indicating a medium effect. “No One” received slightly, but significantly more responsibility in AI ($M = 4.37$, $SD = 14.80$) than in Home ($M = 2.69$, $SD = 10.10$; $t(2158) = 3.07$, $p = .002$). The effect size, as measured by Cohen’s d , was $d = 0.13$, indicating a very small effect. All reported p -values remain significant after Bonferroni correction for multiple comparisons.

These results show that the perceived responsibility of users, developers, and policy-makers differs significantly between the two domains. For AI, responsibility is more diffused and users are seen as less responsible, while developers and policy-makers are seen as more responsible. For “Home,” responsibility gravitates more towards the user. An overview is provided in Fig. 4.

Strategies related to reducing AI energy consumption

Figure 5 illustrates the range of strategies proposed by participants. A majority of the suggestions focus on limiting engagement: 33.5% recommend reducing AI usage, 8.2% suggest total abstinence, and 3.2% propose using non-AI alternatives. Collectively, these restrictive measures account for 44.9% of all responses. In contrast, only a small minority (5.7%) suggests optimizing usage through more conscious selections of providers (1.3%), models (2.1%), or modalities (2.3%). 34.5% of responses were either not sure (9.9%), suggested something else (21.0%), or suggested something irrelevant to the topic at hand (3.6%).




Fig. 4 Boxplots showing group-wise differences between AI and Home conditions across five stakeholder categories: Policy Makers, Vendors, Users, Developers, and No One. “ns” indicates $0.050 < p \leq 1.000$; “*”

indicates $0.010 < p \leq .050$; “**” indicates $0.001 < p \leq .010$; “***” indicates $0.000 < p \leq .001$; “****” indicates $p \leq .000$. Bonferroni correction was applied to all p-values

To further investigate boundary conditions of the mentioned strategies, we ran logistic regressions using self-reported measures of GenAI use, GenAI proficiency, gender, age, and education as predicting variables. The results are shown in Table 6. For the strategies of provider choice, alternatives, and irrelevant responses, the predictors did not significantly explain the likelihood of a participant proposing them. In contrast, the choice of model was significantly predicted by proficiency and age: individuals with higher self-reported proficiency ($B = 0.43$, $p = .030$) and younger participants ($B = -0.04$, $p = .009$) were more likely to suggest this strategy. Modality was significantly predicted by

gender, with men being more likely to propose this strategy ($B = 0.79$, $p = .042$). Prompt optimization was significantly influenced by proficiency ($B = 0.26$, $p < .001$). Reduction of use was influenced by the amount of GenAI use, with more frequent users being more likely to propose this strategy ($B = 0.17$, $p < .041$). Three variables predicted the suggestion to abstain from AI use: usage ($B = -0.85$, $p < .001$), proficiency ($B = -0.25$, $p = .014$), and age ($B = -0.02$, $p = .001$). Older participants were more likely to suggest the “Not sure” category ($B = 0.02$, $p < .001$). Finally, the results indicate that the “Other” category was significantly associated with proficiency ($B = 0.18$, $p = .006$). One noteworthy pattern




Fig. 5 Distribution of pro-environmental strategies derived from majority voting of open-text responses ($N = 1080$ participants; 1345 total codings). Multiple strategies per participant permitted

Table 6 Logistic regression results for pro-environmental strategies. Coefficients are shown with *p*-values in parentheses

Strategy	Provider	Model	Modality	Prompt	Reduce
Intercept	−5.94(0.000)	−4.87(0.000)	−4.20(0.000)	−2.78(0.000)	−0.36(0.226)
Male	0.62 (0.227)	0.59 (0.144)	0.79 (0.042)	0.01 (0.933)	0.10 (0.405)
GenAI usage	0.40 (0.233)	0.11 (0.669)	0.31 (0.209)	0.09 (0.375)	0.17 (0.041)
Proficiency	0.14 (0.565)	0.43 (0.030)	0.02 (0.888)	0.26 (0.000)	−0.09 (0.124)
Age	−0.01 (0.390)	−0.04(0.009)	−0.02 (0.080)	−0.00 (0.541)	−0.00 (0.297)
Pseudo <i>R</i> ²	0.059	0.120	0.053	0.044	0.004
Strategy	Alternative	Abstain	Not sure	Other	Irrelevant
Intercept	−2.52(0.001)	1.96 (0.000)	−2.51(0.000)	−2.67(0.000)	−4.14(0.000)
Male	−0.32 (0.322)	0.13 (0.543)	−0.36 (0.058)	−0.07 (0.609)	−0.08 (0.781)
GenAI usage	−0.03 (0.887)	−0.85(0.000)	−0.04 (0.768)	0.13 (0.145)	0.12 (0.532)
Proficiency	−0.08 (0.566)	−0.25(0.014)	−0.04 (0.639)	0.19 (0.003)	0.03 (0.809)
Age	−0.00 (0.806)	−0.02(0.001)	0.02 (0.001)	0.01 (0.040)	0.01 (0.197)
Pseudo <i>R</i> ²	0.006	0.191	0.025	0.030	0.007

Note. McFadden's pseudo-*R*² is reported
Significant values (*p* < .05) in bold

was that optimizing usage (e.g., model selection and prompt refinement) was more frequently mentioned with higher proficiency levels. Using AI more often was associated with a greater likelihood of recommending a reduction in its use but not for abstaining from it entirely.

Discussion

Discussion of results

The global deployment and utilization of GenAI, particularly the trend towards more computationally intensive inference-

time scaling techniques (Brown et al., 2025), are projected to substantially increase energy consumption and associated carbon emissions. Although organizations control the production and deployment of AI, individual usage patterns can substantially impact overall energy consumption. This study takes this development into perspective and seeks to advance the understanding of public perceptions of energy consumption related to GenAI. This study thus ties into the existing discourse on public perception of energy consumption and the corresponding psychological barriers to energy conservation (Attari et al., 2010; Attari, 2014; Asensio & Delmas, 2015; Dietz, 2014; Gromet et al., 2013), integrates the emerging literature on AI-related energy demand (Luccioni et al.,

Table 7 Overview of hypotheses

#	Proposed hypotheses	Main result
<i>H</i> ₁	Individuals are less aware of GenAI-related energy consumption compared to the energy consumption of household devices	(Supported): Based on an energy ranking task, we observe a slightly poorer performance with GenAI tasks compared to household tasks (<i>p</i> = .037)
<i>H</i> ₂	Individuals perceive less responsibility for GenAI-related energy consumption compared to energy consumption of household devices	(Supported): Respondents perceive users as significantly less responsible for the energy consumption of GenAI than for the energy consumption of household appliances (<i>p</i> < .000)
#	Propositions	Main result
<i>P</i> ₁	Individuals will be more likely to propose solutions that focus on limiting use, such as abstaining from or reducing GenAI usage	The majority of responses indicated a limitation strategy, especially reducing use (33.5%; this strategy was most likely to come from respondents who indicated high usage) and abstaining from GenAI altogether (8.2% of responses)
<i>P</i> ₂	Individuals will be less likely to suggest concrete solutions that focus on optimizing interactions, such as prompt optimization, model selection, and modality choice	The least likely provided responses were prompt optimization, model selection, and modality choice, accounting for a cumulative 5.7% of all responses

2024a, b), and provides empirical data on public perception of energy consumption in AI. Table 7 provides an overview of our proposed hypotheses and our main results.

With regard to H_1 (Individuals are less aware of GenAI-related energy consumption compared to the energy consumption of household devices), we show that, despite the relative novelty of AI, the public is moderately aware of the energy consumption associated with AI. Specifically, we show that individuals are nearly as good at estimating the relative energy consumption of AI tasks as they are in terms of household tasks. One possible explanation for this awareness is the considerable media attention, which has likely sensitized the public to the technology's energy demands and informed individuals about energy-related aspects, such as increased demand in general (Chen, 2025) and the energy costs associated with the generation of different modalities (Luccioni et al., 2024a). Given the growing body of research highlighting AI's energy demands (Luccioni et al., 2024a, b), this awareness is likely to increase in the future. Consequently, we assume that the difference between conventional products and GenAI will further diminish in the near future.

Regarding H_2 (Individuals perceive less responsibility for AI-related energy consumption compared to energy consumption of household devices), we observe a diffusion of responsibility and a shift in perceived responsibility compared to traditional technology, away from users and towards developers and policy-makers. Despite being aware of AI's energy demands, the public exhibits relatively low levels of perceived responsibility. A large share of responsibility for AI's energy consumption is instead attributed to other actors, such as developers and policy-makers. This may be in part because cloud-based AI is not on site. Consistent with CLT, this could create a psychological distance that obscures responsibility.

In our exploratory investigation of user-proposed strategies for mitigating the environmental impact of GenAI, we found that 33.5% of all responses were related to reducing or limiting use. Even though this is well-intentioned, reducing use will be increasingly challenging since most digital tools incorporate some form of GenAI. There was also a considerable number of individuals (15.1%) who suggested prompt optimization. We consider this option valuable because it will encourage conscious use of AI. This also applies to provider and model selection and modality choice, which, unfortunately, only a small percentage of individuals consider (5.7%). Our analysis revealed that individuals with higher GenAI proficiency were more likely to mention strategies that optimize interaction, such as prompt refinement. Since GenAI is relatively novel, it makes sense that those with higher proficiency can think of more specific, effective ways to regulate energy demand during GenAI interactions. Our data also show that there are too many individuals who are unsure about how to reduce energy consumption (9.9%)

or provided irrelevant answers (3.6%). We interpret these results as indicating a pressing need for more education on AI in general, as well as on the relationship between AI and energy consumption. There is also a need for suitable strategies to reduce energy consumption when interacting with GenAI. Additional robustness tests suggest that proficiency, amount of use, and age can explain a small but significant portion of individual responses. We conclude that more efforts are required to inform individuals about GenAI-specific topics, such as downstream consequences of GenAI use.

Implications for theory

Our study expands the growing body of research on Green IS and Green AI (Schwartz et al., 2020; Dedrick, 2010; Watson et al., 2010; vom Brocke et al., 2013) by developing a user-centric, socio-technical perspective on AI's energy consumption. While the environmental footprint of AI is a topic of increasing concern (Luccioni et al., 2024a; Chen, 2025; Strubell et al., 2020), the user's role in this ecosystem has been largely overlooked. We offer three primary contributions to theory.

First, we argue that the technological shift towards inference-time computing necessitates a parallel theoretical shift towards the individual user. As AI computations increasingly occur in response to direct user interactions, the user's agency in shaping AI's energy footprint grows, creating new economic and environmental dynamics. Our work extends the behavioral lens from established domains, such as smart metering (Wendt et al., 2024; Dalén & Krämer, 2017; Sim et al., 2023), to the novel context of AI. By doing so, we provide a theoretical rationale for why user-facing interventions, such as standardized energy labels (Luccioni et al., 2024a; Hugging Face, 2024) or direct feedback mechanisms (Karlin et al., 2015), are not just practical tools but also essential components of a socio-technical understanding of Green AI. A current review on Green AI argues that much research is focusing on the state of the art of Green AI instead of focusing on actual practices (Verdecchia et al., 2023). Using a socio-technical perspective, we offer initial insights into how individuals actually intend to use AI in a sustainable way.

Second, we draw from the concept of psychological distance (Trope & Liberman, 2010; Trope et al., 2007) as a novel theoretical lens for shedding light on user perceptions of AI energy use. Our results are consistent with the idea that the abstract and remote nature of cloud-based AI corresponds with high psychological distance, which is associated with lower awareness and diminished responsibility. This contribution is significant because it offers a theoretical stance that may help theorize about patterns such as those observed here. The concept of construal level (Soderberg et al., 2015)

provides IS researchers with an established framework to theorize about how mental representations shape the subsequent behavior of end users.

Third, our findings contribute to the nascent stream of research on digital responsibility (Recker et al., 2025). We provide strong empirical evidence that perceptions of responsibility are not monolithic but are highly context-dependent and distributed. By contrasting AI with household devices, we demonstrate that responsibility is diffused among various actors (e.g., vendors, developers, policy-makers) in the AI ecosystem. This finding challenges scholars to move beyond domain-general theories of responsibility. Instead, we show that to build a comprehensive theory of digital responsibility, researchers must account for how specific technological contexts and actor constellations fundamentally alter how responsibility is perceived and assigned.

Implications for practice

Since the operation of generative AI incurs significant costs, all measures for greater sustainability contribute to a reduction of CO₂ emissions. For organizations, greater sustainability usually means lower costs. Our study informs organizations in several ways:

First, we have conceptualized how individual strategies can be implemented to reduce AI energy consumption (see Fig. 5). Organizations can further support individuals by offering different models, particularly small models, using sustainable vendors, and implementing strategies to technically implement model routing as a complementary mechanism (Jitkrittum et al., 2025).

Second, our study informs organizations on how to educate employees in this regard. This contributes to the IS literature on AI literacy (see, for instance, Pinski et al., 2024; Pinski & Benlian, 2024) and provides a new dimension—namely, knowledge of the sustainable impact of GenAI—to individual literacy with AI. Moreover, organizations can use our results to explain and discuss different strategies (see Table 5) for empowering employees to make informed decisions about using GenAI at work.

Finally, we contribute to a broader discourse in electronic markets, enhancing digital ecosystems and digital platforms. Designers can leverage these insights to identify strategies and options to reduce energy consumption. They can then translate existing knowledge into actual impact in the daily use of these ecosystems. One practical way is to educate freelancers on online platforms about effective strategies for conserving resources during AI interactions. This is significant, given the size of the labor market, which is in the millions (Kässi et al., 2021). Additionally, cloud providers

of chatbots might use education to reduce costs in free segments, in addition to intelligent model routing.

Implications for policy

Pro-environmental action requires actors to recognize the behavioral consequences and take responsibility for engaging in mitigation (De Groot & Steg, 2009). However, despite being aware of AI's high energy demands, responsibility is diffused, which disrupts this chain of action. This may be partly due to the psychological distance of AI's environmental impact, which makes it easier to diffuse responsibility. To bridge this gap, we propose *proximizing* as a strategy that shifts remote consequences into the immediate context of the user (Brügger et al., 2016; Loy & Spence, 2020). Policy-makers can achieve this by mandating real-time feedback on the energy consumption of AI interactions and investing in educational interventions that translate abstract data into concrete environmental impacts (Bergquist et al., 2023). Regulatory bodies have recognized the ecological footprint of AI and are increasingly mandating more transparency. Notably, the EU AI Act codifies transparency requirements regarding the energy consumption of AI models over the life-cycle (European Union, 2024, see Art. 40(2) & Art. 95(2)). While current directives primarily target institutional reporting, expanding these disclosure requirements to the end-user interface could be an effective intervention.

Awareness of AI's ecological footprint is insufficient to drive change. Users must also be able to choose and implement effective mitigation strategies. Our findings suggest a knowledge gap: most users suggested reducing or abstaining from AI altogether, or simply did not know what to do. As AI will be ubiquitous, abstaining from it will be difficult and often not practical. To address this deficit, educational interventions are needed to empower sustainable engagement. A concrete starting point is teaching users to select task-specific models rather than defaulting to energy-intensive general-purpose models. In many cases, these smaller models provide sufficient performance with a significantly lower carbon footprint (Varoquaux et al., 2025).

Limitations and avenues for future research

This study's findings should be considered in light of its limitations, each of which presents a valuable opportunity for future research.

First, our study was conducted with a representative sample from the USA using the crowdsourcing provider Prolific, which is recognized for high data quality (Albert & Smilek,

2023). Nevertheless, reliance on such providers entails challenges, such as participant habituation, varying degrees of language proficiency, attentional lapses, and potential constraints on demographic diversity (Peer et al., 2017). Besides issues related to single-method bias from a single panel, the generalizability of our findings to other cultural contexts is an open empirical question. Cultural values, environmental norms, and regulatory landscapes can significantly shape perceptions of energy consumption and responsibility. Future cross-cultural research is needed to investigate how these factors may moderate the relationships we identified, providing a more globally nuanced understanding of this phenomenon.

Second, our survey-based design establishes strong associations but does not permit causal claims. We did not, for instance, experimentally manipulate psychological distance. While our findings are consistent with CLT predictions, it is important to emphasize that we do not empirically demonstrate that this theoretical lens is the only valid or primary mechanism. Structural parameters of GenAI, such as opacity, novelty, and limited control, are promising and viable alternative mechanisms. Thus, an important next step is to validate causal mechanisms. Future research could employ experimental designs that systematically manipulate different dimensions of psychological distance to directly test its causal effect on awareness, responsibility, and problem-solving approaches.

Future research could also utilize established scales (Fiedler et al., 2012, 2015) to calculate a psychological distance compound and directly test its correlation with responsibility ratings. Additionally, researchers could measure CLT-based predictions, such as mental representation and construal level (Trope & Liberman, 2010). For instance, the approach by Snefjella and Kuperman (2015) could be adopted to analyze whether users describe AI energy demand using abstract or concrete language.

Third, in line with established research on public perception (Attari, 2014; Attari et al., 2010; Dieckmann et al., 2021), our study primarily measured perceptions and preferred strategies, not actual behavior. A well-known intention-behavior gap often exists in sustainability contexts. Therefore, future research should move beyond self-reports to examine how our findings translate into practice. This could involve behavioral experiments, the analysis of user log data, or field studies to identify the specific conditions under which users adopt more sophisticated, energy-aware strategies when interacting with AI systems.

Fourth, despite collecting qualitative insights, additional in-depth research could provide more insight into specific antecedents related to the perception of energy demand or personal change processed over time. In this regard, longitudinal research would be beneficial for a process perspective on this matter.

Fifth, this study uses two sorting tasks, which allows for the capture of within-category ordering but does not allow for conclusions about cross-domain variations. More specifically, these results do not show how well individuals are able to juxtapose the energy demand of using AI (e.g., image generation) with the energy demand of a washing machine cycle. Future research could address this limitation by using an integrated sorting task.

Sixth, although we sometimes refer to user strategies as effective or ineffective, this classification relies on technical assumptions, as suggested by the current state of the literature, rather than on direct empirical measurements. The effectiveness of strategies depends on the context and technical implementation. For example, the effectiveness of prompt optimization depends on the user's ability to formulate precise prompts, the model's sensitivity to prompt length, and whether the prompt formulation increases the number of required follow-up interactions or not. In this regard, it is also worth acknowledging that not everyone is equally able to implement or leverage all of the strategies considered here. For example, a sustainable option may come with a premium price, which is not feasible for everyone.

Conclusion

The deployment of GenAI demands a vast amount of resources, which increasingly strains energy supplies (see, e.g., Masanet et al., 2020; Ananthaswamy, 2023). The sustainable long-term use of GenAI requires us to understand how to manage it today. Our research contributes to the broader IS conversation on reconciling digitalization with sustainability (Veit & Thatcher, 2023) by examining these questions from a socio-behavioral perspective, focusing on users of GenAI and how they perceive and respond to GenAI's energy footprint. We find that public awareness of GenAI energy consumption is moderate and nearly on par with awareness of household device energy use. However, the responsibility for GenAI's energy consumption is diffused away from individual users and toward other actors, such as developers and policy-makers, which might undermine personal motivation for pro-environmental behavior (Steg & Vlek, 2009). People often struggle with how to mitigate GenAI's energy consumption and often propose avoiding or cutting back on AI usage, which might be a difficult strategy to implement as GenAI becomes more prevalent. Whether GenAI can be aligned with sustainability goals and whether this technology can be sustained in the long term will also depend on how individuals understand and engage with its resource demands. The results of our study emphasize the importance of the social-behavioral perspective in IS research and can be a valuable point of departure for future investigations in this field.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s12525-026-00908-w>.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data Availability The data and analysis scripts supporting the findings of this study are openly available on the Open Science Framework (OSF) at <https://osf.io/8z4qj>.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adamska, M., Smirnova, D., Nasiri, H., Yu, Z., & Garraghan, P. (2025). Green prompting. arXiv preprint [arXiv:2503.10666](https://arxiv.org/abs/2503.10666)
- Albert, D. A., & Smilek, D. (2023). Comparing attentional disengagement between prolific and MTurk samples. *Scientific Reports*, 13(1), 20574.
- Aljbou, J., Wilson, T., & Patel, P. (2024). *Powering intelligence: Analyzing artificial intelligence and data center energy consumption*. White Paper 3002028905, EPRI, Palo Alto, CA.
- Allcott, H. (2011). Consumers' perceptions and misperceptions of energy costs. *American Economic Review*, 101(3), 98–104. <https://doi.org/10.1257/aer.101.3.98>
- Ananthaswamy, A. (2023). In AI, is bigger always better? *Nature*, 615, 202–205. <https://doi.org/10.1038/d41586-023-00641-w>
- Andor, M. A., Gerster, A., Gillingham, K. T., & Horvath, M. (2020). Running a car costs much more than people think—Stalling the uptake of green travel. *Nature*, 580(7804), 453–455. <https://doi.org/10.1038/d41586-020-01118-w>
- Asensio, O. I., & Delmas, M. A. (2015). Nonprice incentives and energy conservation. *Proceedings of the National Academy of Sciences*, 112(6), E510–E515. <https://doi.org/10.1073/pnas.1401880112>
- Attari, S. Z. (2014). Perceptions of water use. *Proceedings of the National Academy of Sciences*, 111(14), 5129–5134. <https://doi.org/10.1073/pnas.1316402111>
- Attari, S. Z., DeKay, M. L., Davidson, C. I., & Bruine de Bruin, W. (2010). Public perceptions of energy consumption and savings. *Proceedings of the National Academy of Sciences*, 107(37), 16054–16059. <https://doi.org/10.1073/pnas.1001509107>
- Attari, S. Z., DeKay, M. L., Davidson, C. I., & De Bruin, W. B. (2011). Reply to Frederick et al.: Anchoring effects on energy perceptions. *Proceedings of the National Academy of Sciences*, 108(8), E24–E24. <https://doi.org/10.1073/pnas.1019040108>
- Bandura, A., Barbaranelli, C., Caprara, G. V., & Pastorelli, C. (1996). Mechanisms of moral disengagement in the exercise of moral agency. *Journal of Personality and Social Psychology*, 71(2), 364.
- Banh, L., & Strobel, G. (2023). Generative artificial intelligence. *Electronic Markets*, 33, 63. <https://doi.org/10.1007/s12525-023-00680-1>
- Bergquist, M., Thiel, M., Goldberg, M. H., & Van Der Linden, S. (2023). Field interventions for climate change mitigation behaviors: A second-order meta-analysis. *Proceedings of the National Academy of Sciences*, 120(13), e2214851120.
- Bick, A., Blandin, A., & Deming, D. J. (2026). The rapid adoption of generative AI. *Management Science*. <https://doi.org/10.1287/mnsc.2025.02523>, advance online publication.
- Bleher, H., & Braun, M. (2022). Diffused responsibility: Attributions of responsibility in the use of ai-driven clinical decision support systems. *AI and Ethics*, 2(4), 747–761.
- Bourzac, K. (2024). Fixing AI's energy crisis. *Nature*. <https://doi.org/10.1038/d41586-024-03408-z>
- Bran, A. M., Cox, S., Schilter, O., Baldassari, C., White, A. D., & Schwaller, P. (2024). Augmenting large language models with chemistry tools. *Nature Machine Intelligence*, 6, 525–535. <https://doi.org/10.1038/s42256-024-00832-8>
- Brown, B., Juravsky, J., Ehrlich, R. S., Clark, R., Le, Q. V., Re, C., & Mirhoseini, A. (2025). Large language monkeys: Scaling inference compute with repeated sampling. <https://openreview.net/forum?id=0xUEBQV54B>
- Brügger, A., Morton, T. A., & Dessai, S. (2016). "proximising" climate change reconsidered: A construal level theory perspective. *Journal of Environmental Psychology*, 46, 125–142.
- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- Burton-Jones, A., & Straub, D. W. (2006). Reconceptualizing system usage: An approach and empirical test. *Information Systems Research*, 17(3), 228–246. <https://doi.org/10.1287/isre.1060.0096>
- Camilleri, A. R., Larrick, R. P., Hossain, S., & Patino-Echeverri, D. (2019). Consumers underestimate the emissions associated with food but are aided by labels. *Nature Climate Change*, 9(1), 53–58. <https://doi.org/10.1038/s41558-018-0354-z>
- Chen, S. (2025). How much energy will AI really consume? The good, the bad and the unknown. *Nature*, 639(8053), 22–24. <https://doi.org/10.1038/d41586-025-00616-z>
- Chen, Q., Wang, J., & Lin, J. (2025). Generative AI exacerbates the climate crisis. *Science*, 387(6734), 587–587. <https://doi.org/10.1126/science.adt5536>
- Dalén, A., & Krämer, J. (2017). Towards a user-centered feedback design for smart meter interfaces to support efficient energy-use choices. *Business & Information Systems Engineering*, 59(5), 361–373. <https://doi.org/10.1007/s12599-017-0489-x>
- Dauner, M., & Socher, G. (2025). Energy costs of communicating with AI. *Frontiers in Communication*, 10, 1572947. <https://doi.org/10.3389/fcomm.2025.1572947>
- De Groot, J. I., & Steg, L. (2009). Morality and prosocial behavior: The role of awareness, responsibility, and norms in the norm activation model. *The Journal of Social Psychology*, 149(4), 425–449.
- Dedrick, J. (2010). Green IS: Concepts and issues for information systems research. *Communications of the Association for Information Systems*, 27(1), 173–184. <https://doi.org/10.17705/1CAIS.02711>
- Dieckmann, N. F., Gregory, R., Satterfield, T., Mayorga, M., & Slovic, P. (2021). Characterizing public perceptions of social and cultural impacts in policy decisions. *Proceedings of the National Academy of Sciences*, 118(24), e2020491118. <https://doi.org/10.1073/pnas.2020491118>
- Dietz, T. (2014). Understanding environmentally significant consumption. *Proceedings of the National Academy of Sciences*, 111(14), 5067–5068. <https://doi.org/10.1073/pnas.1403169111>
- European Union (2024). Regulation (EU) 2024/1689 of the european parliament and of the council of 13 june 2024 laying down har-

- monised rules on artificial intelligence (artificial intelligence act). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- Eyal, T., Liberman, N., Trope, Y., & Walther, E. (2004). The pros and cons of temporally near and distant action. *Journal of Personality and Social Psychology*, 86(6), 781–795. <https://doi.org/10.1037/0022-3514.86.6.781>
- Feuerriegel, S., Maarouf, A., Bär, D., Geissler, D., Schweisthal, J., Pröllochs, N., & Robertson, C. E., Rathje, S., Hartmann, J., Mohammad, S. M. et al. (2025). Using natural language processing to analyse text data in behavioural science. *Nature Reviews Psychology*, 4(2), 96–111. <https://doi.org/10.1038/s44159-024-00392-z>
- Fiedler, K., Jung, J., Wänke, M., & Alexopoulos, T. (2012). On the relations between distinct aspects of psychological distance: An ecological basis of construal-level theory. *Journal of Experimental Social Psychology*, 48(5), 1014–1021. <https://doi.org/10.1016/j.jesp.2012.03.013>
- Fiedler, K., Jung, J., Wänke, M., Alexopoulos, T., & de Molière, L. (2015). Toward a deeper understanding of the ecological origins of distance construal. *Journal of Experimental Social Psychology*, 57, 78–86.
- Fischer, P., Krueger, J. I., Greitemeyer, T., Vogrincic, C., Kastenmüller, A., Frey, D., Heene, M., Wicher, M., & Kainbacher, M. (2011). The bystander-effect: A meta-analytic review on bystander intervention in dangerous and non-dangerous emergencies. *Psychological Bulletin*, 137(4), 517. <https://doi.org/10.1037/a0023304>
- Frederick, S. W., Meyer, A. B., & Mochon, D. (2011). Characterizing perceptions of energy consumption. *Proceedings of the National Academy of Sciences*, 108(8), E23–E23. <https://doi.org/10.1073/pnas.1014806108>
- Fridgen, G., Körner, M. F., Walters, S., Naujoks, E., Hagen, S., Gottstein, N., & Rüb, B. (2021). Not all doom and gloom: How energy-intensive and temporally flexible data center applications may actually promote renewable energy sources. *Business & Information Systems Engineering*, 63(3), 243–256. <https://doi.org/10.1007/s12599-021-00686-z>
- Gromet, D. M., Kunreuther, H., & Larrick, R. P. (2013). Political ideology affects energy-efficiency attitudes and choices. *Proceedings of the National Academy of Sciences*, 110(23), 9314–9319. <https://doi.org/10.1073/pnas.1218453110>
- Gujral, H., Bremer, C., Perera, D., & Easterbrook, S. (2025). Design for digital sufficiency: Understanding user preferences for more sustainable data centers. *ACM Journal on Computing and Sustainable Societies*, 3(4), 1–31.
- Guo, Z., Li, J., & Ramesh, R. (2023). Green data analytics of supercomputing from massive sensor networks: Does workload distribution matter? *Information Systems Research*, 34(4), 1664–1685. <https://doi.org/10.1287/isre.2023.1208>
- Hansen, J., Sato, M., & Ruedy, R. (2012). Perception of climate change. *Proceedings of the National Academy of Sciences*, 109(37), E2415–E2423. <https://doi.org/10.1073/pnas.1205276109>
- Hedwig, M., Malkowski, S., & Neumann, D. (2009). Taming energy costs of large enterprise systems through adaptive provisioning. In *Proceedings of the 30th International Conference on Information Systems (ICIS 2009)*. Association for Information Systems (AIS), Phoenix, Arizona, USA. <http://aisel.aisnet.org/icis2009/140>
- Hermann, J., Rusche, S., Moder, L., Rückert, M., Heller, F., & Weinhardt, C. (2024). Watt's next? Leveraging process flexibility for power cost optimization. *Business & Information Systems Engineering*, 66(5), 541–563. <https://doi.org/10.1007/s12599-024-00888-1>
- Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., de Las Casas, D., Hendricks, L. A., Welbl, J., Clark, A., Hennigan, T., Noland, E., Millican, K., van den Driessche, G., Damoc, B., Guy, A., Osindero, S., Simonyan, K., Elsen, E., Vinyals, O., Rae, J. W., & Sifre, L. (2022). Training compute-optimal large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*. Curran Associates Inc., Red Hook, NY, USA, NIPS '22.
- Hugging Face (2024). AI Energy Score: Benchmarking Energy Efficiency of AI Models. <https://huggingface.github.io/AIEnergyScore/>. Accessed 07 May 2025
- Humlum, A., & Vestergaard, E. (2025). The unequal adoption of chatgpt exacerbates existing inequalities among workers. *Proceedings of the National Academy of Sciences*, 122(1), e2414972121. <https://doi.org/10.1073/pnas.2414972121>
- International Energy Agency (2024). Electricity 2024: Analysis and Forecast to 2026. <https://www.iea.org/reports/electricity-2024>. Accessed 02 May 2025
- Jegham, N., Abdelatti, M., Elmoubarki, L., & Hendawi, A. (2025). How hungry is AI? Benchmarking energy, water, and carbon footprint of llm inference. arXiv preprint [arXiv:2505.09598](https://arxiv.org/abs/2505.09598)
- Jia, R., Steelman, Z. R., & Reich, B. H. (2017). Using mechanical turk data in IS research: Risks, rewards, and recommendations. *Communications of the Association for Information Systems*, 41(1), 14.
- Jitkritum, W., Narasimhan, H., Rawat, A. S., Juneja, J., Wang, C., Wang, Z., Go, A., Lee, C.-Y., Shenoy, P., Panigrahy, R., Menon, A. K., & Kumar, S. (2025). Universal model routing for efficient llm inference. <https://doi.org/10.48550/arXiv.2502.08773>, [arXiv:2502.08773](https://arxiv.org/abs/2502.08773)
- Karlin, B., Zinger, J. F., & Ford, R. (2015). The effects of feedback on energy conservation: A meta-analysis. *Psychological Bulletin*, 141(6), 1205. <https://doi.org/10.1037/a0039650>
- Kässi, O., Lehdonvirta, V., & Stephany, F. (2021). How many online workers are there in the world? A data-driven assessment. *Open Research Europe*, 1, 53. <https://doi.org/10.12688/openreseurope.13639.4>
- Keller, E., Marsh, J. E., Richardson, B. H., & Ball, L. J. (2022). A systematic review of the psychological distance of climate change: Towards the development of an evidence-based construct. *Journal of Environmental Psychology*, 81, 101822. <https://doi.org/10.1016/j.jenvp.2022.101822>
- Ketter, W., Peters, M., Collins, J., & Gupta, A. (2016). A multiagent competitive gaming platform to address societal challenges. *Management Information Systems Quarterly*, 40(2), 447–460. <https://doi.org/10.25300/MISQ/2016/40.2.09>
- Kim, J., Park, J., & Dedrick, J. (2025). An empirical investigation of cloud computing and environmental performance of nations: Implications for shared responsibility in cloud computing. *Journal of the Association for Information Systems*, 26(5), 1239–1267.
- Klesel, M., & Messer, U. (2026). Good for the planet, bad for me? Intended and unintended consequences of ai energy consumption disclosure. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, CHI '26. <https://doi.org/10.1145/3772318.3791015>
- Kogut, T., Ritov, I., Rubaltelli, E., & Liberman, N. (2018). How far is the suffering? The role of psychological distance and victims' identifiability in donation decisions. *Judgment and Decision Making*, 13(5), 458–466. <https://doi.org/10.1017/S1930297500008731>
- Kotlarsky, J., Oshri, I., & Sekulic, N. (2023). Digital sustainability in information systems research: Conceptual foundations and future directions. *Journal of the Association for Information Systems*, 24, 936–952. <https://doi.org/10.17705/1jais.00825>
- Lacoste, A., Luccioni, A., Schmidt, V., & Dandres, T. (2019). Quantifying the carbon emissions of machine learning. arXiv preprint [arXiv:1910.09700](https://arxiv.org/abs/1910.09700)
- Larrick, R. P., & Soll, J. B. (2008). The MPG illusion. *Science*, 320(5883), 1593–1594. <https://doi.org/10.1126/science.1154983>

- Latane, B., & Darley, J. M. (1968). Group inhibition of bystander intervention in emergencies. *Journal of Personality and Social Psychology*, 10(3), 215.
- Lieberman, N., & Trope, Y. (2008). The psychology of transcending the here and now. *Science*, 322(5905), 1201–1205. <https://doi.org/10.1126/science.1161958>
- Lieberman, N., & Trope, Y. (2014). Traversing psychological distance. *Trends in Cognitive Sciences*, 18(7), 364–369. <https://doi.org/10.1016/j.tics.2014.03.001>
- Li, P., Yang, J., Islam, M. A., & Ren, S. (2025). Making AI less ‘thirsty’. *Communications of the ACM*, 68(7), 54–61. <https://doi.org/10.1145/3724499>
- Loock, C. M., Staake, T., & Thiesse, F. (2013). Motivating energy-efficient behavior with green is: An investigation of goal setting and the role of defaults. *MIS Quarterly*, 37(4), 1313–1332. <https://doi.org/10.25300/MISQ/2013/37.4.15>
- Louviere, J. J., & Islam, T. (2008). A comparison of importance weights and willingness-to-pay measures derived from choice-based conjoint, constant sum scales and best-worst scaling. *Journal of Business Research*, 61(9), 903–911. <https://doi.org/10.1016/j.jbusres.2006.11.010>
- Loy, L. S., & Spence, A. (2020). Reducing, and bridging, the psychological distance of climate change. *Journal of Environmental Psychology*, 67, 101388.
- Luccioni, A. S., & Hernandez-Garcia, A. (2023). Counting carbon: A survey of factors influencing the emissions of machine learning. arXiv preprint [arXiv:2302.08476](https://arxiv.org/abs/2302.08476)
- Luccioni, S., Jernite, Y., & Strubell, E. (2024b). Power hungry processing: Watts driving the cost of ai deployment? In *Proceedings of the 2024 ACM conference on Fairness, Accountability, and Transparency* (pp. 85–99). <https://doi.org/10.1145/3630106.3658542>
- Luccioni, A. S., Strubell, E., & Crawford, K. (2025). From efficiency gains to rebound effects: The problem of jevons’ paradox in AI’s polarized environmental debate. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency* (pp. 76–88).
- Luccioni, S., Gamazaychikov, B., Hooker, S., Pierrard, R., Strubell, E., Jernite, Y., & Wu, C.-J. (2024a). Light bulbs have energy ratings—so why can’t ai chatbots? *Nature*, 632(8026), 736–738. <https://doi.org/10.1038/d41586-024-02680-3>
- Luccioni, A. S., Viguier, S., & Ligozat, A. L. (2023). Estimating the carbon footprint of bloom, a 176b parameter language model. *Journal of Machine Learning Research*, 24(253), 1–15. <https://doi.org/10.5555/3648699.3648952>
- Maliakel, P. J., Ilager, S., & Brandic, I. (2025). Investigating energy efficiency and performance trade-offs in llm inference across tasks and dvfs settings. arXiv preprint [arXiv:2501.08219](https://arxiv.org/abs/2501.08219)
- Mankoff, J. C., Blevins, E., Borning, A., Friedman, B., Fussell, S. R., Hasbrouck, J., Woodruff, A., & Sengers, P. (2007). Environmental sustainability and interaction. In *CHI’07 extended abstracts on Human factors in computing systems* (pp. 2121–2124).
- Masanet, E., Shehabi, A., Lei, N., Smith, S., & Koomey, J. (2020). Recalibrating global data center energy-use estimates. *Science*, 367(6481), 984–986.
- Mildenberger, M., Constantino, S. M., Mahdavi, P., Bergquist, P., De Roche, G., Franzblau, E., Martinez-Alvarez, C., & Sturm, I. (2025). How publics in small-island states view climate change and international responses to it. *Proceedings of the National Academy of Sciences*, 122(30), e2415324122. <https://doi.org/10.1073/pnas.2415324122>
- Mytton, D. (2020). Hiding greenhouse gas emissions in the cloud. *Nature Climate Change*, 10(8), 701–701. <https://doi.org/10.1038/s41558-020-0837-6>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- Patterson, D., Gonzalez, J., Hölzle, U., et al. (2022). The carbon footprint of machine learning training will plateau, then shrink. *Computer*, 55(7), 18–28. <https://doi.org/10.1109/MC.2022.3148714>
- Pearson, A. R., Schuldt, J. P., Romero-Canyas, R., et al. (2018). Diverse segments of the US public underestimate the environmental concerns of minority and low-income americans. *Proceedings of the National Academy of Sciences*, 115(49), 12429–12434. <https://doi.org/10.1073/pnas.1804698115>
- Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70, 153–163.
- Pinski, M., & Benlian, A. (2024). AI literacy for users ? A comprehensive review and future research directions of learning methods, components, and effects. *Computers in Human Behavior: Artificial Humans*, 2(1). <https://doi.org/10.1016/j.chbah.2024.100062>
- Pinski, M., Hofmann, T., & Benlian, A. (2024). AI literacy for the top management: An upper echelons perspective on corporate ai orientation and implementation ability. *Electron Markets*, 34(1). <https://doi.org/10.1007/s12525-024-00707-1>
- Recker, J., Chatterjee, S., Sundermeier, J., & Tarafdar, M. (2025). Digital responsibility: Current perspectives and future directions. *Journal of the Association for Information Systems*, 26(5), 1222–1238. <https://doi.org/10.17705/1jais.00966>
- Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54–63. <https://doi.org/10.1145/3381831>
- Shin, M., Kim, J., van Opheusden, B., & Griffiths, T. L. (2023). Super-human artificial intelligence can improve human decision-making by increasing novelty. *Proceedings of the National Academy of Sciences*, 120(12), e2214840120. <https://doi.org/10.1073/pnas.2214840120>
- Sim, J., Lee, J., & Cho, D. (2023). On the effectiveness of smart metering technology adoption: Evidence from the national rollout in the united kingdom. *Journal of the Association for Information Systems*, 24(2), 555–591. <https://doi.org/10.17705/1jais.00786>
- Siy, C., Do, G., Jung, K., Lee, K. H., & Kim, J. (2025). E-litter: Nudging email usage behavior one byte at a time. *Proceedings of the ACM on Human-Computer Interaction*, 9(5), 1–22.
- Sneffella, B., & Kuperman, V. (2015). Concreteness and psychological distance in natural language use. *Psychological Science*, 26(9), 1449–1460.
- Soderberg, C. K., Callahan, S. P., Kochersberger, A. O., Amit, E., & Ledgerwood, A. (2015). The effects of psychological distance on abstraction: Two meta-analyses. *Psychological Bulletin*, 141(3), 525–548. <https://doi.org/10.1037/bul0000005>
- Stanley, M. L., Sinclair, A. H., & Seli, P. (2020). Intellectual humility and perceptions of political opponents. *Journal of Personality*, 88(6), 1196–1216. <https://doi.org/10.1111/jopy.12566>
- Steg, L., & Vlek, C. (2009). Encouraging pro-environmental behaviour: An integrative review and research agenda. *Journal of Environmental Psychology*, 29(3), 309–317. <https://doi.org/10.1016/j.jenvp.2008.10.004>
- Strubell, E., Ganesh, A., & McCallum, A. (2020). Energy and policy considerations for modern deep learning research. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(09), 13693–13696. <https://doi.org/10.1609/aaai.v34i09.7123>
- Tang, C. (2025). Meta’s hyperscale infrastructure: Overview and insights. *Communications of the ACM*, 68(2), 52–63.
- Thomas, M., & Tsai, C. I. (2011). Psychological distance and subjective experience: How distancing reduces the feeling of difficulty. *Journal of Consumer Research*, 39(2), 324–340. <https://doi.org/10.1086/663772>
- Todorov, A., Goren, A., & Trope, Y. (2007). Probability as a psychological distance: Construal and preferences. *Journal of Experimental Social Psychology*, 43(3), 473–482. <https://doi.org/10.1016/j.jesp.2006.04.002>

- Tong, J., Marx, J., Cui, T., & Turel, O. (2025). Bridging the AI accountability gap: The role of collaboration paradigms and reason responsiveness. In *ECIS 2025 Proceedings*, 2. https://aisel.aisnet.org/ecis2025/human_ai/human_ai/2
- Trinh, T. H., Wu, Y., Le, Q. V., He, H., & Luong, T. (2024). Solving olympiad geometry without human demonstrations. *Nature*, 625(7995), 476–482. <https://doi.org/10.1038/s41586-023-06747-5>
- Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review*, 117(2), 440–463. <https://doi.org/10.1037/a0018963>
- Trope, Y., Liberman, N., & Wakslak, C. (2007). Construal levels and psychological distance: Effects on representation, prediction, evaluation, and behavior. *Journal of Consumer Psychology*, 17(2), 83–95. [https://doi.org/10.1016/S1057-7408\(07\)70013-X](https://doi.org/10.1016/S1057-7408(07)70013-X)
- Varoquaux, G., Luccioni, S., & Whittaker, M. (2025). Hype, sustainability, and the price of the bigger-is-better paradigm in AI. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency* (pp. 61–75).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* (pp. 6000–6010). Curran Associates Inc., Red Hook, NY, USA, NIPS'17.
- Veit, D. J., & Thatcher, J. B. (2023). Digitalization as a problem or solution? Charting the path for research on sustainable information systems: Dj veit, jb thatcher. *Journal of Business Economics*, 93(6), 1231–1253.
- Venkatesh, V., Brown, S. A., & Bala, H. (2013). Bridging the qualitative-quantitative divide: Guidelines for conducting mixed methods research in information systems1. *MIS Quarterly*, 37(1), 21–54.
- Verdecchia, R., Sallou, J., & Cruz, L. (2023). A systematic review of green AI. *WIREs Data Mining and Knowledge Discovery*, 13(4), e1507. <https://doi.org/10.1002/widm.1507>
- vom Brocke, J., Watson, R. T., Dwyer, C., Elliot, S., & Melville, N. (2013). Green information systems: Directives for the IS discipline. *Communications of the Association for Information Systems*, 33, 509–520. <https://doi.org/10.17705/1CAIS.03330>
- Watson, R. T., Boudreau, M. C., & Chen, A. J. (2010). Information systems and environmentally sustainable development: Energy informatics and new directions for the is community. *MIS Quarterly*, 34(1), 23–38. <https://doi.org/10.2307/20721413>
- Watson, R. T., Ketter, W., Recker, J., & Seidel, S. (2022). Sustainable energy transition: Intermittency policy based on digital mirror actions. *Journal of the Association for Information Systems*, 23(3), 631–638. <https://doi.org/10.17705/1jais.00752>
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E. H., Le, Q. V., & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*. Curran Associates Inc., Red Hook, NY, USA, NIPS '22.
- Wendt, C., Kosin, D., Adam, M., & Benlian, A. (2024). Sustainable energy consumption behaviour with smart meters: The role of relative performance and evaluative standards. *Information Systems Journal*, 34(6), 1995–2023. <https://doi.org/10.1111/isj.12520>
- Zampou, E., Mourtos, I., Pramataris, K., & Seidel, S. (2022). A design theory for energy and carbon management systems in the supply chain. *Journal of the Association for Information Systems*, 23(1), 329–371. <https://doi.org/10.17705/1jais.00725>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.