Evaluating Visual Nudges and User Trade-Offs for Sustainable Computing in AI Interfaces

Alain Starke[1,2,*], Samira Dauth[3,†], Tilda Lundkvist[3,†] and Bruce Ferwerda[3]

[1]Amsterdam School of Communication Research, University of Amsterdam, Amsterdam, Netherlands
[2]Department of Information Science and Media Studies, University of Bergen, Bergen, Norway
[3]Department of Computer Science and Informatics, Jönköping University, Jönköping, Sweden

Abstract

The increase in energy and water consumption of artificial intelligence (AI) services, such as GPT, poses significant sustainability challenges. To overcome these challenges, user-centered strategies are necessary to reduce the environmental impact without compromising the user experience. This study investigates how visual nudges influence users' decision making in AI interfaces and explores the trade-offs users make between processing time, output quality, and environmental impact. A between-subjects online experiment (N = 93) using mock-up AI interface scenarios and a choice-based conjoint design was conducted. Participants selected among processing options, generating a specific result of a prompt, with one group exposed to a visual nudge and the other not. The options to choose varied in processing time, output quality and environmental impact. Results show that the visual nudge increases the selection of environment friendlier options, particularly trading off output quality rather than processing time. The study also found evidence of a compromise effect, where users frequently chose moderate options to balance sustainability with usability. Environmental attitudes, measured via the New Ecological Paradigm (NEP) scale, significantly predicted sustainable choices, though they did not moderate the nudge's effect. These findings contribute to sustainable human-computer interaction (HCI) by demonstrating that visual nudges can support greener computing behaviors, especially when aligned with acceptable performance trade-offs.

Keywords

Digital nudging, Nudges, Sustainable computing, AI, AI interface, User experience, Conjoint analysis, NEP scale

1. Introduction

Humans interact with artificial intelligence through AI interfaces that can be used for tasks such as generating information, giving advice, and making decisions [1]. For these interfaces, the deep learning model GPT (Generative Pre-trained Transformer) is used, which is trained to process, analyze and generate text [2]. An example of an AI interface is the chatbot, such as ChatGPT, which in 2025 is the most widely used application worldwide with almost 800 million weekly users [3].

In 2 years, 4.2 to 6.6 billion cubic meters of water are expected to be needed to operate AI systems. To illustrate the impact, this is 4-6 times the annual water withdrawal in Denmark [4]. Today, 2.9 Wh (watt-hour) of electricity is required per request on ChatGPT, which is about 10 times more than what is required for a Google search [5]. Compared to the most used led light bulb (i.e., 10W), it would take 17.4 minutes to consume the same amount of energy [6].

Although individual prompts have a minor impact, the collective use of AI interfaces is likely to have a large ecological footprint. One reason is that most models are not optimized to consider both energy efficiency and model performance [7, 8], particularly when low performance is possible [9]. We hypothesize that if users are given the opportunity to forgo on model performance at the advantage of energy efficiency, a significant number of users will do so.

One possible solution is to use nudges in the design of AI interfaces. Nudges are design changes in an interface that predictably affect user preferences and behavior without limiting choices [10]. For

Joint Proceedings of the ACM UMAP Workshops 2026, UMAP 2026, June 8–11, 2026, Gothenburg, Sweden
*Corresponding author.
†These authors contributed equally.
a.d.starke@uva.nl (A. Starke); bruce.ferwerda@ju.se (B. Ferwerda)
0000-0002-9873-8016 (A. Starke); 0000-0003-4344-9986 (B. Ferwerda)

example, providing energy consumption feedback or setting eco-friendly defaults can encourage users to make more sustainable choices. Although nudges have been used successfully in some digital systems, including recommender systems [11, 12], their potential to promote sustainability in AI-powered tools such as GPT has not been fully explored.

Most of the research regarding sustainable AI focuses on system-level improvements rather than user behavior. Studies have evaluated the carbon footprint of algorithms and highlighted the need for more sustainable design. However, few studies address how user interventions like nudges can influence decision making in AI-powered interfaces [13]. There is a knowledge gap in how digital nudges could be used to inform users about the environmental impact of using AI. Another inherent part of decision making is trade-offs, where choosing one option often comes at the expense of another [14]. It is therefore important to understand how users navigate trade-offs between processing time, output quality, and environmental impact when interacting with AI interfaces.

In this paper, we explore how nudges can be used to promote sustainable behavior when interacting with an AI interface. We focus on evaluating visual cues to see how effective they are in encouraging energy-efficient decisions. Additionally, we investigate users' willingness to make trade-offs between processing time, output quality, and environmental impact. Specifically, this paper contributes an online experiment on decision making in a ChatGPT-like interface, examining whether emphasizing the environmental impact of processing options through visual eco-nudges affects user choices. The study also investigates how trade-off preferences vary across prompt complexity and how environmental attitudes relate to sustainable choices. To guide our research, we formulate the following research questions:

RQ1: To what extent are sustainable computing decisions in an AI interface affected by visual nudges?

RQ2: How do users make trade-offs between processing time and output quality, regarding the sustainability of computing decisions in AI interfaces?

RQ3: How does user awareness of environmental impacts play a role in the effectiveness of nudges?

2. Related Work

Nudge theory provides insight into human decision making by subtly shaping the decision environment without restricting options [10]. It recognizes that human choices are often influenced by cognitive biases, habits, and contextual factors. Central to nudge theory is the concept of choice architecture (i.e., similar to an interface), which illustrates how structuring choices in specific ways can predictably influence decisions [15, 10]. Nudges aim to facilitate better decision making by increasing the appeal, accessibility or salience of preferable options [16].

Visual nudges extend this concept by using visual design elements such as colors, shapes, symbols, or images to guide individuals toward desired behaviors [17, 18]. These visual cues help overcome cognitive biases and decision fatigue, promoting informed and deliberate decisions. An illustrative example is the traffic light imagery used on food packaging to represent environmental impacts, employing green, yellow, and red colors to symbolize low, moderate, and high impacts, respectively [19, 20, 21].

Incorporating nudges within recommender systems aligns naturally with their common goal of influencing user choices without restricting freedom [22]. Traditionally, recommender systems focus on suggesting items based on predicted user preferences [23]. However, recent research emphasizes the addition of environmental impact disclosures to recommendations, as recommender systems have been shown to have a substantial environmental impact [24, 25]. Visual nudges could encourage eco-friendly choices in such a context. For instance, Cossatin et al. [26] found that visual nudges in clothing recommender systems effectively increased purchases of environmentally friendly products. This suggests that similar strategies could be effective within AI systems, where recommendation interfaces can encourage users toward more sustainable computing options, balancing user preferences and environmental considerations.

Trade-offs are an inherent aspect of decision-making, where selecting one option often involves sacrificing another [27]. For example, opting for high-quality recommendations may worsen the impact on a user's privacy [14]. Such trade-offs are particularly relevant in AI systems, where optimizing models for energy efficiency often means accepting reductions in speed or accuracy [13]. Understanding how users navigate these trade-offs is crucial for designing recommender systems capable of effectively balancing user preferences with sustainability goals.

Moreover, user attitudes towards the environment may influence how sustainability nudges operate. The Value-Belief-Norm theory [28] suggests that individuals who highly value environmental protection feel a moral obligation to act sustainably. Similarly, the New Ecological Paradigm (NEP) scale measures ecological worldviews, reflecting beliefs about nature's limitations and human-nature interdependence [29]. In adaptive interfaces, such individual differences can be relevant because the same sustainability cue may not be equally meaningful to all users. We therefore treat environmental concern as a user characteristic that may shape sustainable choices and may inform future personalization of sustainability-aware AI interfaces.

3. Method

3.1. Procedure and Participants

We conducted an online experiment via Qualtrics XM. When opening the survey, participants were informed about the purpose of the study, how the data would be used, and that their responses would remain anonymous. They were also informed that participation was voluntary and that they could withdraw at any time. No personal identifying data were collected. To avoid biasing participants' choices, the nudging manipulation itself was not explicitly disclosed before the choice tasks. The survey consisted of four parts: demographic characteristics, AI-interface preferences, transparency questions, and environmental considerations.

Participants were recruited through social media and university channels and were required to have prior experience using ChatGPT. The final sample consisted of $N = 93$ respondents. The gender distribution was 53 males (57%) and 40 females (43%). In terms of age, the largest proportion of participants (57%) fell within the age group of 25–34, followed by 32.3% aged 18–24. Most participants were students (53.8%), followed by respondents working full-time (35.5%) and part-time (5.4%). In terms of geographic distribution, the most frequently reported country of residence was Germany (59.1%), followed by Sweden (26.9%) and France (9.7%).

3.2. Research Design

Participants were randomly assigned to either a nudge group or a control group. The experiment consisted of nine choice questions: three questions for each of the three prompt types. In every choice question, participants saw three alternative AI processing options and selected the one they would prefer when interacting with an AI chatbot. Each option combined three attributes: processing time, output quality, and environmental impact. The attribute levels varied between slow/moderate/fast processing time, low or reduced/moderate/high output quality, and low/moderate/high environmental impact. The same prompt types and option structure were used in both conditions. In the nudge condition, the three options were shown side by side as chatbot mock-ups with a visual nudge underneath each prompt. In the control condition, participants saw the same type of prompt and option information, but without the visual nudge; one prompt image without the nudge was used for the three options. The order of questions and the selection of option combinations were randomized to reduce order and set effects.

The design of the nudges is shown in Figure 1. We selected a visual nudge because visual cues can draw attention, simplify information, and create associations such as green symbols for eco-friendly behavior. The colors follow a traffic-light metaphor, where green represents low impact, yellow moderate impact, and red high environmental impact. This metaphor was chosen because traffic-light colors provide an immediate and familiar heuristic for interpreting environmental information. The icons use plants in three stages of life to connect the cue to the environment: a living plant for low environmental impact,

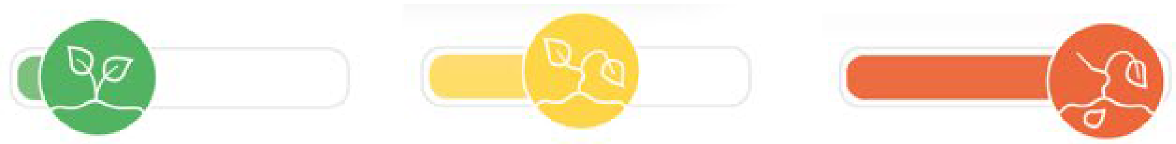

(a) Living plant: low environmental impact (b) Drooping plant: moderate environmental impact (c) Dead plant: high environmental impact

Figure 1: Visual nudge design representing the three stages of life by using a plant.

Table 1
All possible combinations of options in the choice-based conjoint analysis. The different prompts reflect complexity, where dinner recipe is the least complex and image generation the most complex. Processing time would also be indicated in seconds, while quality also described the level of detail. 'Mod' indicates Moderate.

	Processing Time	Output Quality	Environmental Impact
Dinner Recipe	Slow/Mod/Fast	Low/Mod/High	Low/Mod/High
Summary of PDF	Slow/Mod/Fast	Low/Mod/High	Low/Mod/High
Image Generation	Slow/Mod/Fast	Reduced/Mod/High	Low/Mod/High

a drooping plant for moderate environmental impact, and a dead plant for high environmental impact. The plant illustrations were designed using Gestalt principles, where combinations of simple shapes are perceived as coherent visual objects. We did not compare alternative visualizations in this study; rather, we selected this design as a simple, familiar, and interpretable first implementation of a sustainability cue in an AI interface.

We used three prompts that differed in complexity. The first prompt presented to the participants was "Can you make me an easy recipe that takes 15 minutes to make?" This prompt was designed to be the least complex of the three prompts used in the study. The processing time could be slow (20 seconds), moderate (15 seconds) or fast (10 seconds). The output quality ranged from low with low detail in the instructions and containing 3 bullet points, to moderate with 5 bullet points and additional information in a short paragraph and finally high with 7 bullet points and additional information in a clear and well-structured paragraph.

The intermediate prompt presented to participants was "Can you create a summary and list all key points from this PDF, so it's easy to understand?" For this prompt, the processing time varied between slow (55 seconds), moderate (40 seconds), and fast (25 seconds). Low output quality was described as including a few points (3 points) from the PDF and presented with minimal structure and lower detail. Moderate quality was described to include some relevant points (5 points) from the PDF and presented with general structure and moderate detail. The high output quality included 8 key points from the PDF that were presented in a clear and structured format.

The final prompt presented to the participants stated: "Can you generate an image according to the instructions in this PDF?" This task was designed to represent a more demanding AI operation compared to the prior prompts. The processing time was therefore longer and was categorised into slow (2:30 minutes), moderate (2 minutes) and fast (1:30 minutes). Output quality was described as reduced, moderate, or high in terms of complexity and resolution.

This resulted in a 2 (between-subjects: nudge vs. control) × 3 (within-subjects: prompt type) mixed research design.

3.3. Measures and Analysis

We used the features of chosen scenarios as dependent variables. Each chosen scenario could be expressed in terms of its processing time (i.e., slow (-1), moderate (0), or fast (1)), output quality (i.e., low / reduced (-1), moderate (0), or high (1)), and environmental impact (low (-1), moderate (0), or high (1)). We predicted whether features of chosen scenarios differed across conditions, as well as across prompt

Table 2
Specific option profiles used in the choice-based conjoint study. '-' indicates low or reduced, '0' moderate, and '+' high. For each choice question, three of the nine option profiles were randomly selected and displayed simultaneously as the alternatives from which participants chose one preferred option.

Option	1	2	3	4	5	6	7	8	9
Processing Time	-	0	-	+	0	+	-	0	+
Output Quality	-	-	0	-	0	0	+	+	+
Environmental Impact	-	-	-	0	0	0	0	0	+

Table 3
Summary of selected attribute levels across control and nudge conditions.

Selected attribute level	Control	Nudge
Slow processing time	34.11%	31.78%
Moderate processing time	37.73%	39.56%
Low output quality	7.75%	15.11%
High output quality	52.97%	46.67%
Low environmental impact	17.83%	23.78%
High environmental impact	10.85%	8.44%

type and user characteristics. These choices were made in a trade-off context, similar to a choice-based conjoint analysis, which is a well-established method to understand decision-making processes [30].

The final part of the experiment inquired on users' level of environmental concern, which was used a moderating factor. We used the New Ecological Paradigm (NEP) scale, which consisted of fifteen items [31]. In this scale, agreement with the odd-numbered items (after reverse coding) reflected support for the dominant social paradigm (DSP), while agreement with the even-numbered items reflected support for the new environmental paradigm (NEP). The overall score from the NEP survey was assumed to align with the user's environmental concern.

4. Results

4.1. Nudging Effect

The following sections present the collected data across all prompt types and a comparison between the control and nudge groups to explore whether the nudging intervention influenced participants' preferences. Overall, the nudge condition led to more environmentally friendly choices, mainly through greater acceptance of lower output quality rather than slower processing time. Across both conditions, participants most often selected moderate environmental impact and high output quality, suggesting a preference for compromise rather than extreme sustainability-oriented choices.

The data were combined from all nine choice sets and separated by the two condition groups. Table 3 summarizes the selected attribute levels across the control and nudge conditions. A slow processing time was selected in 31.78% of cases in the nudge group and in 34.11% in the control group. A moderate processing time was chosen in 39.56% in the nudge group and 37.73% in the control group. In terms of output quality, low quality was selected in 15.11% for the nudge group and in 7.75% for the control group of all cases. A high output quality was selected in most cases, with 46.67% for the nudge group and 52.97% for the control group. In terms of environmental impact, the nudge group selected low-impact options in 23.78% of all decisions, moderate impact in 67.78%, and high impact in 8.44%. The control group selected low environmental impact in 17.83%, moderate in 71.32%, and high impact in 10.85% of cases.

A regression analysis was conducted to examine whether participants' condition group influenced their choices. There were significant differences for output quality and environmental impact between

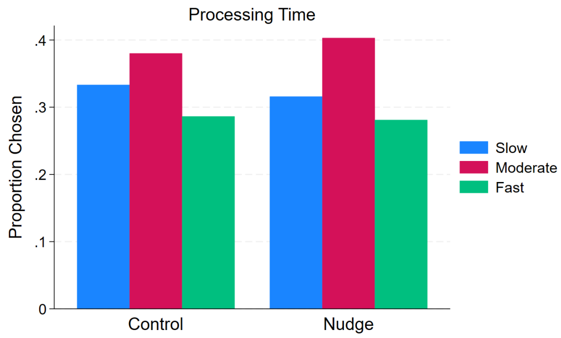

(a) Chosen processing times

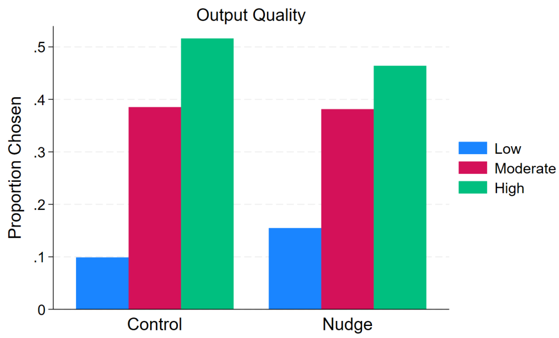

(b) Chosen output quality

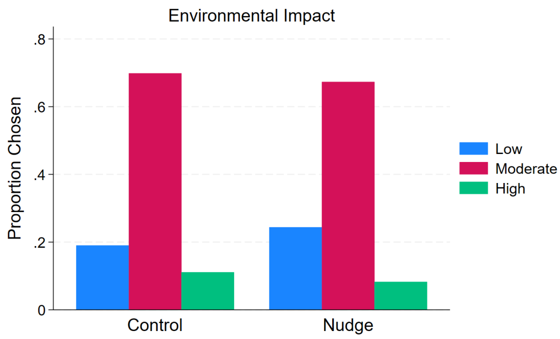

(c) Chosen impacts

Figure 2: Chosen prompting scenarios in terms of the underlying features across nudging and control conditions.

the two groups, while processing time did not show any significant difference. The observed negative coefficient for output quality ($p = .004$, $B = -0.137$) indicates that the nudge group was associated with more choices for lower output quality compared to the control group. As the environmental impact coefficient was in the negative area as well ($p = .026$, $B = -0.084$), a lower environmental impact was chosen in the nudge group more often than in the control group. More detailed logistic regression analyses showed that participants in the nudge condition were 2.21 times more likely to choose low quality output compared to high quality output than those in the control condition.

An overview of the different choices can be inspected in Figure 2. Overall, the figure shows that moderate environmental impact was chosen most often, while high output quality was also selected frequently.

4.2. Trade-offs

Overall, a slow processing time was selected in 32.86% of cases, moderate in 38.71%, and fast in 28.43% of all cases. For output quality, low quality was selected the least, with 11.71% of all decisions made, moderate was chosen in 38.71% of cases, and high output quality was selected the most, in 49.58% of cases. A low environmental impact was chosen in 21.03% of all cases. A moderate environmental impact was preferred and was selected in 69.41% of cases, while a high environmental impact was chosen the least (9.56%).

To examine more closely which options were chosen together and which trade-offs were made preferably, a crosstabs analysis was conducted. The results demonstrate that participants traded processing time rather than output quality. The populations for the processing time choices were more evenly distributed across the different options than for output quality, where almost half of the population chose high output quality ($n = 415$). The combination chosen most was high output quality and moderate processing time ($n = 187$), presenting a willingness to trade time. Furthermore, many participants chose slow processing time with high output quality ($n = 148$). The low-quality output, regardless of processing time, was chosen the least ($n = 98$), underlining the willingness to rather trade off processing time. The majority of participants were not willing to choose an extreme trade-off with a slow processing time and a low output quality ($n = 21$).

4.3. Trade-Offs for Different Prompt Types

Additionally, to investigate whether there were different trade-off preferences across prompt difficulty, a crosstabs analysis was conducted for the choices made over the three prompts.

For the simple recipe prompt, most choices were made for moderate output quality ($n = 125$), and low output quality was chosen the least ($n = 36$). The most chosen combination included high output quality and slow processing time ($n = 49$), closely followed by high output quality and moderate processing time ($n = 47$). Overall, participants traded both output quality and processing time for a better environmental impact.

For the PDF summary prompt, most choices were made for high output quality ($n = 157$). The most chosen combinations were high output quality with moderate processing time ($n = 72$) and slow processing time ($n = 51$). These results show a clear preference for trading time rather than output quality.

For the image generation prompt, the most chosen combination included high output quality and moderate processing time ($n = 68$). The second most chosen combination was moderate output quality and fast processing time ($n = 51$). For a complex task like image generation, participants were willing to trade off options, but rather processing time than output quality.

4.4. Environmental Concern

The scores on the NEP scale in this sample ($N = 93$) ranged from 39 to 64, out of 75, with a mean score of 50.61 and a standard deviation of 5.435. Both groups displayed similar distributions of NEP scores, suggesting that baseline environmental attitudes were relatively consistent between conditions.

A series of linear regression analyses were conducted to test whether environmental attitudes (NEP scores) moderated the effect of the nudge condition group on the option selection behavior. For the selection of processing times (slow, moderate, fast), output quality (low, moderate, high), and environmental impact (low, moderate, high), no significant interaction effects between the nudge condition and the NEP scores were found (all $p > .25$), indicating no evidence of moderation.

However, environmental attitudes influenced choices in expected ways. For processing time, each increase in NEP score increased the odds of selecting slow over fast processing time by 7% and moderate over fast processing time by 4%. No significant effects were found for output quality. For environmental impact, each increase in NEP score increased the odds of selecting a low environmental impact option over a high environmental impact option by 15%, and the odds of selecting a moderate environmental impact option over a high environmental impact option by 8.6%. These results indicate that participants with stronger pro-environmental attitudes were more likely to choose environmentally friendlier options, independent of the nudge.

4.5. Transparency

After selecting their preferred AI-interface options, participants were asked to share their level of agreement with two transparency statements about their experiences. The first statement was "It was easy to understand what the environmental impact for each scenario was." The Anova results showed no statistical difference between the two condition groups. The mean values were 3.63 for the control group and 3.88 for the nudge group.

The second item stated: "I considered the environmental impact during my decision-making process." Again, the Anova analysis showed no significant difference between the condition groups. The mean values were 3.74 for the control group and 3.56 for the nudge group.

A correlation analysis showed that for the item "It was easy to understand what the environmental impact for each scenario was," no statistically significant correlations were found. For the second item, “I considered the environmental impact during my decision-making process,” significant positive relationships were found with the selection of slow processing time ($r = 0.363$, $p < 0.001$), low environmental impact ($r = 0.511$, $p < 0.001$), and total NEP scores ($r = 0.313$, $p = 0.002$).

5. Discussion

5.1. Effectiveness of the Visual Nudge

Overall, the results indicate that introducing a visual cue as a nudge influenced user choices toward more sustainable options. Participants in the nudge group were more likely to choose lower environmental impact options compared to the control group. These lower environmental impact choices were mainly influenced by the acceptance of lower output quality in the nudge group. Users in the nudge group were about as likely to choose slow, moderate, or fast processing as those in the control condition, with no statistically significant differences observed. This suggests that many users continued to prioritise shorter processing times for AI tasks regardless of the nudge.

The findings show that visual nudges can support greener computing behaviors, particularly when trade-offs remain acceptable. At the same time, the effect should be interpreted as modest rather than transformative.

5.2. Trade-Off Preferences

The trade-offs between environmental impact, processing time, and output quality illustrate the complexity of sustainable interaction design. Participants often compromised. They were willing to trade processing time more often than output quality, and they frequently selected moderate options. This suggests that users balance convenience and sustainability and are willing to compromise when the cost in user experience remains acceptable.

At the same time, participants were generally not willing to trade off both processing time and output quality to an extreme at the same time. This suggests that there is a limit to how far users prioritise sustainability over a satisfying experience and result.

5.3. The Role of Environmental Concern

Participants’ environmental attitudes shaped their decision-making mostly independent of the nudge. Results showed that participants with higher NEP scores were more likely to choose environmentally friendlier options regardless of condition group. This indicates that personal awareness and values can drive sustainable behavior in AI interfaces even in the absence of a nudge.

However, the data did not show a statistically significant interaction between environmental awareness and the nudge condition on the choices made. In other words, the effectiveness of the visual nudge was not influenced by the participants’ environmental attitude.

5.4. Limitations

This study has several limitations. First, participants were asked to imagine using an AI chatbot while interacting with mock-up chatbot interfaces. Although the scenarios were designed to resemble realistic AI use, this is not identical to actual interaction with an implemented AI system, and the results may therefore not fully represent real behavior. Second, the processing-time and output-quality values were hypothetical and were not calibrated against measured energy or water consumption from deployed AI systems. The attribute levels were intentionally kept within a plausible everyday-use range, but more extreme contrasts, such as very short versus very long processing times, may lead to different choices. Third, an important real-life behavior was not taken into account: users can send a prompt multiple times when they are unsatisfied with the output. A low-impact option that requires repeated prompting may therefore be less sustainable in practice than a higher-impact option that succeeds immediately. Fourth, the generalizability of the findings may be influenced by the sample population, because the study mainly included students and younger participants and does not fully represent all AI-interface users. Fifth, the environmental survey relied on self-report data, which may not fully reflect real-world behavior. Finally, because the effectiveness of nudging is highly context dependent, results in a mock-up AI interface could differ from results in an implemented AI interface.

6. Conclusion

The findings indicate that visual nudges can influence users' choices toward more environmentally friendly options, but the effect is modest. Participants in the nudge group were more likely to choose sustainable options than those in the control group, particularly low-impact options alongside lower-quality options. However, the nudge did not guide participants to choose slower processing-time options compared to the control group.

Regarding trade-offs, participants often compromised. They opted for the middle ground when facing decisions that varied in output quality, processing time, and environmental impact. More specifically, users were more willing to trade processing time than output quality, and they were generally not willing to make extreme trade-offs on both dimensions at the same time.

Environmental attitudes shaped participant decision making independent of nudging effectiveness. Participants with higher environmental concern consistently made greener choices, but environmental concern did not moderate the nudge's effect. Overall, the study provides empirical evidence that many users are open to being guided toward more sustainable computing practices, as long as the trade-offs remain reasonable.

6.1. Practical Implications

The results of this study could have practical implications for the industries of artificial intelligence and sustainability. Most users were willing to be guided towards sustainable computing practices to a certain extent. The trade-offs available should be reasonable, because that can lead to users sacrificing processing time or output quality for an environmental benefit. This means that companies that run AI chatbots could inform users about the environmental impact of their usage and give them the option to make trade-offs. If users choose lower-impact settings instead of the standard experience, this could reduce the amount of water and energy required by data centres running AI programs.

6.2. Scientific Implications

This study advances scientific understanding in the combination of digital nudging, sustainable human-computer interaction, and user experience. By demonstrating that visual nudges or just the provision of environmental information can modestly influence users' decisions in an AI interface context, the research extends behavioral economics and nudging theory into a relatively underexplored area. The usage of a choice-based conjoint analysis in a UX setting offers a framework for quantifying user preferences and trade-offs. Moreover, the integration of individual difference measures, particularly the New Ecological Paradigm (NEP) scale, adds theoretical depth by showing that pro-environmental attitudes are significantly associated with sustainable choice behavior. This suggests that sustainability should not just be treated as a system-level goal but as a user-centred design factor.

6.3. Future Research

In this study, only the visual nudge type was explored. For future research, it would be relevant to explore other nudge types in the scenario of influencing sustainable decision-making when interacting with AI interfaces. It would also be relevant to explore diverse populations by including a larger sample size from a wider range of countries. To reach more accurate results, studies could be performed that observe participants in realistic scenarios over a longer period. Real AI platforms could also be used instead of prototypes. Another important direction is to explore how users act when they are unsatisfied with an AI output, and to what extent they resend prompts, since sending a lower-impact prompt multiple times could be less sustainable than sending a higher-impact prompt once.

6.4. Implications for Adaptive AI Interfaces

Although this study focused on AI-interface choices rather than recommender-system use directly, the findings are relevant for adaptive and personalized interfaces that present users with computational options. Environmental impact could be treated as one criterion in such interfaces, alongside processing time and output quality. The observed trade-off patterns suggest that users may be more receptive to sustainability-oriented options when these preserve acceptable output quality and avoid extreme sacrifices in usability.

The role of environmental concern also points to a personalization opportunity. Participants with higher NEP scores were more likely to choose environmentally friendlier options regardless of condition, suggesting that adaptive systems could allow users to express sustainability preferences or learn from previous choices. However, such personalization should be framed as an extension of the current findings rather than as a direct result tested in this study.

6.5. Design Considerations

The findings suggest several design considerations for eco-conscious AI interfaces. First, environmental impact information should be made visible at the point of choice, because participants did respond to sustainability cues. Second, comparative visual cues, such as traffic-light colors and simple environmental iconography, may help users interpret impact levels without requiring detailed technical knowledge. Third, sustainability options should preserve reasonable usability: participants were generally willing to compromise, but not to accept extreme losses in both processing time and output quality. These considerations should be treated as preliminary because only one visual nudge design was tested.

6.6. Ethical Considerations

Although visual nudges have proven to be effective in steering choices, they also bring ethical nuances around autonomy and perceived manipulation. If an eco-option noticeably slows down an AI task without any explanation, users may feel misled or frustrated, undermining trust. To mitigate this, interfaces should accompany every sustainability cue with a concise rationale. For example, one could add "This option uses 30% less energy by loading one extra second", which would be an explanation that helps users understand the benefit-cost balance. Moreover, ethical transparency demands that higher-impact, higher-performance choices remain equally visible and accessible; nudges should inform rather than coerce.

6.7. Future Research

Future research should test additional nudge types and compare them with the visual cue used in this study. Studies should also include more diverse populations and examine realistic AI use over longer periods. In particular, future work should use implemented AI platforms and measured energy or water estimates to calibrate the trade-offs more accurately. Another important direction is to study repeated prompting behavior: if users choose a lower-impact option but need to resend prompts several times, this may offset or even reverse the intended sustainability benefit.

Declaration on Generative AI

During the preparation of this work, the author(s) Grammarly in order to: Grammar and spelling check

References

[1] A. Sundar, T. Russell-Rose, U. Kruschwitz, K. Machleit, The ai interface: Designing for the ideal machine-human experience, Computers in Human Behavior 165 (2025) 108539.

[2] G. Yenduri, et al., Gpt (generative pre-trained transformer) – a comprehensive review on enabling technologies, potential applications, emerging challenges, and future directions, arXiv preprint arXiv:2305.10435v2 (2023).

[3] S. Singh, Chatgpt statistics 2025 – dau &mau data (worldwide), 2025. URL: https://www.demandsage.com/chatgpt-statistics/.

[4] P. Li, J. Yang, M. Islam, S. Ren, Making ai less "thirsty": Uncovering and addressing the secret water footprint of ai models, arXiv preprint arXiv:2304.03271 (2023).

[5] International Energy Agency, Electricity 2024: Analysis and forecast to 2026 (2024).

[6] J. Marsh, How many watts does a light bulb use?, 2024. URL: https://www.energysage.com/electricity/house-watts/how-many-watts-does-a-light-bulb-use/.

[7] P. J. Maliakel, S. Ilager, I. Brandic, Investigating energy efficiency and performance trade-offs in llm inference across tasks and dvfs settings, arXiv preprint arXiv:2501.08219 (2025).

[8] J. Stojkovic, E. Choukse, C. Zhang, I. Goiri, J. Torrellas, Towards greener llms: Bringing energy-efficiency to the forefront of llm inference, arXiv preprint arXiv:2403.20306 (2024).

[9] G. Wilkins, S. Keshav, R. Mortier, Hybrid heterogeneous clusters can lower the energy consumption of llm inference workloads, in: Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems, 2024, pp. 506–513.

[10] R. H. Thaler, C. R. Sunstein, Nudge: Improving decisions about health, wealth, and happiness, Constitutional Political Economy 19 (2008) 356–360.

[11] A. S. Oonk, A. D. Starke, M. C. Willemsen, Goal setting and signposting in a rasch-based recommender system to promote household energy conservation, Energy Efficiency 18 (2025) 1–22.

[12] A. Starke, M. Willemsen, C. Snijders, Promoting energy-efficient behavior by depicting social norms in a recommender interface, ACM Transactions on Interactive Intelligent Systems (TiiS) 11 (2021) 1–32.

[13] G. Spillo, A. De Filippo, C. Musto, M. Milano, G. Semeraro, T. d. Noia, P. Lops, J. Basilico, Towards sustainability-aware recommender systems: Analyzing the trade-off between algorithms performance and carbon footprint, in: Proceedings of the 17th ACM Conference on Recommender Systems, 2023, pp. 856–862.

[14] S. Berning, V. Dunning, D. Spagnuelo, T. Veugen, J. van der Waa, The trade-off between privacy & quality for counterfactual explanations, in: Proceedings of the 17th International Conference on Availability, Reliability and Security, 2024, pp. 1–9.

[15] E. J. Johnson, S. B. Shu, B. G. Dellaert, C. Fox, D. G. Goldstein, G. Häubl, R. P. Larrick, J. W. Payne, E. Peters, D. Schkade, et al., Beyond nudges: Tools of a choice architecture, Marketing letters 23 (2012) 487–504.

[16] S. Thomas, A. Arno, The efficacy of nudge theory strategies in influencing adult dietary behaviour: a systematic review and meta-analysis, BMC Public Health 16 (2016).

[17] S. Banker, M. Miller, G. Voichek, D. Goor, T. Makov, Prosocial nudges and visual indicators increase social distancing, but authoritative nudges do not, Proceedings of the National Academy of Sciences (2022).

[18] A. D. Starke, M. C. Willemsen, C. Trattner, Nudging healthy choices in food search through visual attractiveness, Frontiers in Artificial Intelligence 4 (2021) 621743.

[19] A. Angelsen, A. D. Starke, C. Trattner, Healthiness and environmental impact of dinner recipes vary widely across developed countries, Nature Food 4 (2023) 407–415.

[20] J. Morton, A guide to color symbolism (color voodoo #1), COLORCOM (1997).

[21] L. Muller, A. Lacroix, B. Ruffieux, Environmental labelling and consumption changes: A food choice experiment, Environmental and Resource Economics 73 (2019) 871–897.

[22] M. Jesse, D. Jannach, Digital nudging with recommender systems: Survey and future directions, Computers in Human Behavior Reports 3 (2021) 100052.

[23] D. Jannach, M. Zanker, A. Felfernig, G. Friedrich, Recommender systems: an introduction, Cambridge University Press, 2010.
[24] G. Spillo, A. G. Valerio, F. Franchini, A. De Filippo, C. Musto, M. Milano, G. Semeraro, Recsys carbonator: Predicting carbon footprint of recommendation system models, in: International Workshop on Recommender Systems for Sustainability and Social Good, Springer, 2024, pp. 98–110.
[25] T. Vente, L. Wegmeth, A. Said, J. Beel, K. Verbert, P. Lops, M. Zanker, B. London, From clicks to carbon: The environmental toll of recommender systems, in: 18th ACM Conference on Recommender Systems, 2024, pp. 580–590.
[26] A. Cossatin, N. Mauro, L. Ardissono, Promoting green fashion consumption through digital nudges in recommender systems, IEEE Access 12 (2024) 6812–6829.
[27] Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services, Trade-off, 2025. URL: https://www.ipbes.net/glossary/trade.
[28] R. Jadkowski, Exploring driving forces of greener shopping behavior (2021).
[29] M. Anderson, New ecological paradigm (nep) scale, Berkshire Encyclopedia of Sustainability 6 (2012) 260–262.
[30] D. Raghavarao, J. B. Wiley, P. Chitturi, Choice-based conjoint analysis: models and designs, CRC Press, 2011.
[31] R. E. Dunlap, K. D. Van Liere, The “new environmental paradigm”, The journal of environmental education 9 (1978) 10–19.