

**The following paper was researched and written
as part of LSE GROUPS 2025.**

LSE GROUPS takes place during the final fortnight of the LSE Spring Term. Undergraduate students are placed in small groups; these are cross-year, interdisciplinary, and group members do not know one another in advance. Each group must then devise its own research question and carry out every stage of a small-scale research project, in less than two weeks.

LSE GROUPS is part of the LSE commitment to students learning through enquiry and developing the skills needed for knowledge creation.

The overall theme of LSE GROUPS 2025 was *Visions for the Future*

This paper was submitted on the final Thursday afternoon of the project. Students then presented their work at a conference, on the closing Friday.

[More information on LSE GROUPS, and other papers.](#)

Papers are presented as submitted by the students, without corrections.

London School of Economics and Political Science

Knowing but not acting: Exploring the disconnect between environmentalism and AI consumption

Group 2

**Vrinda Parekh, Suha Farook, Satyasree Vangoor, Mahira Haque, Phool Ashraf,
Charikleia Giagkozoglou, Vera Povaiah**

ABSTRACT

Despite thorough documentation of generative AI's environmental impact, end-users remain largely unaware of the environmental costs. The widespread adoption of generative AI by university students is particularly concerning in light of this awareness gap. Existing models consider the drivers of AI usage through a framework of consumer-specific constructs, performance, effort expectancy, social influence, and facilitating conditions. However, this framework neglects sustainability-oriented factors. Drawing on environmental behaviour research highlighting the Value-Action Gap – a model explaining the disconnect between pro-environmental attitudes and behaviours through psychological and structural barriers – this study proposes an extension of existing AI use models by uncovering previously excluded relevant factors. Using data collected from 26 semi-structured interviews of university students from the UK, US, and the Netherlands, this exploratory study reveals how increased awareness of Generative AI's environmental impacts relates to reduced student consumption of it. Specifically, students revealed a preference for limiting their personal use of the technology in the face of environmental consequences, while academic use remained largely unchanged. Discussion of these findings illustrates potential directions for future research and university policies surrounding AI and environmental education.

Keywords: Environmentalism, AI consumption, Exploratory research, Generative AI, University students, Qualitative Research

Acknowledgements:

We are extremely grateful to our supervisor Kristof Horvath, who made this work possible. His continuous guidance, time, and expertise has greatly enhanced the quality of this work.

Table of contents

Introduction	3
Literature Review	3
Methodology	6
Discussion/Empirical Analysis	8
Limitations and Ethical Considerations.....	16
Conclusion.....	17

Introduction

As of 2025, 92% of university students report using AI, a dramatic increase from 66% in 2024 (HEPI, 2025). As this technology becomes increasingly embedded in university life, there is an increased imperative to understand the factors that contribute to AI's rapid adoption. Previous explorations of these factors have primarily drawn from consumer theory and existing technological adoption models. However, one critical factor remains underexplored: environmentalism. Despite growing discourse on AI's environmental consequences, the extensive literature discussing student environmental behaviour has not been sufficiently integrated into discussions of AI adoption. This gap in understanding is even more concerning in the face of consistent evidence that young people are "very concerned about the environment" (Hernandez-Arriaza, et al., 2023). To offer a more comprehensive understanding of student's decisions to use AI, this exploratory study answers the following question:

How do university students perceive the environmental impacts of Generative AI, and how might these perceptions relate to their attitudes toward climate change and their usage of AI?

Literature Review

Generative AI refers to a subset of artificial intelligence capable of creating new content based on patterns learned from large datasets. This includes tools such as ChatGPT, Copilot, and Claude, which are increasingly integrated into students' academic routines (Dilmegani, 2023, as cited in Chan & Colloton, 2024). In this study, *environmental impact* refers to the measurable effects that AI deployment and usage exert on natural ecosystems, specifically through energy consumption, carbon emissions, and freshwater usage, particularly during the operation and cooling of data centres. These environmental costs largely remain invisible to end-users.

The rapid integration of generative AI into higher education has intensified concern over its broader environmental implications. According to the 2025 HEPI survey, 92% of students now use AI, up from 66% in 2024. Furthermore, 88% of students report using generative AI for assessment-related tasks, and 18% have directly incorporated AI-generated text into academic work (HEPI, 2025). These figures reveal the extent to which generative AI has become embedded in student routines, yet this uptake is occurring largely without awareness of the associated environmental toll.

Training large language models like GPT-3 demands considerable energy: its training required 1,287 megawatt hours (MWh) of electricity and produced 502 tonnes of CO₂ equivalent emissions (Patterson et al., 2022). Even beyond initial training, the ongoing use of generative AI is resource-intensive – each ChatGPT query is estimated to consume ten times the energy of a standard Google search (IEA, 2024). Generative AI is typically hosted in vast data centres, which significantly contribute to carbon emissions and freshwater depletion. For instance, in 2023 alone,

Google's data centres withdrew 29 billion litres of freshwater, consuming over 23 billion litres through evaporative cooling, 80% of which was potable (Li et al., 2025).

UTAUT2 Framework and Its Extension

To explore students' adoption of generative AI, this study draws on the Unified Theory of Acceptance and Use of Technology 2 (UTAUT2). Developed by Venkatesh et al. (2012), UTAUT2 extends the original UTAUT model by incorporating consumer-specific constructs – hedonic motivation, price value, and habit – alongside the original factors: performance expectancy, effort expectancy, social influence, and facilitating conditions (Venkatesh et al., 2003). UTAUT2 has demonstrated robust explanatory power in various higher education contexts (Zhu et al., 2024; Gansser & Reich, 2021; Tseng et al., 2019; Zacharis & Nikolopoulou, 2022).

However, despite its utility, UTAUT2 does not explicitly account for sustainability-oriented behaviours. There is a notable absence in the literature linking environmental considerations with AI usage, both within and outside of this framework. In response to this gap, this research adopts an exploratory stance, suggesting that it may be valuable to consider additional constructs such as *Environmental Awareness (EA)* and *Environmental Concern (EC)*, as outlined below.

Environmental awareness refers to an individual's understanding of the ecological consequences of using generative AI tools. *Environmental concern* denotes the emotional and moral response to this awareness, such as guilt, anxiety, or perceived responsibility. By incorporating these constructs, this research aims to explore whether – and how – students' sustainability-related perceptions shape their engagement with AI tools.

This extension is not intended as a definitive revision of UTAUT2, but as a working model for inquiry. Rather than testing a fixed set of causal pathways, this study is open to identifying possible associations between environmental attitudes and technology use, as well as uncovering other relevant factors not currently accounted for in existing frameworks. Semi-structured interviews will provide the primary data, allowing space for participants to raise issues or considerations beyond those predefined in UTAUT2 or environmental psychology.

Environmental Psychology

To support this theoretical exploration, we turn to environmental psychology, which offers insight into the internal and contextual factors shaping pro-environmental behaviours (PEBs). Steg et al. (2014) define PEBs as actions that enhance environmental quality, whether or not they stem from explicit environmental intent. Janmaimool and Khajohnmanee (2019) distinguish between direct PEBs (yielding tangible environmental improvements) and indirect PEBs (which influence

broader systems and practices). In this framework, reducing or moderating AI usage – even passively – could be considered a form of direct PEB.

The literature identifies several internal drivers of PEBs. Intentions are often considered the strongest predictor of environmentally relevant behaviour (Klöckner, 2013), acting as a bridge between environmental attitudes and low-cost actions (Kollmuss & Agyeman, 2002). Social norms, personal values, and perceived moral obligations also play a significant role (Steg et al., 2014). In particular, *perceived effectiveness* – the belief that one’s actions can make a meaningful difference – has been found to influence whether individuals act on their environmental knowledge (Vicente-Molina et al., 2013).

Although *environmental knowledge* is often assumed to correlate with PEBs, this relationship is not always straightforward. For example, in a Thai university study, students who completed an environmental course expressed stronger environmental values, but their actual behaviours remained largely unchanged (Janmaimool & Khajohnmanee, 2019). This is consistent with other findings that *awareness alone does not guarantee behavioural change* (Kollmuss & Agyeman, 2002; Vicente-Molina et al., 2013).

This disjunction is commonly referred to as the *attitude-behaviour gap* or *value-action gap*. It highlights how rational decision-making models often underestimate the influence of structural, emotional, and habitual constraints (Blake, 1999). Barriers such as routine habits, emotional detachment from environmental problems, and feelings of inefficacy can prevent individuals from acting on their stated values (Kim et al., 2018; Janmaimool & Khajohnmanee, 2019). Additionally, infrastructure and institutional support – or lack thereof – can significantly mediate pro-environmental engagement.

Given that many individuals express concern for the environment but do not consistently make choices aligned with these values, this study is particularly interested in exploring how students *perceive* the environmental implications of their AI use and whether these perceptions inform decision-making. This line of inquiry is timely, especially in a context where AI usage is rapidly increasing while its environmental impacts remain largely abstract or invisible to end-users.

Research Orientation

This research does not seek to establish causal relationships, but to explore the *potential relevance* of environmental attitudes within existing models of technology adoption. The working expectation is that environmental awareness and concern may be associated with students’ use of generative AI tools, but the strength and nature of this association remain open to investigation.

Crucially, the study also remains open to the emergence of additional, previously unconsidered factors that may shape student decision-making. To this end, qualitative interviews are employed not only to probe how students interpret environmental aspects of AI usage but also to identify other values, motivations, or constraints that may influence adoption.

By situating this research at the intersection of UTAUT2 and environmental psychology, the study contributes to a more holistic understanding of technology adoption in a time of ecological urgency. It invites reflection on the limits of current theoretical frameworks and encourages future research to take seriously the environmental dimensions of digital life.

Methodology

This study draws on qualitative data collected through semi-structured interviews with university students in the United Kingdom, the United States and the Netherlands. In the context of a rapidly evolving technological landscape, where generative AI tools are being adopted at unprecedented rates and with little consensus around best practices or ethical standards, a qualitative, exploratory approach was essential. Rather than attempting to quantify behaviour or test pre-established relationships, this study aims to explore how students make sense of generative AI, including how they reflect on its value, consequences, and appropriateness in both academic and personal contexts.

The decision to use semi-structured interviews reflects the complexity and fluidity of the topic. Students' use of generative AI involves subtle, context-specific decision-making shaped by a wide range of factors, including academic norms, personal ethics, peer influence, technological curiosity, and broader social narratives. Furthermore, much of the public discourse around generative AI is marked by hype, anxiety, and speculation. In such a setting, fixed survey instruments may be ill-suited to capture the uncertainty, ambivalence, or improvisational reasoning that often accompanies students' actual use of these tools.

Semi-structured interviews provided a flexible yet focused format for exploring students' interpretations and experiences. This method allowed us to begin with guiding themes – such as motivations for AI use, perceived risks, and awareness of environmental impact – while remaining open to emergent topics and unexpected perspectives. The conversational structure created space for participants to introduce their own frames of reference, values, and dilemmas, offering insight into how meaning is constructed rather than assumed.

This open-ended design was especially well-suited to our research objective: to understand how students weigh the potential benefits and trade-offs of using generative AI in daily life. In particular, we were interested in how – or whether – environmental considerations enter into their decision-making processes around their usage of generative AI, and how such considerations interact with other factors like academic pressure, convenience, or social acceptability. Given the

absence of prior research on these intersections, the study leaves room to identify not only expected associations (e.g., between environmental concern and usage patterns) but also novel variables or influences not currently accounted for in existing frameworks such as UTAUT2.

A prompt was provided to measure the change in AI usage among participants, factually detailing the common environmental consequences of AI:

- A ChatGPT query consumes about five times more electricity than a simple web search.
- By 2026, AI could use as much electricity as a small country like Belgium
- AI requires water to cool down its servers. Increased AI demand for water is forecast to amount to four to six times the annual water consumption of Denmark by 2030.
- A short conversation of 20-50 questions and answers with ChatGPT costs half a litre of fresh water.
- Generative AI count accounts for up to 5 million metric tons of e-waste by 2030

Figure 1: Prompt Shown to Participants

Qualitative interviews were also essential given the potentially sensitive nature of the topic. Issues such as AI-assisted cheating, dependency on technology, or disregard for ecological consequences may generate discomfort or moral conflict. A dialogical, non-judgmental interview setting allowed participants to explore their own contradictions, hesitations, and evolving views, providing a richer understanding of how students navigate the ethical and emotional complexity surrounding AI use. This format encouraged disclosure of attitudes or behaviours that might remain unspoken in more structured or evaluative research contexts.

Participants were recruited using a combination of convenience and snowball sampling. Drawing on institutional networks and peer referrals, we engaged students from diverse academic disciplines at UK, US and Dutch universities. This sampling strategy, while non-random, was appropriate for the aims of early-stage exploratory research, which prioritises depth of insight over statistical generalisability. We intentionally sought a range of individuals with differing levels of usage of AI, as their experiences offered a relevant and situated entry point into the phenomena under study.

In total, we conducted 26 semi-structured interviews. These interviews provided detailed accounts of students' academic and non-academic uses of generative AI, including instances of creative experimentation, emotional support, productivity management, and strategic shortcuts. By paying attention to these informal or non-institutionalised uses, the study moves beyond the dominant focus on academic misconduct or classroom integration. This broader lens makes it possible to uncover the logics, justifications, and tensions that shape how students engage with generative AI in everyday life.

Overall, the use of qualitative interviews enabled an exploratory, grounded investigation into how students understand and enact their relationship with generative AI – and what role, if any, environmental considerations play in those practices. The method allowed us to surface latent themes, identify gaps in existing theoretical models, and generate propositions for further research, rather than confirm pre-existing assumptions.

Discussion/Empirical Analysis

The following graphs display the breakdown of various demographic characteristics of our participants, as well as data relating to their use of AI.




Figure 2: Self-Identified Gender of Participants




Figure 3: Participant's Year of Study at University

AI Platforms

"What generative AI platforms do you use?"

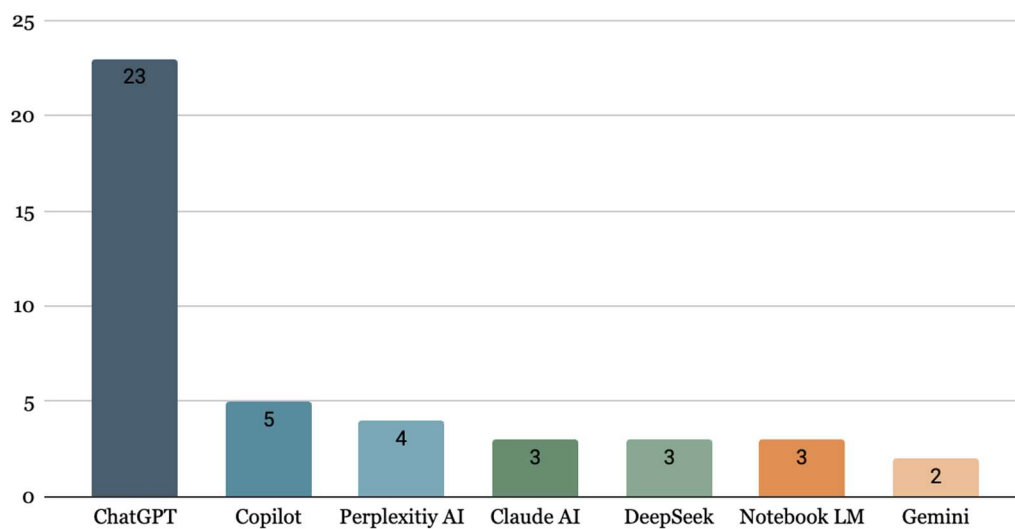


Figure 5: AI Platforms Used by Participants (Multiple Response)

Frequency of Academic Usage

"How often do you use AI for academic purposes?"




Figure 6: Self-reported Frequency of Academic Usage of AI

Frequency of Personal Usage

"How often do you use AI for personal purposes?"




Figure 7: Self-reported Frequency of Personal Usage of AI

Primary Findings

Across the 26 interviews conducted, our findings suggest that awareness of the environmental impact of AI does influence student usage, though this effect varies according to participants' climate attitudes. Students who expressed strong concern about climate change were generally more likely to report that the environmental consequences of AI influenced their behaviour – nine out of twelve in this group indicated that their AI usage had already been shaped by these

concerns. However, after being exposed to an informational prompt outlining the environmental cost of generative AI, only three of these students said they would further reduce their consumption.

Among the eight students who expressed a more moderate level of climate concern (‘Care’), three said the prompt would lead them to reduce their usage, while others, previously unaware of AI’s environmental footprint, noted that the information would likely influence them. It is worth noting, however, that two participants in this group said the statistics presented in the prompt would not change their behaviour. Within the subgroup who expressed the least concern (‘slightly care’), two out of four participants reported they would decrease their usage in response to the prompt, while one remained unaffected. The final participant in this category, who had been unaware of the environmental impact prior to the prompt, said their behaviour would be influenced as a result. This suggests that even among students with relatively low concern for climate issues, targeted awareness efforts can prompt behavioural reflection and, in some cases, change.

Post-Prompt AI Usage Reduction

"Does the prompt reduce your AI usage, if so how?"

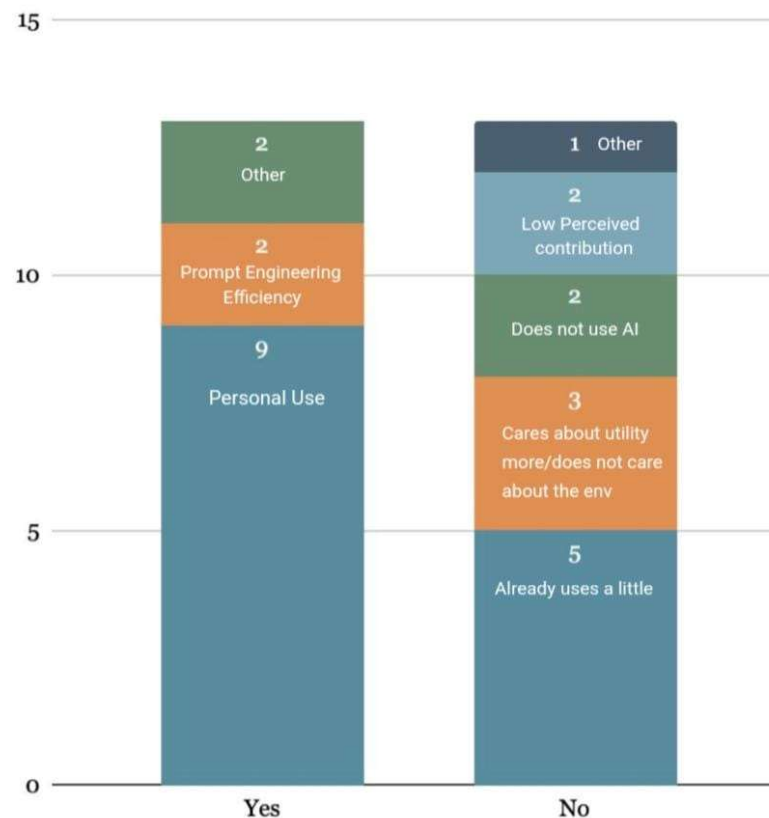


Figure 8: Post-Prompt AI Usage Reduction

Another notable theme was the distinction students made between personal and academic AI use. Many of those who expressed willingness to reduce their usage were referring specifically to personal activities – such as using ChatGPT for travel planning, advice, or casual information-seeking – rather than academic applications. Due to the environmental impact associated with AI usage, some users may choose to minimise their prompts – omitting niceties like “please” and “thank you” – to reduce computational load and energy consumption. This pragmatic approach reflects a growing awareness of the ecological cost of even seemingly small interactions with generative AI systems. One participant mentioned sharing what they learned with their family, including their mother and grandparents, to reduce broader household usage.

Purpose For Academic Use




Figure 9: Purposes for Academic Use of AI (Multiple Response)

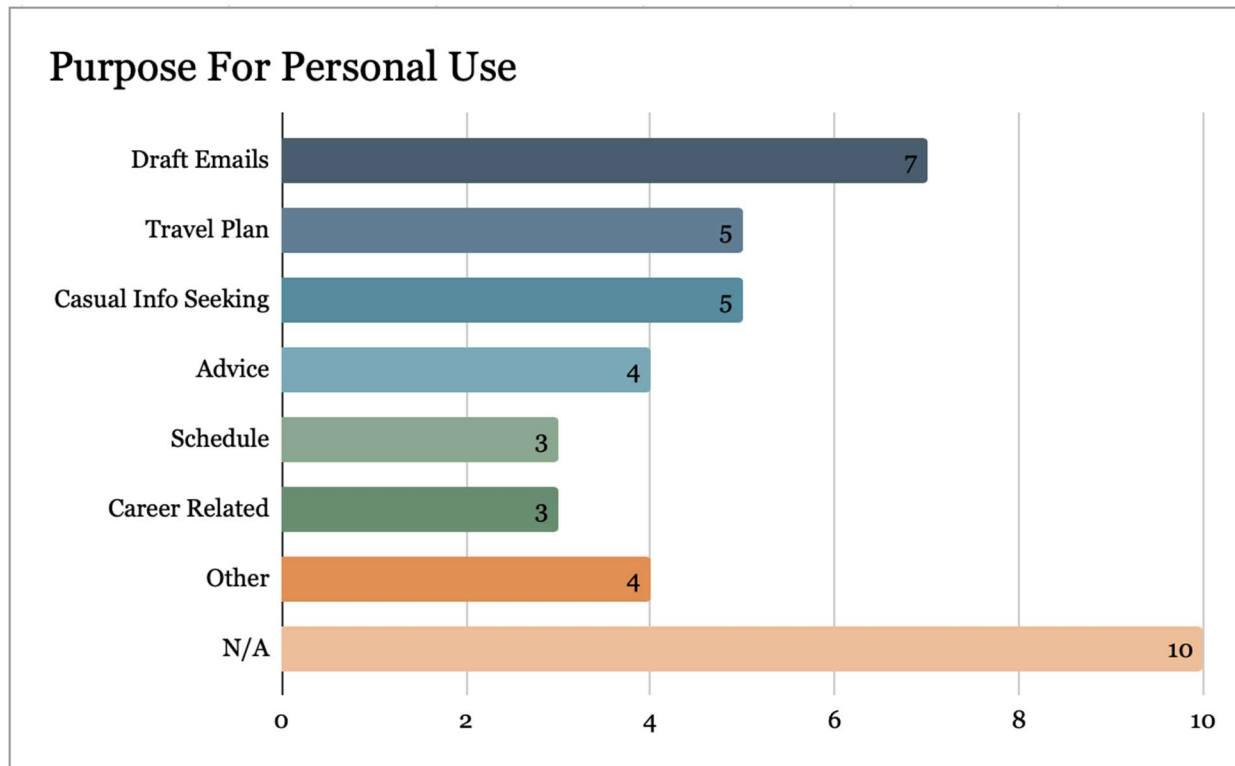


Figure 10: Purposes for Personal Use of AI (Multiple Response)

These findings suggest that while awareness of AI's environmental impact can reduce consumption among students, the degree of reduction is heavily moderated by pre-existing climate attitudes and knowledge. Students who already care deeply about climate change may be less susceptible to further behaviour change, while those previously unaware appear more responsive to new information. Importantly, all four participants who had no prior knowledge of AI's environmental implications stated that the prompt affected them, underscoring the critical role of informational exposure in shaping responsible consumption.

Academic Utility

A central theme that emerged from the interviews was the perceived academic indispensability of generative AI tools. Guided by the UTAUT2 framework, this aligns closely with the construct of *performance expectancy* – the belief that technology will enhance one's ability to perform academically. Many students described AI not as an optional enhancement, but as a necessary adaptation to the evolving academic landscape. One participant noted that “the bar is set higher” due to widespread AI adoption, reflecting a belief that abstaining from AI use could place them at a competitive disadvantage.

Participants further reinforced this notion by contrasting the accessibility of AI tools with traditional academic support. For instance, one student remarked that it was “far easier” to access AI than to visit campus and consult a lecturer, suggesting that convenience is a major driver of adoption. Others described AI as a critical tool for managing “academic overload,” with usage framed as a rational, pragmatic response to institutional pressures.

Furthermore, our study identified distinct patterns in students’ use of generative AI across academic disciplines. In essay-based courses, students primarily used AI to define terms, understand concepts, assist with essay structure, correct grammar, and summarise readings. Similarly, students in quantitative fields used AI for conceptual understanding and terminology but also relied on it to generate code for assignments. A computer science student noted that AI is often encouraged, describing how students “just copy over the code and then kind of use ChatGPT to understand because it's such a dense network of files.”

In contrast, a music student reported not using AI for performance-based coursework, stating, “AI hasn’t yet transformed the music industry as it has in many other fields,” though she uses it for academic writing. Students also emphasised that institutional and departmental policies significantly shape AI usage. For example, an anthropology student explained that while her department prohibits AI for essay writing, it permits its use for understanding theories and literature. A finance student, meanwhile, highlighted AI's usefulness in exam preparation, adding, “they [professors/departments] can’t really mandate how I use ChatGPT to learn stuff.” Many students pointed to inconsistencies between university-wide and departmental policies, with some admitting they do not always adhere to official guidelines. While most did not view AI as essential for academic success, they acknowledged that widespread use among peers creates pressure to adopt it to remain competitive.

Taken together, these findings imply that for many students, academic utility tends to override environmental considerations. Calls for environmentally conscious behaviour are often outweighed by immediate academic demands, suggesting that awareness-based interventions may have limited impact when they conflict with perceived academic necessity.

Perceived Individual Contribution

Another key theme across interviews was the belief that individual behaviour makes little difference in addressing the environmental impact of AI. Sixteen of the 26 students stated explicitly that governments and institutional actors – rather than individuals – should bear the primary responsibility for mitigating climate change. Corporations and international organisations were also frequently named as the main culprits, with participants expressing scepticism about the impact of personal action.

Several students noted that their own contributions would be negligible compared to the structural drivers of climate harm. One student responded to the environmental prompt by saying that even if they stopped using AI, “it wouldn’t matter,” as meaningful change would depend on fossil fuel regulation by corporations. Another participant mentioned that they regularly eat meat, which they perceived as more harmful, suggesting that this minimised the relative importance of reducing AI use. Comments such as “if one person stops, it doesn’t make a difference” were common, indicating a sense of limited individual agency and responsibility. These findings reveal a cognitive barrier to behaviour change: where responsibility is externalised, motivation to act is diminished.

Academic and Informational Integrity

Several participants expressed scepticism toward the academic value and informational reliability of generative AI, with some stating that concerns about accuracy had a greater impact on behaviour than environmental issues. Several students mentioned avoiding AI in academic contexts due to doubts about factual correctness and originality. One described AI as “useless for essay writing,” while others raised fears of plagiarism, misinformation, and academic misconduct. These concerns – perceived as immediate and institutionally relevant – often carried more weight than abstract environmental consequences.

Participants also voiced broader ethical concerns, particularly about AI’s role in creative industries. Some worried that generative AI could devalue human creativity or displace artistic labour. For example, one student acknowledged AI’s utility but hesitated to fully embrace it due to its perceived effect on academic culture. These insights suggest that while environmental concerns may be cognitively recognised, it is often the more proximate risks to academic integrity and professional ethics that drive student behaviour.

Visualisation

Finally, participants repeatedly identified the abstract nature of AI’s environmental impact as a challenge to meaningful engagement. Many reported a lack of emotional connection or urgency, not due to apathy, but because the consequences felt remote or intangible. One participant remarked, “you can’t really visualise these damages it’s doing to the environment,” which led them to feel disconnected from the actual effects of their AI usage.

Several students suggested that more vivid, personal, or emotionally resonant representations of environmental harm could strengthen the case for reduced consumption. One participant said they would be more likely to stop using AI if they could “visualise the actual impact,” while another recommended using “personal stories” to communicate environmental consequences. These

insights point to a significant perceptual gap: when environmental harm is framed in abstract or statistical terms, it is more easily dismissed or deprioritised. As one student summarised, “the more we find out the impacts, the more I’d be disincentivised, but it will have to be a big thing.” This underscores a key mechanism behind our findings – without tangible reference points, environmental awareness alone often fails to motivate behavioural change.

Taken together, our findings indicate that while raising awareness of AI’s environmental impact can influence student behaviour, the overall effectiveness of such interventions is limited by several intersecting factors. Students’ pre-existing climate attitudes, perceptions of individual responsibility, and the prioritisation of academic utility all shape how environmental information is received and acted upon. Although some participants demonstrated a willingness to reduce AI usage – particularly in personal, non-academic contexts – the perceived indispensability of AI for academic success often outweighed environmental considerations. Furthermore, where students viewed environmental harm as abstract or psychologically distant, behavioural change was less likely. In contrast, immediate concerns such as academic integrity and information reliability exerted a stronger influence on usage patterns. These insights suggest that behavioural shifts regarding AI consumption are more likely when interventions are personalised, visually compelling, and aligned with students’ lived realities. Future efforts to encourage more sustainable AI use must therefore go beyond awareness-raising alone, engaging more deeply with how students weigh trade-offs between ethics, utility, and agency in their everyday lives.

Limitations and Ethical Considerations

This study faces some limitations worth addressing. First, the findings are based on semi-structured interviews in which participants self-reported their use of generative AI and predicted behavioural changes in response to environmental information. Self-reported data are inherently susceptible to social desirability bias, whereby individuals may respond in ways they believe are socially acceptable rather than entirely truthful. Second, sample size and composition limit the generalisability of the findings. The study included 26 participants, of whom 17 self-identified as female, 7 as male, and 2 as non-binary. The sample was also disproportionately composed of students from the social sciences – particularly economics and international relations – due to the constraints of convenience sampling. This imbalance, combined with the relatively small sample size, restricted the ability to draw meaningful comparisons across demographic or disciplinary groups.

In addition, several participants appeared to struggle with fully internalising statistical information about AI’s environmental impact. Some noted that visual or narrative representations – such as images, videos, or personal stories – might have elicited stronger emotional engagement and behavioural reflection. This suggests a potential limitation in the communicative approach

used and highlights a promising direction for future research. Studies exploring the effects of visual or multimodal environmental messaging on AI-related behaviour could yield further insights. Lastly, while the interviews collected information about university AI policies, the study did not establish any direct relationship between these institutional factors and students' usage patterns. Further investigation into the influence of formal policy environments on student behaviour would offer a valuable extension to the present work.

Conducting interviews on topics related to individual behaviour and values, such as attitudes towards AI and environmental responsibility, involved a number of ethical considerations. Participants may be asked to disclose personal opinions, habits, and experiences, which can provoke discomfort or raise concerns around privacy. To address these issues, the research followed established ethical guidelines for qualitative inquiry involving human subjects. All participants were fully informed about the purpose and scope of the study and provided written consent prior to taking part. They were reminded of their right to withdraw at any stage without penalty, and were assured that their responses would be anonymised and kept strictly confidential. Interview questions were designed to be respectful and non-intrusive, and participants were encouraged to skip any questions they did not feel comfortable answering. By implementing these safeguards, the study sought to uphold participants' autonomy, ensure transparency, and foster a safe and respectful interview environment.

Conclusion

This study explored the ways in which university students' perception of the environmental impacts of generative AI relates to their attitudes towards their usage of AI. We found that the degree to which students would adjust their usage of AI is heavily moderated by their pre-existing environmental knowledge and attitudes. Participants who were exposed to new information about environmental impact during our study became more responsive and susceptible to behaviour change. However, while raising awareness of AI's environmental impact can be shown to positively influence student behaviour, informational exposure is limited by the prioritisation of academic utility and low perceived individual contribution to the environment. Furthermore, an individual's psychological distance from environmental consequences may limit the effectiveness of informational campaigns. Student concerns around academic and informational integrity further limit such efforts. These findings suggest that informational campaigns, while valuable, may be insufficient in isolation. Future interventions must frame environmental consequences as more tangible, immediate, and personally relevant, rather than relying solely on abstract facts or metrics. If AI is to be integrated responsibly into higher education, then sustainability must be treated as a shared ethical obligation, not as an afterthought.

Bibliography

- Chan, C. K. Y., & Colloton, T. (2024). *Generative AI in Higher Education : the ChatGPT effect* (1st ed.). Routledge.
- Gansser, O.A. and Reich, C.S. (2021). A New Acceptance Model for Artificial Intelligence with Extensions to UTAUT2: An Empirical Study in Three Segments of Application. *Technology in Society*, 65(1), p.101535.
doi:<https://doi.org/10.1016/j.techsoc.2021.101535>.
- HEPI (2025). *Student Generative AI Survey 2025*. <https://www.hepi.ac.uk/2025/02/26/student-generative-ai-survey-2025/>. <https://doi.org/10.2307/41410412>
- Hernandez-Arriaza, M., Roque, I. M. S., Blanc, G. A., & Shorey, S. (2023). Gen-Z university students' attitudes toward current global trends: environment, globalisation, immigration, feminism and capitalism. *Globalisation, Societies and Education*, 1–16.
<https://doi.org/10.1080/14767724.2023.2265840>
- International Energy Agency (2024). *Electricity 2024 - Analysis and Forecast to 2026*. [online] IEA. Available at: <https://iea.blob.core.windows.net/assets/6b2fd954-2017-408e-bf08-952fdd62118a/Electricity2024-Analysisandforecastto2026.pdf>.
- Janmaimool, P., & Khajohnmanee, S. (2019). Roles of Environmental System Knowledge in Promoting University Students' Environmental Attitudes and Pro-Environmental Behaviors. *Sustainability*, 11(16), 4270. <https://doi.org/10.3390/su11164270>
- Kim, M.-S., Kim, J., & Thapa, B. (2018). Influence of Environmental Knowledge on Affect, Nature Affiliation and Pro-Environmental Behaviors among Tourists. *Sustainability*, 10(9), 3109-. <https://doi.org/10.3390/su10093109>
- Klockner, C. A. (2013). A comprehensive model of the psychology of environmental behaviour—A meta-analysis. *Global Environmental Change*, 23(5), 1028–1038.
<https://doi.org/10.1016/j.gloenvcha.2013.05.014>
- Kollmuss, A., & Agyeman, J. (2002). Mind the Gap: Why do people act environmentally and what are the barriers to pro-environmental behavior? *Environmental Education Research*, 8(3), 239–260. <https://doi.org/10.1080/13504620220145401>
- Li, P., Yang, J., Islam, M. A., & Ren, S. (2023). Making AI Less “Thirsty”: Uncovering and Addressing the Secret Water Footprint of AI Models. *arXiv.Org*.

- Li, Y., Yang, D., & Liu, S. (2024). The impact of environmental education at Chinese Universities on college students' environmental attitudes. *PloS one*, 19(2), e0299231. <https://doi.org/10.1371/journal.pone.0299231>
- Patterson, D., Gonzalez, J., Holzle, U., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D. R., Texier, M., & Dean, J. (2022). The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink. *Computer* (Long Beach, Calif.), 55(7), 18–28. <https://doi.org/10.1109/MC.2022.3148714>
- Steg, L., Bolderdijk, J. W., Keizer, K., & Perlaviciute, G. (2014). An Integrated Framework for Encouraging Pro-environmental Behaviour: The role of values, situational factors and goals. *Journal of Environmental Psychology*, 38, 104–115. <https://doi.org/10.1016/j.jenvp.2014.01.002>
- Tseng, T. H., Lin, S., Wang, Y. S., & Liu, H. X. (2019). Investigating teachers' adoption of MOOCs: the perspective of UTAUT2. *Interactive Learning Environments*, 30(4), 635–650. <https://doi-org.lse.idm.oclc.org/10.1080/10494820.2019.1674888>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User Acceptance of Information Technology: Toward a Unified View. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>
- Venkatesh, V., Thong, J. Y. L., & Xu, X. (2012). Consumer Acceptance and Use of Information Technology: Extending the Unified Theory of Acceptance and Use of Technology. *MIS Quarterly*, 36(1), 157–178.
- Vicente-Molina, M. A., Fernández-Sáinz, A., & Izagirre-Olaizola, J. (2013). Environmental knowledge and other variables affecting pro-environmental behaviour: comparison of university students from emerging and advanced countries. *Journal of Cleaner Production*, 61, 130–138. <https://doi.org/10.1016/j.jclepro.2013.05.015>
- Zacharis, G. and Nikolopoulou, K. (2022). Factors predicting University students' behavioral intention to use eLearning platforms in the post-pandemic normal: an UTAUT2 approach with 'Learning Value'. *Education and Information Technologies*, 27(1). [doi:https://doi.org/10.1007/s10639-022-11116-2](https://doi.org/10.1007/s10639-022-11116-2).
- Zhu, W., Huang, L., Zhou, X., Li, X., Shi, G., Ying, J., & Wang, C. (2024). Could AI Ethical Anxiety, Perceived Ethical Risks and Ethical Awareness About AI Influence University Students' Use of Generative AI Products? An Ethical Perspective. *International Journal of Human–Computer Interaction*, 41(1), 742–764. <https://doi-org.lse.idm.oclc.org/10.1080/10447318.2024.2323277>