Human Factors in Design, Engineering, and Computing, Vol. 238, 2026, 383–393
https://doi.org/10.54941/ahfe1008271

The Invisible Footprint: Awareness, Attitudes, and Designable Solutions for Sustainable Human-AI Interaction

Jasmine McKenzie[1,2], Alaina E. Smith[1,2], Mark Holden[1], Wesley Junkins[1], and Juan E. Gilbert[1]

[1]University of Florida, Gainesville, FL 32607, USA
[2]Virginia State University, Petersburg, VA 23806, USA

ABSTRACT

What people know about an AI system shapes their relationship with it just as much as what that system can do. As generative AI (GenAI) becomes embedded in student life, an open question in human-AI interaction is how new information, including a system's environmental footprint, reshapes that relationship. This study tested whether a short micro-learning intervention could shift students' attitudes, subjective norms, perceived behavioral control, and intentions toward more sustainable use of GenAI. Using a two-arm design, 125 undergraduate and graduate students across four universities were randomly assigned to view a five-minute video on GenAI's water, energy, and hardware costs either before or after completing surveys measuring these constructs. Following the intervention, students showed large gains in perceived knowledge and behavioral intention, and a moderate gain in perceived behavioral control; attitudes and subjective norms did not change significantly. A regression model explained 72% of the variance in intention, with attitudes being the strongest predictor and knowledge showing no independent effect. These findings suggest that informing students about GenAI's environmental costs can measurably shift how they relate to the technology. That attitude change drives intention to act, not knowledge alone, which is a finding with direct implications for human-AI interaction design. The results extend responsible-AI discourse to include environmental stewardship, pointing to concrete strategies such as usage feedback, eco-mode defaults, and AI-literacy curricula to embed sustainability in everyday human-AI interaction.

Keywords: Generative AI, Environmental sustainability, Human-AI interaction, Theory of planned behavior

INTRODUCTION

What people know about an artificial intelligence (AI) system shapes their relationship with it just as much as what that system can do. As generative AI (GenAI) becomes embedded in students' academic routines, an open question is how new information, including information users don't expect, like a system's environmental footprint, reshapes that relationship. GenAI offers a clear case. With ChatGPT users aged 18–34 accounting for an estimated 53–60% of its user base, these tools have shifted from a passive tool to a seemingly ubiquitous presence in students' lives, used as everyday aids for

Received September 24, 2026; Revised October 23, 2026; Accepted October 31, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

tasks like writing code, finding information, solving problems, and drafting essays (Attewell, 2024; Digital Education Council, 2024; Duarte, 2026; Omnibound, 2025; Singh, 2026). Much of the resulting research has focused on GenAI's visible tradeoffs, such as productivity gains, academic integrity concerns, and effects on student development (Li and Jan, 2023; Lund et al., 2025; Yavich, 2025). Far less attention has been paid to a cost that is structurally invisible to the people generating it: running and training these large-scale AI models consumes enormous amounts of electricity, requires significant water for cooling data centers, and depends on the constant production and disposal of hardware (Strubell, Ganesh, and McCallum, 2019; Zewe, 2025). Unlike issues such as plagiarism or algorithmic bias, this cost leaves no trace on the interface itself.

The scale of that cost is well documented. Training a single large model can require weeks of computation across thousands of graphical processing units (GPUs) (Kandemir, 2025), which is enough electricity to power more than 100 homes for a year (Zewe, 2025). Each prompt after training adds further inference costs. For example, a single generative AI query consumes roughly five to ten times more electricity than a standard web search (de Vries, 2023). Powering these systems still heavily relies on fossil fuels. One estimate put the carbon emissions from training a single large model at over 626,000 pounds, which is comparable to the lifetime emissions of several cars (Strubell, Ganesh and McCallum, 2019). An increasingly visible issue is water consumption. Data centers tend to rely on water-based cooling systems to prevent the hardware from overheating (Kandemir, 2025; Skibell, 2026). As a result, training and operating AI models can require millions of gallons of mostly fresh water, which is a finite, non-recyclable resource in this context (Gioino, 2026). Generally, salt water has not been used because it can damage hardware, making freshwater dependence unavoidable with current infrastructure (Fox, 2026). Unfortunately, high water consumption for AI systems is already creating competition with local communities in drought-prone regions across the US (Gioino, 2026; Skibell, 2026). Finally, there are also the problems of electronic waste and resource extraction. Building powerful AI chips requires rare minerals like gallium and germanium (Kandemir, 2025; Mineral Resources Program, 2025). Mining these materials often harms ecosystems, pollutes water, and creates unsafe working conditions for miners (Zapp et al., 2022). Since hardware tends to become obsolete quickly, e-waste is difficult to recycle and harmful if not disposed of properly (Kandemir, 2025).

Responsible AI research typically centers around fairness, bias, transparency, explainability, privacy, and safety; however, environmental sustainability sits outside of this center. When sustainability is discussed, it is usually framed as an engineering problem, such as designing more efficient models or greener data centers (Kandemir, 2025), overlooking the people generating the demand in the first place. In contrast, human-AI interaction research more often looks at how people relate to an AI system, including how much they trust it and how that trust responds to new information (Lee and See, 2004). That relationship is dynamic, so it can move in either direction when new evidence emerges, which means that trust can sometimes

be eroded and sometimes grow stronger (Dietvorst, Simmons, and Massey, 2015; Logg, Minson, and Moore, 2019). The environmental cost of AI is new information, and end-user awareness of it remains almost entirely understudied despite the fact that billions of individual interactions are what aggregate into these systems' total footprint. This work treats responsible AI as encompassing environmental stewardship at the user level and not only the engineering level.

To test whether awareness can shift behavior intentions, this work is grounded in the Theory of Planned Behavior (TPB), which models behavioral intention as a function of three antecedents. Attitudes reflect how positively or negatively a person evaluates a behavior (Bosnjak, Ajzen, and Schmidt, 2020). Subjective norms are the perceived social pressure or approval from important people in someone's world, such as peers or instructors (Bosnjak, Ajzen, and Schmidt, 2020). Perceived behavioral control is one's confidence in their ability to perform a behavior (Bosnjak, Ajzen, and Schmidt, 2020). A fourth construct, perceived knowledge, is a person's self-assessed sense of how informed they are about a given issue. This isn't part of core TPB, but it's well established as shaping both attitudes and perceived control. That makes it the entry point into our model: intervention → knowledge → attitudes, norms, and control → intention. TPB is the dominant framework in technology adoption and sustainability-behavior research, but it has not been explored much in relation to GenAI use, a gap this study addresses directly.

This study tested whether a short micro-learning intervention could shift students' attitudes, perceived behavioral control, and intentions toward more sustainable use of GenAI. In doing so, it extends the idea of responsible AI beyond fairness, ethics, and accessibility to include environmental sustainability, while also contributing to a growing body of research on human-AI interaction that examines how new information about a system reshapes people's relationship to it.

METHODS

This study was designed using a mixed design with both between-subjects and within-subject groups across two conditions. Undergraduate and graduate students aged 18-34 were recruited across four universities. Students were randomly assigned to a condition with roughly 60 participants per group.

Group 1 used a within-subjects design with pre- and post-surveys. Students completed a pre-survey that measured their current use of generative AI, attitudes, social norms, behavioral intentions, knowledge, and perceived self-control over GenAI. After completing the pre-survey, they viewed a micro-learning module on the environmental costs of GenAI, including water, energy, and hardware. Participants then completed a post-survey to assess any shifts in the measures captured in the pre-survey. The post-survey also included a Planned Behavior Index, a Cost-Benefit Tradeoff scale, and two knowledge-check items to assess the actions students intended to adopt, whether they perceived eco-friendly AI use as beneficial, and how much they retained from the video.

In Group 2, participants were not exposed to the pre-survey, so they first watched the environmental costs video and then completed the same post-survey as those in Group 1. This design tested whether the pre-survey questions in Group 1 primed participants to give more socially desirable responses, a common concern in attitudinal research. The full study lasted ~15 minutes for both conditions, and all participants received course credit. This study was deemed exempt by the University of Florida's exempt auto-determination tool (Protocol #: ET00048038).

A total of 125 students participated across four universities (Group 1: n = 63; Group 2: n = 62). One Group 1 participant did not complete all pre-survey items and was excluded from paired-sample analyses. Most were undergraduates (79.2%), with the remainder enrolled in graduate programs (20.8%), and 92% were between 18 and 34 years old. Of the 112 participants who reported race and ethnicity, 40.2% identified as White, 37.5% as Asian, 16.1% as Black or African American, and 5.4% as another race; 13 participants declined to report their race and ethnicity.

Micro-Learning Intervention

Each participant reviewed a micro-learning intervention on the environmental effects of generative AI. The intervention was a 5-minute video covering energy demands, water usage for cooling, e-waste, and resource extraction. The video also included individual-level mitigation strategies for generative AI use. One strategy was the intentional use of these systems. Instead of using GenAI for all inquiries, students could turn to a traditional search engine. A second was prioritizing text generation over image or video generation because it consumes more energy. A third was prompt optimization, which involves providing comprehensive, detailed instructions in the initial prompt to reduce the number of iterations. A fourth was batching, which involves combining questions or prompts to reduce the need to prompt the model repeatedly. Finally, there was avoiding pleasantries, because even though they are small, they can add up quickly in terms of the energy required to process small additions like "please" and "thank you" from billions of people. These changes were intended to be approachable for the average user. The goal was to get students to consider how large-scale use of these strategies could help alleviate some of the stress GenAI places on the environment. The video was posted privately on YouTube and delivered via embedded link in Qualtrics. Closed captions were also available.

Measures

Table 1 summarizes the measures used across the pre- and post-surveys. Attitudes, subjective norms, perceived behavioral control, and perceived knowledge were each measured before and after the intervention on 7-point Likert scales. The pre-survey used resource-neutral framing, while the post-survey used explicit environmental framing, minimizing priming while still capturing the same underlying constructs. The post-survey also included a six-item Planned Behavior Index (PBI), with each item summed, a three-item Cost-Benefit Tradeoff scale, and two multiple-choice knowledge-check items.

Baseline AI use frequency (school and non-school prompt counts) and tool preferences were also captured in the demographics.

Table 1: Pre- and post-survey measures by construct.

Construct	Pre – Arm 1 Only	Post – Both Arms	Format
Attitude	Neutral/resource framing	Explicit environmental framing	7-pt Likert (4 items)
Subjective norms	Neutral	Environmental	7-pt Likert (3 items)
Perceived behavioral control	Neutral	Environmental	7-pt Likert (3 items)
Behavioral intention	Neutral	Environmental	7-pt Likert (4 items)
Perceived knowledge	Neutral	Environmental	7-pt Likert (2 items)
Planned behaviors index	—	Yes/No (6 items)	Sum 0–6
Cost–benefit tradeoff	—	7-pt Likert (3 items, 1 reverse-scored)	Mean
Knowledge checks	—	2 multiple choice items	0–2 correct

Analysis

Analyses were conducted in IBM SPSS. The descriptive statistics summarized the frequency of AI use, prompt counts, tool preferences, and demographic composition. Within Group 1, paired-samples t-tests compared pre- and post-intervention scores on all five constructs, with Cohen's dz reported for effect size and 95% confidence intervals reported for the mean difference. Independent-samples t-tests compared post-intervention scores between Group 1 and Group 2 to assess priming effects. A multiple regression examined post-intervention attitudes, subjective norms, perceived behavioral control, perceived knowledge, and cost-benefit tradeoff as predictors of behavioral intention. Internal consistency for each multi-item scale was assessed using Cronbach's alpha.

RESULTS

Internal reliability for each scale was assessed using Cronbach's alpha. Post-intervention scales demonstrated good to excellent reliability: attitudes ($\alpha = .871$), subjective norms ($\alpha = .767$), perceived behavioral control ($\alpha = .901$), behavioral intention ($\alpha = .930$), and perceived knowledge ($\alpha = .874$). Pre-intervention scales showed acceptable to good reliability for most constructs: attitudes ($\alpha = .720$), perceived behavioral control ($\alpha = .764$), behavioral intention ($\alpha = .804$), and perceived knowledge ($\alpha = .750$). Pre-intervention subjective norms showed lower consistency ($\alpha = .496$). The cost-benefit tradeoff scale, collected only post-intervention, also demonstrated low reliability ($\alpha = .464$). For subjective norms, this likely reflects the intentionally neutral pre-survey framing, since post-intervention norms showed markedly improved reliability ($\alpha = .767$) once environmental framing was introduced. Given its low reliability, the three cost-benefit tradeoff items are reported descriptively rather than as a composite.

Descriptive Statistics and AI Use Patterns

Pre-intervention means for Group 1 ranged from 4.75 (perceived knowledge) to 5.75 (attitudes) on the 1–7 scale, indicating that students entered the study with moderately positive orientations toward resource-conscious AI use. Post-intervention means across all 125 participants were uniformly higher, ranging from 5.24 (subjective norms) to 5.98 (perceived knowledge). This pattern is consistent with the intervention producing improvement across all measured constructs, even when individual changes did not reach statistical significance. Among Group 1 participants, GenAI use was already frequent at baseline: 57.2% reported using it a few times per week, and 27.0% reported daily use. School- and non-school-related prompt volume in the preceding week varied widely, with a substantial share of students reporting both light use (1–5 prompts) and heavy use (21 or more) in both contexts.

Knowledge Check

To confirm that participants retained information from the micro-learning module, two multiple-choice knowledge-check items were included in the post-survey. Most participants answered both questions correctly: 117 of 125 (93.6%) identified that increased paper printing was not a cited environmental impact of large-scale AI systems, and 113 of 125 (90.4%) identified batching-related prompts as a strategy that reduces environmental impact. These results indicate that the micro-learning module successfully conveyed its core content.

Pre-Post Changes in Group 1

Paired-samples t-tests assessed whether the micro-learning module produced a significant change across the five constructs (see Table 2). Perceived knowledge showed the largest shift, increasing by 1.36 points on average, a large effect. Behavioral intention also increased significantly by 0.87 points on average, and perceived behavioral control improved by 0.80 points on average, indicating a medium effect. Attitudes and subjective norms did not change significantly.

Table 2: Pre-to post-intervention changes in group 1 (n = 63).

Construct	Pre M (SD)	Post M (SD)	t(61)	p	dz	95% CI
Perceived knowledge	4.75 (1.30)	6.11 (0.92)	−7.464	<.001**	0.95	[0.998, 1.728]
Behavioral intention	5.02 (1.28)	5.88 (1.16)	−6.874	<.001**	0.87	[0.618, 1.124]
Perceived behavioral Control	5.03 (1.21)	5.81 (1.24)	−3.638	<.001**	0.46	[0.358, 1.233]
Attitudes	5.75 (0.89)	5.83 (1.25)	−0.519	.606	—	[−0.082, 0.534]
Subjective norms	5.23 (0.83)	5.44 (1.14)	−1.466	.148	—	[−0.219, 0.372]

Note. n = 62 for paired analyses due to one participant's incomplete pre-survey data.

Survey Priming Check

To assess whether pre-survey exposure primed Group 1 to respond differently on post-measures, independent-samples t-tests were conducted to compare post-intervention scores between Group 1 (n = 63) and Group 2 (n = 62). Significant differences emerged for perceived behavioral control (Group 1 M = 5.81, Group 2 M = 5.05), t(123) = 3.202, p = .002, d = 0.573, and behavioral intention (Group 1 M = 5.88, Group 2 M = 5.37), t(123) = 2.139, p = .034, d = 0.383, suggesting the pre-survey amplified post-intervention responses on these two constructs beyond what the video alone produced. No significant differences were found for perceived knowledge (p = .093), attitudes (p = .539), or subjective norms (p = .056), indicating that knowledge gains were driven by the intervention itself rather than by priming.

Predictors of Intention

A multiple regression examined attitudes, subjective norms, perceived behavioral control, perceived knowledge, and cost-benefit tradeoff as predictors of post-intervention behavioral intention (see Table 3). The overall model was significant and explained 72% of the variance in intention. Attitudes emerged as the strongest predictor: students who more strongly valued eco-conscious AI use reported substantially higher intentions to change their behavior. The cost-benefit tradeoff was the only other significant predictor, suggesting that students who saw eco-friendly AI use as academically compatible were more likely to intend to act. Subjective norms, perceived behavioral control, and perceived knowledge did not independently predict intention once the other predictors were accounted for.

Table 3: Predictors of post-intervention behavioral intention (multiple regression).

Predictor	B	SE	β	t	p
(Constant)	–1.178	.473	—	–2.493	.014
Attitudes	.711	.085	.599	8.322	<.001**
Cost-benefit tradeoff	.222	.074	.174	2.984	.003**
Subjective norms	.112	.067	.097	1.656	.100
Perceived knowledge	.099	.077	.071	1.277	.204
Perceived behavioral control	.079	.063	.079	1.245	.215

DISCUSSION

The micro-learning module produced change, but not equally. The participants in Group 1 showed large gains in perceived knowledge (d = 0.95) and behavioral intention (d = 0.87), and a moderate gain in perceived behavioral control (d = 0.46). Therefore, the brief, five-minute exposure to GenAI's water, energy, and hardware costs was enough to shift what students knew, how capable they felt of changing their use, and how likely they said they were to do so. Attitudes and subjective norms, however, did not change significantly. This is interesting because attitudes turned out to be

the strongest predictor of intention, so the intervention shifted what students knew and what they felt able to do without shifting how they felt about GenAI itself. Subjective norms also do not change significantly, likely because the environmental costs of AI aren't something that students hear from their peers or professors yet. Without a shared social reference point, a single exposure has little room to move perceived norms. Therefore, one short video can inform students and make them feel more in control. However, it can't change how they feel about GenAI or what they think others expect of them on its own.

Why Knowledge Didn't Predict Intention

Despite the large knowledge gain, this did not independently predict behavioral intention. Attitudes were by far the strongest predictor ($\beta = .599$), with cost-benefit tradeoff being the only other significant contributor ($\beta = .174$). With perceived control and social norms, the model explained 72% of the variance in intention ($R^2 = .720$). This is consistent with what environmental behavior research shows: telling people more facts does not automatically change what they do (Kollmuss and Agyeman, 2002). For GenAI, this suggests that students may already have sufficient baseline awareness, so additional facts alone won't change their intention. Instead, what moves them is how they feel about that information, not just having it. That means future interventions may need to target attitudes directly rather than only targeting knowledge and hoping that attitudes follow.

The Academic Cost Barrier

Even when attitudes and intentions shift, translating them into reduced GenAI use is complicated by the academic context in which it occurs. GenAI is embedded in the time-pressured, outcome-driven structure of coursework, where deadlines and grading pressure reward efficiency over deliberation. The fact that the cost-benefit tradeoff significantly predicted intention shows that students are already weighing environmental costs against these academic pressures. However, the math rarely favors cutting back when falling behind feels far more immediate than an environmental cost that feels far away. This mirrors a well-documented gap between environmental attitudes and behavior (Kollmuss and Agyeman, 2002). Recent work on GenAI specifically shows the same pattern: even when students were shown real-time carbon feedback, their overall platform usage didn't drop, in part because they prioritized reliability, cost, and speed over environmental impact (Andersen et al., 2026).

Human-AI Interaction Design and Educational Implications

These findings carry direct implications for how GenAI systems are designed. If attitude change and information alone drive intention, then static, one-time disclosures are unlikely to be sufficient on their own. The eco-feedback literature in human-computer interaction offers a more promising direction:

real-time, ambient feedback about resource consumption has been shown to shift attitudes and behavior more durably than one-time information delivery, by keeping costs visible at the moment of use (Froehlich, Findlater, and Landay, 2010). Recent consumer research suggests disclosure of AI's environmental footprint, paired with credible and feasible action options, can turn even intensive AI users into a receptive audience for sustainable consumption (Apaolaza et al., 2026), lending some indirect support to feedback-based design directions like these. Translated to GenAI interfaces, this could take the form of per-query energy or water estimates, eco-mode defaults that route simple queries to smaller models, or prompt-batching nudges that consolidate redundant queries. Since subjective norms did not shift in this study, another design direction worth exploring is surfacing social or institutional benchmarks, such as showing users how their usage compares to a peer or department average. This kind of comparison could reach the social influence that a purely informational intervention could not reach on its own.

Beyond interface design, these results speak to formal AI-literacy education. Universities have moved quickly to establish GenAI policies for coursework, but most address academic integrity, not environmental impact. Embedding environmental costs into existing AI-literacy curricula through orientation modules, library workshops, or course-level AI-use policies offers a low-cost way to reach students who respond measurably to even brief exposure, as the results show. However, since knowledge gain alone didn't predict intention, curricula that only convey facts are unlikely to be sufficient. Instead, universities may need to pair information with activities designed to prompt attitudinal reflection, and not just awareness, to move beyond informing students toward actually shifting how they relate to their own GenAI use.

LIMITATIONS, FUTURE WORK, AND CONCLUSION

One limitation of this study was that all measures were self-reported, and intention was measured immediately after the intervention rather than being linked to observed behavior, leaving open the question of whether these shifts persist or translate into reductions in actual use. The between-arm comparison also suggested some priming from the pre-survey on Group 1's post-intervention control and intention scores. Group 2, who were unaffected by the priming, offers a more conservative estimate of the intervention's standalone effect. Finally, the intervention consisted of a single passive video and did not test the interactive or curricular strategies proposed above. Future work should track whether these shifts persist and translate into measured behavior change, test interactive eco-feedback interfaces directly, and extend this design into more institutions. Together, these findings suggest that a brief intervention can meaningfully shift knowledge, control, and intention. However, moving from awareness to sustained behavior change will likely require strategies that engage attitudes directly, not information delivery alone.

REFERENCES

Andersen, L.B., Herklotz, M., Liu, A., Goeke, M., Juelich, M., Kern, C. and Kreuter, F. (2026). From awareness to action? The impact of CO2 emission feedback on student LLM usage. Proceedings of the Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems, pp. 1–11. doi: 10.1145/3772363.3798840.

Apaolaza, V., Paredes, M.R., Eletxigerra, A. and Hartmann, P. (2026). Is AI use intensity associated with consumers' pro-climate intentions? The roles of climate anxiety, climate commitment, and awareness of AI's environmental impact. International Journal of Consumer Studies, 50(5). doi: 10.1111/ijcs.70248.

Attewell, S. (2024). How students are currently using generative AI. [online] Jisc: National Centre for Artificial Intelligence. https://nationalcentreforai.jiscinvolve.org/wp/2024/03/26/how-students-are-currently-using-generative-ai/

Bosnjak, M., Ajzen, I. and Schmidt, P. (2020). The theory of planned behavior: Selected recent advances and applications. Europe's Journal of Psychology, 16(3), pp. 352–356. doi: https://doi.org/10.5964/ejop.v16i3.3107.

de Vries, A. (2023). The growing energy footprint of artificial intelligence. Joule, 7(10). doi: 10.1016/j.joule.2023.09.004.

Dietvorst, B.J., Simmons, J.P. and Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. Journal of Experimental Psychology: General, 144(1), pp. 114–126. doi: 10.1037/xge0000033.

Digital Education Council (2024). How Students Use AI: The Evolving Relationship Between AI and Higher Education. Digital Education Council. https://www.digitaleducationcouncil.com/dec-insights/how-students-use-ai-the-evolving-relationship-between-ai-and-higher-education

Duarte, F. (2026). Number of ChatGPT users (January 2026). Exploding Topics. https://explodingtopics.com/blog/chatgpt-users

Fox, A. (2026). Why do data centers need fresh water? Friends of the Earth. https://foe.org/blog/why-do-data-centers-need-fresh-water/

Froehlich, J., Findlater, L. and Landay, J. (2010). The design of eco-feedback technology. Proceedings of the 28th international conference on Human factors in computing systems - CHI '10. doi: 10.1145/1753326.1753629.

Gioino, C. (2026). America's data centers are thirsty. Rural towns are paying the price—From tanked water pressure to stolen desert groundwater. Fortune. https://fortune.com/2026/05/13/data-center-georgia-arizona-water-wars/

Kandemir, M. (2025). Why AI uses so much energy—And what we can do about it. Penn State Institute of Energy and the Environment. https://iee.psu.edu/news/blog/why-ai-uses-so-much-energy-and-what-we-can-do-about-it

Kollmuss, A. and 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), pp. 239–260. doi: 10.1080/13504620220145401.

Lee, J.D. and See, K.A. (2004). Trust in automation: Designing for appropriate reliance. Human Factors: The Journal of the Human Factors and Ergonomics Society, 46(1), pp. 50–80. doi: 10.1518/hfes.46.1.50_30392.

Li, E.Y. and Jan, A. (2023). Impact of artificial intelligence (AI) in enhancing productivity and reducing stress among students. In: ICEB 2023 Proceedings (Chiayi, Taiwan). https://aisel.aisnet.org/iceb2023/33/

Logg, J.M., Minson, J.A. and Moore, D.A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. Organizational Behavior and Human Decision Processes, 151(1), pp. 90–103. doi: 10.1016/j.obhdp.2018.12.005

Lund, B.D., Lee, T.H., Mannuru, N.R. and Arutla, N. (2025). AI and academic integrity: Exploring student perceptions and implications for higher education. Journal of Academic Ethics, 23, pp. 1545–1565. doi: 10.1007/s10805-025-09613-3

Mineral Resources Program (2025). Key minerals in data centers infographic. [online] USGS. https://www.usgs.gov/media/images/key-minerals-data-centers-infographic

Omnibound (2025). ChatGPT user statistics 2025-26: 56+ Data points on users, growth, countries, and revenue. Omnibound.ai. https://www.omnibound.ai/blog/chatgpt-user-statistics

Singh, S. (2026). ChatGPT statistics for 2026: Comprehensive facts and data. Demand Sage. https://www.demandsage.com/chatgpt-statistics/

Skibell, A. (2026). A data center drained 30m gallons of water unnoticed — Until residents complained about low water pressure. POLITICO. https://www.politico.com/news/2026/05/08/georgia-data-centers-water-00909988

Strubell, E., Ganesh, A. and McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics, pp. 3645–3650. doi: 10.18653/v1/P19-1355.

Yavich, R. (2025). Will the use of AI undermine students independent thinking? Education Sciences, 15(6), p. 669. doi: 10.3390/educsci15060669.

Zapp, P., Schreiber, A., Marx, J. and Kuckshinrichs, W. (2022). Environmental impacts of rare earth production. MRS Bulletin, 47(3), pp. 267–275. doi: 10.1557/s43577-022-00286-6.

Zewe, A. (2025). Explained: Generative AI's environmental impact. MIT News. https://news.mit.edu/2025/explained-generative-ai-environmental-impact-0117