AI Use, Environmental Awareness and Moral Disengagement Among Gen Z Business Students

Dovilé Daveluy and Adib Bensalem
ESCE International Business School, Paris, France

dovile.daveluy@esce.fr
adib.bensalem@esce.fr

Abstract: Generative artificial intelligence (AI) has become a defining feature of how Generation Z business students learn, work, and produce knowledge. It is also among the most resource-intensive consumer technologies ever deployed, requiring substantial water and electricity for both training and use. Yet the environmental costs of AI remain largely invisible to its end users, and existing AI literacy frameworks systematically exclude environmental considerations from their account of responsible AI use. Drawing on four focus groups and 42 open-response questionnaires conducted at an international business school, this study examines how environmentally aware Gen Z business students reconcile intensive AI use with their stated environmental commitments. Using reflexive thematic analysis of 395 double-coded segments, we identify a four-stage moral disengagement architecture for AI sustainability. The architecture begins with productivity primacy with AI use framed as professionally unavoidable, proceeds through rationalisation and collective disengagement, then responsibility delegation to governments and corporations, and culminates in technological solutionism, in which future innovation is expected to resolve environmental concerns. Each stage closes a specific cognitive and moral gap left open by the previous one, producing an internally coherent structure that allows pro-environmental values to coexist with high-intensity AI use without moral incongruence. The findings also reveal asymmetry in AI education. Participants demonstrated awareness of cognitive and academic risks associated with AI, largely reflecting institutional messaging, while environmental risks remained fragmented, informal, and weakly connected to behavioural decision-making. Although many students expressed concern about sustainability, environmental considerations rarely entered their ethical evaluation of AI use. The study extends Bandura's moral disengagement theory by demonstrating its sequential, architecture-like operation in a high-uptake technology context. It also proposes Green AI Literacy as a missing dimension of existing AI literacy frameworks. Finally, it provides preliminary insights into how higher education institutions might better integrate environmental sustainability into AI literacy education.

Keywords*:* Green AI literacy, Environmental ethics, Moral disengagement, Generation Z, Business education, Sustainability

1. Introduction

A single query to a generative AI model consumes water in the milliliter range and emits carbon in the gram range (Sidorkin, 2025). Scaled across the hundreds of millions of users now interacting with these tools daily, the aggregate footprint already rivals the annual emissions of medium-sized nations (Guidi et al, 2024; Yang et al, 2024). Yet for most users, especially Generation Z, the cohort that has integrated these tools most deeply into everyday life, AI remains experientially weightless. Nothing is extracted, no waste is visible, no meter ticks on the screen. The "black box" mythology of AI (Brevini, 2020) operates not only at the level of algorithms, but at the level of matter where the physical infrastructure of generative AI is doubly invisible.

This invisibility is reproduced inside educational frameworks preparing users to engage with AI. Foundational AI literacy frameworks (Long and Magerko, 2020; Ng et al., 2021; Pinski and Benlian, 2024) and more recent multidimensional models (Hackl et al., 2026) consistently include cognitive, practical, critical and ethical competencies, yet environmental sustainability remains absent. The OECD/European Commission AI Literacy Framework draft (OECD-EC, 2025) acknowledges environmental costs but does not operationalise them as a distinct literacy domain. AI users are taught to evaluate fairness, accountability, and academic integrity, but not ecological footprint.

This omission is especially consequential for Gen Z, a cohort of digital natives (Chen, & Ha, 2023; Yadav, & Rai, 2017; Zhu, 2023), characterised by sustainability values (Jezierski, 2025), and a persistent attitude–behaviour gap when convenience and utility dominate decision-making (Evanilay et al., 2024; Williams and Hodges, 2020). This tension is emerging in AI use, when environmentally aware students rely heavily on resource-intensive tools without a framework for evaluating the trade-off. How they reconcile this tension remains empirically underexplored. System-level scholarship on AI's footprint dominates the literature (Berthelot et al, 2024; Verdecchia et al, 2023), while user-level evidence is scarce.

This study examines how Gen Z business students reconcile intensive use of generative AI with their stated environmental values, using a qualitative design. Three research questions guide the inquiry:

(RQ1) How knowledgeable are Gen Z business students about the environmental impacts of AI?

(RQ2) How do they navigate the trade-off between AI-driven productivity benefits and environmental costs?

(RQ3) How can higher education curricula empower them to be more eco-conscious users of AI?

The paper makes three contributions. First, it identifies a four-stage moral disengagement architecture for AI sustainability that systematically neutralises the values-to-behaviour link predicted by Stern's (2000) Value-Belief-Norm theory. Drawing on Bandura's (1999, 2002) mechanisms of moral disengagement and Morozov's (2013) critique of technological solutionism, we show that these tactics operate sequentially, with each stage closing a moral gap left open by the previous one. The study extends moral disengagement theory by showing how disengagement can function as a structured architecture with an invariant entry point and final stage rather than a static set of parallel mechanisms. Second, we propose Green AI Literacy as a missing dimension of existing AI literacy frameworks. Finally, while the study does not primarily focus on educational interventions, the findings provide preliminary insights into how higher education institutions might better incorporate environmental sustainability into AI literacy education.

2. Background

2.1 AI's Environmental Footprint and the Rebound Paradox

AI's environmental costs are real, growing, and structurally embedded. At the micro level, large generative models consume water in milliliter ranges per query (Sidorkin, 2025). At infrastructure scale, data centers are projected to consume up to 8% of U.S. electricity demand by 2030 (Jha et al, 2025), with the aggregate carbon footprint of major AI systems already comparable to medium-sized nations (Yang et al, 2024). A life-cycle perspective further reveals that operational measurements underestimate true impact by overlooking hardware production, rare metal extraction, and embodied emissions (Berthelot et al, 2024). The empirical evidence is inconclusive. Under some conditions, AI deployment is associated with reductions in carbon emissions through energy efficiency and industrial upgrading (Ding et al, 2023; Wang, Li and Li, 2024). But a critical pattern complicates these findings. Following Jevons' (1865) observation, efficiency gains in resource-intensive technologies typically generate rebound effects. As technologies become more efficient and convenient, they increase overall demand, causing aggregate consumption to rise (Gillingham, 2016; Sorrell, 2009). Alnafrah (2025) provides the first systematic evidence that the same dynamic operates for AI. Efficiency gains are more than offset by expanded consumption, with a threshold beyond which environmental damage accelerates. The implication is that AI's environmental trajectory is shaped by how users, organisations and curricula frame the productivity–sustainability trade-off, not only by technical model design.

2.2 Generation Z's Environmental Values and the Attitude–behaviour gap

Generation Z is consistently characterised as proclaiming sustainability values (Gomes et al., 2023; Jezierski, 2025). Yet research also documents a persistent attitude–behaviour gap, most visibly in fast fashion, where environmental awareness coexists with high consumption when affordability, immediacy and peer norms dominate decision-making (Evanilay et al, 2024; Niinimäki et al, 2020). Williams and Hodges (2020) frame this as a structural feature of Gen Z behaviour rather than a personal failure. This generation has grown up immersed in digital media, with worldviews and behaviors significantly shaped by social media content (Chen & Ha, 2023; Confetto et al., 2023; Yadav & Rai, 2017; Zhu, 2023). This digital environment not only shapes consumption patterns but also influences how emerging technologies are perceived, adopted, and integrated into everyday life. Consequently, many conditions sustaining the attitude–behaviour gap in consumer contexts are also present—and potentially amplified—in AI use. AI tools are free at the point of use, deliver visible productivity gains, are normalised through peer adoption, and confer competitive advantage in professional settings (Chan and Hu, 2023). At the same time, AI's environmental costs are even less visible to end users than those of fast fashion: no garment to touch, no label to read, no supply chain narrative in mainstream media.

2.3 The Missing Environmental Dimension

AI literacy has emerged as a priority concept across education and policy (Long and Magerko, 2020; Ng et al, 2021; Pinski and Benlian, 2024). Dominant frameworks converge on cognitive, technical, ethical and societal dimensions, with ethics typically described in terms of fairness, transparency, accountability, privacy and safety. The environmental ethics of AI (resource extraction, water consumption, energy demand, electronic waste and unequal geographical burdens) are not foregrounded in these frameworks. Even Hackl et al.'s (2026) AI Literacy Heptagon, a comprehensive recent synthesis, lists sustainability only at the periphery.

Two recent developments hint at change: the OECD-EC (2025) framework explicitly requires learners to weigh the environmental cost of AI use, and Cheng et al. (2025) demonstrate that AI literacy affects green absorptive capacity. Both are partial since the former is under-operationalised, while the latter operates at the organisational, not individual level. This paper proposes Green AI Literacy as the missing individual-level dimension.

2.4 Theoretical Framing

To explain how environmentally aware users sustain intensive AI use without experiencing moral incongruence, we draw on three complementary traditions. Bandura's (1999, 2002) *moral disengagement theory* specifies psychosocial mechanisms (moral justification, advantageous comparison, distortion of consequences, displacement and diffusion of responsibility, among others) by which individuals selectively deactivate their internal moral standards while continuing to act inconsistently with them. The framework has become central in business ethics for explaining participation in organisational wrongdoing (Detert, Treviño and Sweitzer, 2008; Moore, 2008). Stern's (2000) *Value-Belief-Norm* (VBN) theory specifies the causal chain through which biospheric values shape beliefs about consequences, ascription of responsibility, personal norms, and ultimately behaviour, while identifying where the chain can break. Morozov's (2013) critique of technological solutionism describes how present moral problems are reframed as future technical solutions. We argue that these mechanisms operate not as alternatives but as a sequential architecture in which each stage closes a moral gap left open by the previous one.

3. Methodology

This study adopts a qualitative, multi-method research design grounded in an interpretivist epistemology. Issues raised by the research questions invoke forms of ethical sensemaking best surfaced through dialogic and reflective formats that allow participants to articulate value tensions in their own terms (Lim, 2024). To triangulate across discursive formats and reduce limitations of any single qualitative instrument (Milfont, 2008), the study combines two complementary data sources: four focus group discussions and 42 individual qualitative questionnaires.

Participants were recruited from business programmes at ESCE International Business School in France. The questionnaire sample comprised 42 participants. The four focus groups were composed of students who had also completed the questionnaire, yielding a total of 42 unique participants across both data sources rather than two separate samples, enabling within-subject triangulation. Three focus groups comprised generalist business students, while the fourth (TMIE) included students with substantial ESG exposure. The TMIE group is referenced only where contrasts with the other groups provide analytical insight. The focus group guide progressed from descriptions of AI use, through elicitation of environmental knowledge, to projective and normative scenarios. The open-response questionnaire combined five Likert items (T1–T5) with seven scenario questions (Q1–Q7). The analytical centerpiece is the Alex scenario (Q1), a structured vignette in which participants evaluate a fictional master's student who used generative AI to produce substantial portions of his thesis. Q4 directly asked whether environmental cost should figure in this ethical evaluation.

Data were analysed using Braun and Clarke's (2006; 2021) six-phase reflexive thematic analysis. Coding was conducted independently by both authors across the entire corpus, yielding 395 individually tagged segments. Disagreements were systematically tracked, and consensus codes were established through structured discussion. The nine final themes (T1–T9) emerged from the consensus codebook.

4. Findings

The thematic analysis produced nine themes organised into three families mapped onto the research questions: knowledge (T1–T3, RQ1), the four-stage trade-off architecture (T4–T7, RQ2), and curricular vision (T8–T9, RQ3).

4.1 Fragmented Awareness Without Ascribed Responsibility (T1–T3)

Baseline factual knowledge of AI's environmental footprint was limited and fragmented. Many participants disclosed having no prior awareness about AI's environmental impact *("I totally was not aware of that until I started looking at this questionnaire")*. Participants' knowledge mostly concerned water usage and energy consumption, typically expressed in vague terms ("*AI… produces carbon emissions. This is the only thing I know"; "I think that AI uses a lot of water just to cool the servers*"). Knowledge sources were overwhelmingly informal. Social media, memes and documentaries dominated, while formal education was a

source only within the TMIE group. Crucially, even when awareness was present, it lacked behavioural salience, as participants frequently described environmental impact as distant and having minimal bearing on their choices ("*Honestly, I do not really feel it in a tangible way*"; "*Out of sight, out of mind*"). The Q4 coding probed whether environmental cost should figure in the evaluation of Alex's vignette: 37 codes treated it as a background fact; only 9 - as a genuine ethical factor. These findings confirm Brevini's (2020) argument that AI is mythologized as immaterial and "clean," while its environmental infrastructure remains largely hidden. Moreover, while surface-level knowledge is present, moral weight does not attach to it. In Stern's (2000) terms, the consequence-belief node of the VBN chain fails to crystallise.

4.2 The Four-stage Moral Disengagement Architecture (T4–T7)

Analysis identified a stable four-stage structure through which participants reconcile their AI use with their environmental commitments. Each stage resolves a moral challenge left unresolved by the preceding one (see Figure 1).

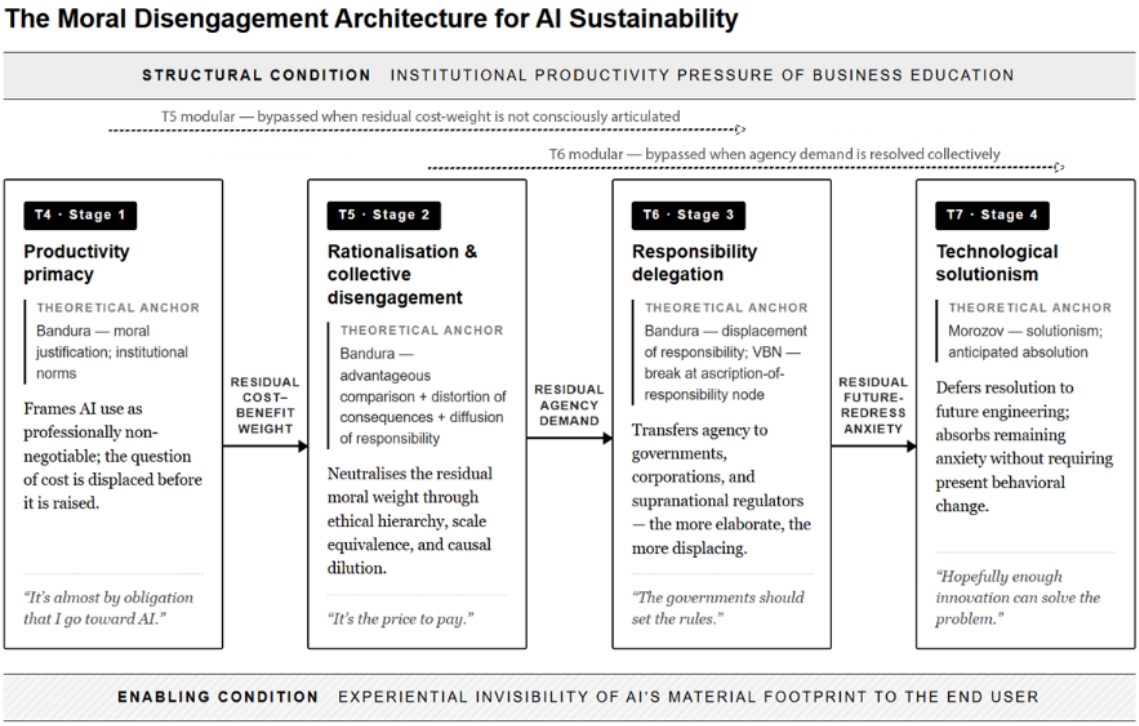


Figure 1: The moral disengagement architecture for AI sustainability

T4 — Productivity primacy. The architecture's entry point is the framing of AI use as professionally and competitively unavoidable. The vocabulary is consistent: AI is necessary, obligatory, expected ("*It's almost by obligation that I go toward AI*"; "*Whoever masters this technology will have an advantage over everyone else*"). Theoretically, this stage operates as moral justification (Bandura, 1999): a worthy purpose displaces the question of environmental cost. Notably, the source of justification is not personal preference, but the institutional norms — a socially distributed cognitive operation rather than an individual one.

T5 — Rationalisation and collective disengagement. A second cluster of responses minimizes the moral weight of acknowledged environmental cost. The most direct form of rationalization was cost acceptance: *"It's the price to pay."* Participants also engaged in advantageous comparison, framing environmental impact as less important than other ethical considerations. Participants overwhelmingly evaluated Alex's behavior through the lens of academic integrity rather than environmental cost ("*The ethical breach of academic integrity carries more weight than the carbon footprint*"). Environmental consequences were further minimized through scale equivalence, whereas AI-related impacts were presented as equivalent to other digital activities ("*everyday usage either for this or for other things generates pollution*"). Finally, participants downplayed their own contribution by emphasizing the collective scale of AI use, thereby reducing their own responsibility ("*If I change or stop using AI, others are still going to use AI*").

T6 — Responsibility delegation. The third stage transfers ethical agency from the individual user to institutional actors. T6 is the most discursively elaborated theme (39 codes in the focus group data, nearly double the codes associated with personal strategies under T9). Participants spoke extensively about who should act ("G*overnments should set the rules*"; "*Companies can adapt and try to build greener AI*"). Theoretically, T6 corresponds to Bandura's displacement of responsibility at the individual level and diffusion at the collective level, and to the failure of Stern's (2000) ascription-of-responsibility node in the VBN chain. The sophistication of the institutional vocabulary itself functions as the displacement mechanism: the more elaborate the attribution of agency to institutional actors, the less personal agency remains.

T7 — Technological solutionism. The fourth stage relocates resolution to a future technical horizon ("*Maybe we can find another solution to cool the servers*"; "*Hopefully enough innovation can solve the problem*"). Strikingly, no participant cited a documentary or course as the source of the "underwater data centers" narrative; it appeared independently across multiple groups, consistent with technological solutionism circulating in mainstream tech media and Gen Z's digital ecosystems. Optimism here is socially and culturally constructed rather than institutionally taught. Theoretically, T7 reflects Morozov's (2013) technological solutionism, which defers the moral obligation to act to future technological solutions. Unlike the preceding stages, this mechanism does not require deactivation of environmental concern but redirects it toward anticipated absolution.

Taken together, the four stages form a coherent architecture with two invariant anchors. T4 (productivity primacy) functions as the universal entry gate: in every instance of architecture-stage reasoning across both data sources, T4 appeared before any other stage, with zero violations. T7 (technological solutionism) functions as the final absorber: it appeared exclusively in focus group data (11 utterances), consistent with a discursive move that becomes necessary only once prior stages have been deployed and challenged. Between these anchors, T5 (rationalisation) and T6 (responsibility delegation) are modular: participants may deploy either or both depending on which moral challenges remain salient after the entry stage. The architecture does not require all four stages to be traversed in strict sequence, but it is functionally complete only when all four are available. By the final stage, users can simultaneously hold pro-environmental values and sustain high-intensity AI use without experiencing moral incongruence.

4.3 Vision for Green AI Literacy (T8–T9)

Participants characterised their AI education as overwhelmingly productivity-oriented ("*They train us to be more performant*"; "*They teach us how to save time with AI*") and identified environmental impact as a curricular gap ("*The school hasn't prepared me for the environmental cost of AI*"). T9, the largest theme in the dataset (98 codes), captures demand for an education that integrates efficiency with environmental responsibility ("*We should be taught how to use it efficiently, without creating negative secondary effects*"; "*Students should know when AI is necessary and when it is not*"). These demands map directly onto the three dimensions of the Green AI Literacy construct: cognitive, affective, and behavioural.

The TMIE focus group exhibited deeper knowledge under T1 and more nuanced ethical reasoning under Q4, suggesting that curricular intervention is consequential, even though our design does not support causal attribution.

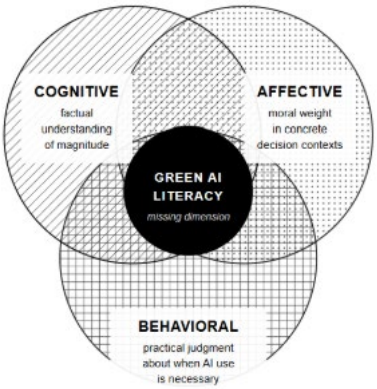


Figure 2: The Green AI Literacy construct

Across the data, students articulated significant institutionally induced anxiety about AI's effects on cognition and integrity, but only fragmentary, self-constructed concern about its environmental effects. This cognitive–

environmental asymmetry between institutionally taught risks and those learned by default from social media is depicted in Figure 3.

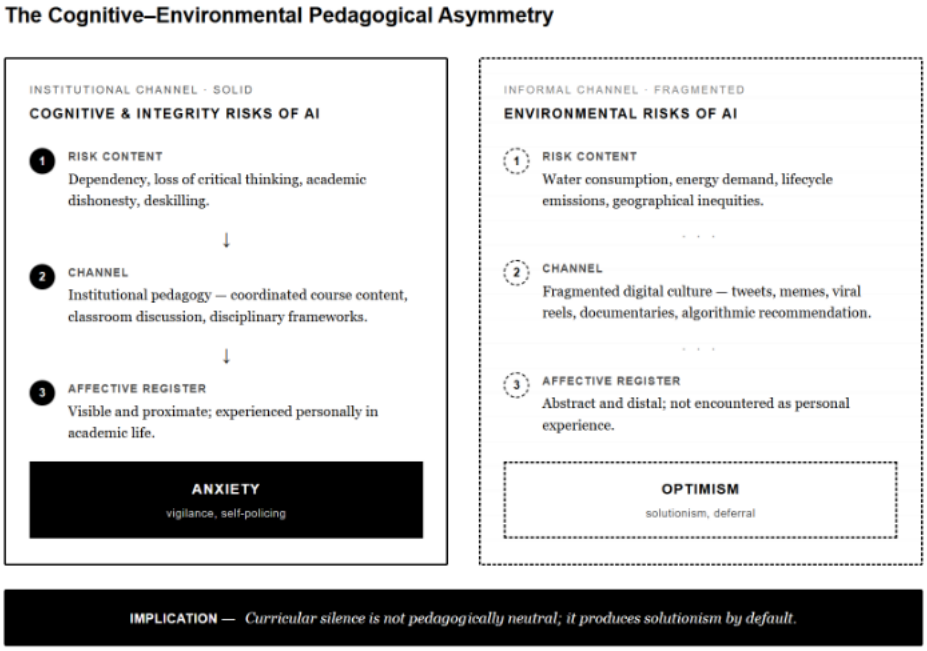


Figure 3: The cognitive–environmental pedagogical asymmetry

5. Discussion

The findings extend Bandura's moral disengagement framework in a direction the literature has not pursued. Bandura's mechanisms have been applied across business ethics contexts to identify which mechanism operates in any given setting. Our findings document an empirical setting in which the mechanisms operate not as alternatives but as a sequential structure: each resolves an open cognitive gap, and the architecture is functionally complete only when all operations are available together. This is the paper's principal contribution to theory: moral disengagement in high-uptake technology contexts operates through a structured architecture with an invariant entry gate (T4) and final absorber (T7), not merely as a menu of co-available mechanisms. It also specifies where Stern's (2000) VBN chain breaks under conditions of high institutional pressure - at the ascription-of-responsibility node, where Bandura's displacement and the institutional sophistication of business education intersect.

The findings support Green AI Literacy as a missing dimension of AI literacy. Building on D'Attoma and Ieva's (2024) multidimensional environmental consciousness construct and aligned with the Affective-Behavioural-Cognitive-Ethical (ABCE) scaffolding of Ng et al (2023), we define Green AI Literacy as the integrated capacity to recognise, evaluate, and incorporate AI environmental consequences into behavioral decisions. It comprises three interrelated dimensions: a *cognitive* dimension (factual understanding of magnitude and mechanism), an *affective* dimension (moral weight attached to decisions), and a *behavioural* dimension (practical judgement about when AI use is necessary). Rather than replacing existing AI literacy dimensions, Green AI Literacy complements them by addressing an environmental blind spot in current frameworks. Empirically, participants demonstrated fragmented but often accurate awareness of AI's environmental impacts (T1), and many articulated concrete strategies for more responsible use (T9). The primary breakdown occurred at the affective level, where environmental knowledge rarely translated into ethical concern or behavioural restraint. Green AI Literacy therefore addresses not only what users know about AI's environmental footprint, but also whether they recognise its relevance for their behaviour.

The implications extend beyond education. If the cognitive architecture identified in this study carries forward into participants' future managerial roles, the environmental ethics of corporate AI adoption may be governed by the same four-stage structure. Therefore, educational and corporate interventions are likely to be more effective if designed to interrupt one or more stages of this architecture.

6. Implications and Limitations

Implications. We suggest that Green AI Literacy should be embedded as a core competency within business education rather than acquired informally through fragmented external sources. The cognitive dimension can be taught through magnitude calibration (lifecycle assessments, per-query versus aggregate framings, the rebound paradox). The affective dimension requires engagement with ethical hierarchy that subordinates environmental costs. The behavioural dimension can be supported through tools that make environmental cost visible at the point of use. None of these interventions are technically demanding; they require the recognition that environmental costs belong within institutional AI conversations.

Limitations. The sample is drawn from a single French international business school. Therefore, documented architecture may reflect that institutional culture. Focus groups and self-reported questionnaires remain vulnerable to social desirability bias (Milfont, 2008). The study captures reasoning at a single point in time. The focus on business students excludes other professional cohorts whose reasoning architectures may differ in instructive ways. While the TMIE contrast is suggestive, the study was not designed as a comparative intervention, and the associated effects cannot be causally attributed. Future research should pursue quantitative validation of the four-stage architecture, cross-cultural and cross-institutional replication, intervention studies converting the TMIE finding into causal evidence, and longitudinal follow-up into professional roles.

Ethics Declaration*:* Ethical clearance was obtained through ESCE International Business School. All participants provided written informed consent using the ESCE 2026 consent form, allowing recording, transcription and anonymised publication of quoted material. Participation was voluntary, with no compensation, and participants retained the right to withdraw at any stage. Identifying information was stripped at intake and replaced with anonymised identifiers prior to analysis. All quoted materials have been anonymised; transcripts are retained by the researchers.

AI Declaration: The authors used Claude AI (Anthropic, Opus 4.7) to support code aggregation, literature identification, drafting, and structural editing. All theoretical claims, coding decisions, empirical analyses, findings, and interpretations are the authors' own. Data collection, coding, and analytical synthesis were conducted entirely by the authors. No empirical material was generated by AI, and all references suggested by Claude were verified against original sources. Full disclosure is consistent with the paper's commitment to responsible and transparent AI use.

The environmental impact of AI-assisted drafting cannot be determined precisely. Anthropic neither publishes nor displays per-token environmental metrics. Based on estimates in 2.1 (Sidorkin, 2025), the session consumed approximately 1.5–5.5 liters of water and generated 0.25–3 kg CO_2-equivalent emissions. Until such metrics are visible at the point of use, even researchers studying Green AI Literacy cannot confidently assess the footprint of their AI use.

References

Alnafrah, I. (2025) "The two tales of AI: A global assessment of the environmental impacts of artificial intelligence from a multidimensional policy perspective", *Journal of Environmental Management*, Vol. 392, p 126813.

Bandura, A. (1999) "Moral disengagement in the perpetration of inhumanities", *Personality and Social Psychology Review*, Vol. 3, No. 3, pp 193–209.

Bandura, A. (2002) "Selective moral disengagement in the exercise of moral agency", *Journal of Moral Education*, Vol. 31, No. 2, pp 101–119.

Berthelot, A., Caron, E., Jay, M. and Lefèvre, L. (2024) "Estimating the environmental impact of generative-AI services using an LCA-based methodology", *Procedia CIRP*, Vol. 122, pp 707–712.

Braun, V. and Clarke, V. (2006) "Using thematic analysis in psychology", *Qualitative Research in Psychology*, Vol 3, No. 2, pp 77–101.

Braun, V. and Clarke, V. (2021) *Thematic Analysis: A Practical Guide*, Sage.

Brevini, B. (2020) "Black boxes, not green: Mythologizing artificial intelligence and omitting the environment", *Big Data & Society*, Vol. 7, pp 1-5.

Chan, C.K.Y. and Hu, W. (2023) "Students' voices on generative AI: Perceptions, benefits, and challenges in higher education", *International Journal of Educational Technology in Higher Education*, Vol. 20, No.1, p.43.

Chen, P. and Ha, L. (2023) "Gen Z's social media use and global communication", *Online Media and Global Communication*, Vol. 2, pp 301–303.

Cheng, J., Xu, N.-R., Khan, N.U. and Singh, H. (2024) "The impacts of artificial intelligence literacy, green absorptive capacity, and green information system on green innovation", *Corporate Social Responsibility and Environmental Management*, Vol. 32, No.2, pp 2375-2389.

Confetto, M.G., Covucci, C., Addeo, F. and Normando, M. (2023) "Sustainability advocacy antecedents: How social media content influences sustainable behaviours among Generation Z", *Journal of Consumer Marketing*, Vol 40, No. 6, pp 758-774.

D'Attoma, I. and Ieva, M. (2024) "A new composite index to assess environmental consciousness using survey data and big data: Empirical evidence from European consumers", *Socio-Economic Planning Sciences*, Vol. 95, p.102038.

Detert, J.R., Treviño, L.K. and Sweitzer, V.L. (2008) "Moral disengagement in ethical decision making: A study of antecedents and outcomes", *Journal of Applied Psychology*, Vol. 93, No. 2, pp. 374–391.

Ding, T., Li, J., Shi, X., Li, X. and Chen, Y. (2023) "Is artificial intelligence associated with carbon emissions reduction? Case of China", *Resources Policy*, Vol. 85, p.103892.

Evanilay, E., Lee, M.D., Millenia, M. and Murwani, I.A. (2024) "Why Generation Z still buy fast fashion although environmental awareness is high?", *Scholars Journal of Economics, Business and Management*, Vol. 11, No. 9, pp. 268–278.

Gillingham, K., Rapson, D. and Wagner, G. (2016) "The rebound effect and energy efficiency policy", *Review of Environmental Economics and Policy*, Vol. 10, No. 1, pp. 68–88.

Gomes, S., Lopes, J. and Nogueira, S. (2023) "Willingness to pay more for green products: A critical challenge for Gen Z", *Journal of Cleaner Production*, Vol. 390, p. 136092.

Guidi, G.M., Dominici, F., Gilmour, J., Butler, K., Bell, E., Delaney, S. and Bargagli Stoffi, F.J. (2024) "Environmental burden of United States data centers in the artificial intelligence era", [online], *arXiv preprint arXiv:2411.09786*.

Hackl, V., Mueller, A.S. and Sailer, M. (2026) "The AI literacy heptagon: A structured approach to AI literacy in higher education", *Computers and Education: Artificial Intelligence*, Vol. 8, p. 100540.

Jevons, W.S. (1865) *The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal Mines*, Macmillan.

Jezierski, D. (2025) "Generation Z's pro-environmental behaviour", *Humanities and Social Sciences Quarterly*, Vol. 32, No. 2, pp. 67–84.

Jha, R., Jha, R. and Islam, M. (2025) "Forecasting US data center CO_2 emissions using AI models: Emissions reduction strategies and policy recommendations", *Frontiers in Sustainability*, Vol. 5, p. 1423456.

Li, P., Yang, J., Islam, M.A. and Ren, S. (2025) "Making AI less 'thirsty': Uncovering and addressing the water footprint of AI models", *Communications of the ACM*, Vol. 68, No. 7, pp. 54–62.

Lim, W.M. (2024) "What is qualitative research? An overview and guidelines", *Australasian Marketing Journal*, Vol. 33, No. 2, pp. 199–229.

Long, D. and Magerko, B. (2020) "What is AI literacy? Competencies and design considerations", in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–16.

Milfont, T.L. (2009) "The effects of social desirability on self-reported environmental attitudes and ecological behaviour", *The Environmentalist*, Vol. 29, No. 3, pp. 263–269.

Moore, C. (2008) "Moral disengagement in processes of organizational corruption", *Journal of Business Ethics*, Vol. 80, No. 1, pp. 129–139.

Morozov, E. (2013) *To Save Everything, Click Here: The Folly of Technological Solutionism*, PublicAffairs.

Ng, D.T.K., Leung, J., Chu, S.K.W. and Qiao, M.S. (2021) "Conceptualizing AI literacy: An exploratory review", *Computers and Education: Artificial Intelligence*, Vol. 2, p. 100041.

Ng, D.T.K., Wu, W., Leung, J., Chiu, T. and Chu, S.K.W. (2024) "Design and validation of the AI literacy questionnaire: The affective, behavioural, cognitive and ethical approach", *British Journal of Educational Technology*, Vol. 55, No. 3, pp. 1082–1104.

Niinimäki, K., Peters, G., Dahlbo, H., Perry, P., Rissanen, T. and Gwilt, A. (2020) "The environmental price of fast fashion", *Nature Reviews Earth & Environment*, Vol. 1, No. 4, pp. 189–200.

OECD and European Commission (2025) *Empowering Learners for the Age of AI: An AI Literacy Framework for Primary and Secondary Education (Review Draft)*, [online], AILitFramework_ReviewDraft.pdf.

Pinski, M. and Benlian, A. (2024) "AI literacy for users: A comprehensive review and future research directions of learning methods, components, and effects", *Computers in Human Behavior: Artificial Humans*, Vol. 2, No. 1, p. 100062.

Sidorkin, A.M. (2025) *Environmental Impact of Generative AI: Carbon and Water Footprint*, National Institute on Artificial Intelligence in Society, [online], AI-EDU Arxiv.

Sorrell, S. (2009) "Jevons' paradox revisited: The evidence for backfire from improved energy efficiency", *Energy Policy*, Vol. 37, No. 4, pp. 1456–1469.

Stern, P.C. (2000) "New environmental theories: Toward a coherent theory of environmentally significant behavior", *Journal of Social Issues*, Vol. 56, No. 3, pp. 407–424.

Verdecchia, R., Sallou, J. and Cruz, L. (2023) "A systematic review of green AI", *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, Vol. 13, No. 4, p. e1507.

Wang, Q., Li, Y. and Li, R. (2024) "Ecological footprints, carbon emissions, and energy transitions: The impact of artificial intelligence (AI)", *Humanities and Social Sciences Communications*, Vol. 11, No. 1, pp. 1-18.

Williams, A. and Hodges, N.J. (2020) "Generation Z and socially responsible fashion consumption: Exploring the value-action gap", *International Textile and Apparel Association Annual Conference Proceedings*, Vol. 77, No. 1, Iowa State University Digital Press, [online].

Yadav, G.P. and Rai, J. (2017) "The Generation Z and their social media usage: A review and a research outline", *Global Journal of Enterprise Information System*, Vol. 9, pp. 110–116.

Yang, Y., Wang, J., Liu, Y., Yu, P., Wang, D., Zheng, P. and Zhang, M. (2024) "Revisit the environmental impact of artificial intelligence: The overlooked carbon emission source?", *Frontiers of Environmental Science & Engineering*, Vol. 18, No. 12, p. 158.

Zhu, Y. (2023) "The impact of media on Generation Z and possible future social changes", *Communications in Humanities Research*, Vol. 22, No.1, pp. 36-41.