

Knowing but Not Doing: The Awareness-Practice Gap in Responsible and Green AI Use

Sirje Virkus, Muhammad Shaheer Alam

School of Digital Technology

Tallinn University

Narva mnt 25, 10120 Tallinn, Estonia

{sirje.virkus, shaheer}@tlu.ee

Abstract. *Generative artificial intelligence is now routine in university work, but responsible use remains inconsistent. This paper examines how students in AI-related programmes at Tallinn University understand and practise Green AI and Responsible AI. In this exploratory study of 57 students and six interview participants from two AI-related programmes at Tallinn University, participants generally recognised ethical, privacy and environmental concerns, but reported that efficiency and convenience sometimes took priority over verification, transparency, privacy protection and sustainable use. The findings suggest that universities may need to move beyond awareness raising by embedding responsible AI routines, institutional guidance and supportive tool design into academic practice.*

Keywords. Green AI, Responsible AI, Generative AI, higher education, student behaviour, awareness-practice gap, AI ethics, sustainability

1 Introduction

Generative artificial intelligence (GenAI) tools such as ChatGPT, virtual assistants and writing support tools are now embedded in the ordinary workflow of many university students. They are used to summarise readings, draft text, generate ideas, explain difficult concepts and accelerate repetitive academic tasks (Athanassopoulos et al., 2026; Deng et al., 2025; Xiao et al., 2025). This uptake has produced an important pedagogical tension. On one side, artificial intelligence (AI) tools can improve access, productivity and confidence. On the other side, they introduce risks for academic integrity, privacy, intellectual dependence, bias, accountability and environmental sustainability (Francis et al., 2025).

This paper is based on the study that examined this tension through the combined lens of Green AI and Responsible AI. Green AI refers to the development and use of AI technologies that minimise environmental impact by reducing energy consumption and carbon emissions (Tabbakh et al., 2024). Responsible AI addresses the ethical and

societal conditions under which AI should be designed and used, including fairness, accountability, transparency, privacy and human oversight (Raman et al., 2024). In higher education these two areas intersect because student behaviour is not merely a private matter. Repeated daily prompting, uncritical reliance on outputs and casual sharing of data can accumulate into broader institutional and societal consequences.

A central problem is that knowledge of AI risks does not automatically produce responsible action. Students may know that AI outputs should be checked, that private data should not be uploaded, or that unnecessary use has environmental costs, while still bypassing these practices when under pressure. The empirical question is therefore not only whether students have heard of Green AI or Responsible AI, but how such awareness appears in actual academic routines and what conditions enable or inhibit responsible behaviour.

This paper focuses on four research questions: 1) What level of awareness do students report regarding Green AI and Responsible AI? 2) How do students report applying Green AI and Responsible AI principles in academic work? 3) What gaps emerge between students' awareness and reported practice? and 4) What contextual factors influence responsible AI behaviour?

This paper contributes to higher education AI research in three ways: 1) It connects Responsible AI and Green AI through the concept of routine academic practice. 2) It provides exploratory evidence of an awareness-practice gap among students in two AI-related programmes at one university. 3) It identifies contextual factors that participants associated with this gap, including time pressure, institutional ambiguity and tool design. The paper contributes to discussions on AI literacy by arguing that literacy must include situated behavioural competence, not only conceptual familiarity.

2 Related Work

The rapid development of GenAI in higher education has generated substantial interest in its pedagogical

opportunities, institutional implications and ethical risks. Data from the Web of Science underscores this trend: between 1981 and May 2026, 35,815 publications were indexed under ‘artificial intelligence and education.’ The higher education sector mirrors this growth, contributing 6,323 publications to the field.

Recent studies show that students use GenAI tools for a wide range of academic purposes, including writing support, summarisation, programming assistance, idea generation and personalised explanation (Athanasopoulos et al., 2026; Deng et al., 2025; Xiao et al., 2025). These uses can support productivity, accessibility and confidence, particularly when students face complex tasks or time constraints. At the same time, the literature identifies significant concerns, including plagiarism, over-reliance, reduced critical engagement, inaccurate or fabricated outputs, unequal access and uncertainty about acceptable academic practice (Kasneci et al., 2023; Dwivedi et al., 2023). A recurring issue is that student adoption often develops faster than institutional policy and classroom guidance, leaving students to interpret expectations through peers, lecturers and tool affordances (Nikolic et al., 2024; Stracke et al., 2026).

Responsible AI provides an important framework for addressing these challenges. Foundational AI ethics frameworks emphasise principles such as fairness, accountability, transparency, privacy, explicability, human oversight and non-maleficence (European Commission, 2019; Floridi et al., 2018; Floridi & Cowls, 2019; Jobin et al., 2019; OECD, 2026). In higher education, these principles become relevant not only at the level of institutional governance but also in students’ everyday academic practices. Responsible student use of AI may involve checking AI-generated claims, recognising possible bias, avoiding the upload of sensitive data, disclosing AI assistance where required, and retaining human judgement and authorship over academic work (European Commission, 2024; Wang et al., 2024). However, a persistent challenge is operationalisation. Ethical principles are often easier to endorse in general terms than to apply consistently in concrete, time-pressured academic situations (Mittelstadt, 2019; Morley et al., 2023).

Green AI adds a further dimension to responsible AI use by foregrounding the environmental costs of AI systems. Green AI research has moved from a model-centred concern with reducing the computational cost of training toward a broader lifecycle perspective. Early work challenged accuracy-at-any-cost AI and promoted reporting energy and compute use. Subsequent research developed optimisation methods such as pruning, quantisation, compression and efficient algorithms. More recent work expands the field to hardware, data centres, renewable energy and system-level design, while also shifting attention from training to inference and large-scale real-world use. This evolution shows that Green AI is no longer only a

technical efficiency agenda, but an emerging sociotechnical concern involving developers, infrastructure providers, institutions and users (Tabbakh et al., 2024). This broader lifecycle perspective is particularly relevant for higher education, where frequent low-stakes AI use may appear individually insignificant but can become consequential when repeated across large student populations. For students, however, the environmental cost of AI use is usually invisible, indirect and difficult to estimate, which may weaken its influence on everyday decisions. At the student-use level, Green AI is operationalised in this study through observable or reportable choices that may reduce avoidable computational demand: selecting a lightweight tool for simple tasks, avoiding unnecessary regeneration, combining related requests, limiting redundant prompts, reusing adequate outputs, and considering energy efficiency when choosing between tools.

The relationship between awareness and behaviour is a useful lens for understanding this problem. Environmental and ethical behaviour research has long shown that knowledge, concern or positive attitudes do not automatically translate into consistent action. The value-action gap, or attitude-behaviour gap, describes the mismatch between what individuals know or value and what they actually do (Blake, 1999; Kollmuss & Agyeman, 2002). Similar patterns have been identified in ethical consumption, where expressed ethical intentions often fail to produce corresponding behaviour (Carrington et al., 2014). Behavioural theories suggest that action is shaped not only by attitudes, but also by subjective norms, perceived behavioural control, habits, convenience and feedback mechanisms (Ajzen, 1991; Stern, 2000; Verplanken & Wood, 2006). Applied to AI use, students may recognise that outputs should be verified, sensitive data should be protected and unnecessary AI use should be avoided, yet still fail to act accordingly when deadlines, unclear expectations or convenient interfaces make responsible behaviour more difficult.

This perspective is especially important because much of the existing research on generative AI in higher education focuses on perceptions, attitudes, benefits, risks, intentions or institutional policies (Almassaad et al., 2024; Chan & Hu, 2023; Fošner, 2024). Related work has also developed AI ethics guidelines for education, emphasising transparency, privacy, accountability, academic integrity, human oversight, and responsible pedagogical integration (European Commission, 2022; Holmes & Miao, 2023). In parallel, Green AI and sustainable computing research has examined the environmental implications of AI systems, including energy consumption, carbon emissions, infrastructure and technical efficiency (Dalvi-Esfahani et al., 2020; Raman et al., 2024; Schwartz et al., 2020; Tabbakh et al., 2024; Verdecchia et al., 2023). These bodies of work are valuable, but they have often developed separately. Less attention has been given to how ethical and environmental

responsibility intersect in students' everyday academic AI use.

Therefore, the key gap addressed in this study is not simply whether students are aware of Responsible AI or Green AI, but whether such awareness is translated into reported academic behaviour. Existing studies often examine what students think about AI, how they perceive its benefits and risks, or how institutions should regulate its use. Less is known about the relationship between students' awareness and their routine practices, including prompting, verification, disclosure, privacy protection and tool selection. This gap is especially important among students in AI-related programmes, who are frequent users of AI tools and potential future developers, analysts or decision-makers. This exploratory study addresses this gap by examining how students enrolled in two AI-related programmes at Tallinn University report understanding and enacting Responsible AI and Green AI in academic work, and by identifying the contextual factors that enable or constrain responsible and sustainable behaviour.

3 Methodology

This study used an exploratory mixed-methods design to examine how students understand and apply principles of Green AI and Responsible AI in their academic use of AI tools. A mixed-methods approach was appropriate because the study aimed not only to describe students' levels of awareness and reported practices, but also to understand why responsible behaviours may or may not occur in everyday academic contexts. The quantitative component provided a descriptive overview of students' AI use, awareness and self-reported practices, while the qualitative component offered deeper insight into students' reasoning, experiences and perceived barriers to responsible AI use.

The study was conducted among students enrolled in AI-related programmes at Tallinn University. The participants included bachelor's-level Informatics students and master's students in the Artificial Intelligence for Sustainable Societies programme. These groups were selected because they are likely to use AI tools regularly in academic work and may also become future developers, analysts or decision-makers in AI-related fields. The survey sample consisted of 57 respondents, including 28 bachelor's students and 29 master's students. In addition, six students participated in semi-structured interviews. The sample was exploratory and context-specific rather than statistically representative, but it was suitable for identifying patterns of awareness, practice and contextual influence among frequent academic users of AI tools.

Data were collected through an online questionnaire and semi-structured interviews in spring 2026. The questionnaire examined students' frequency of AI use,

types of AI tools used, academic use cases, awareness of Green AI and Responsible AI, knowledge of AI-related energy consumption, environmentally responsible AI use, privacy-related behaviour, verification practices, disclosure of AI use and perceptions of institutional guidance. Several questions used scenario-based items to examine how students responded to concrete academic situations involving privacy, academic integrity, sustainability and tool selection. This helped move beyond general attitudes by asking students to evaluate practical choices in realistic contexts.

The questionnaire contained 46 items organised into five domains: AI-use frequency and purposes; Responsible AI awareness; Green AI awareness; reported responsible practices; and perceptions of institutional and tool-level support. Attitudinal and self-reported behavioural items used a five-point Likert scale ranging from 1 ('strongly disagree'/'never') to 5 ('strongly agree'/'always'), as appropriate. Knowledge and scenario items used categorical response options with one or more responses identified as reflecting the responsible choice.

The items were developed from the study's conceptual framework and informed by literature on Responsible AI, Green AI, privacy, verification, disclosure and the awareness-behaviour gap. The instrument was reviewed by two experts for content relevance.

The semi-structured interviews were used to explore selected survey themes in greater depth. Participants were recruited voluntarily through course-related communication channels and student networks within the relevant programmes. Bachelor's students were primarily approached through course communication channels, while master's students were contacted through programme-related communication and peer networks. The six interview participants had also completed the survey and were selected purposively based on their familiarity with AI tools and willingness to reflect on their personal experiences. Interview questions focused on students' academic use of AI tools, verification of AI-generated outputs, disclosure of AI assistance, handling of sensitive or private data, consideration of environmental impact and perceptions of institutional expectations. The interviews also examined contextual factors such as time pressure, convenience, trust in AI outputs, uncertainty about university rules and the design of AI tools. Before participation, interviewees received a consent form and provided written informed consent. Participants were informed about the purpose of the research, the voluntary nature of participation, anonymity and secure data handling. This qualitative component was important for interpreting the awareness-practice gap, as it allowed students to explain why responsible practices may be difficult to maintain in routine academic work.

Survey data were analysed descriptively. The analysis focused on identifying patterns in students' AI

use, awareness of Green AI and Responsible AI, and self-reported responsible practices. Bachelor's and master's students were compared across Green AI awareness, Responsible AI awareness, privacy-related behaviour and environmentally responsible AI use. Group means and response distributions were examined to identify possible differences between the two academic levels. Because several constructs were measured using single items and the subgroup sizes were small, the comparisons were restricted to descriptive statistics and should not be interpreted as evidence of statistically reliable group differences. The comparative findings were therefore treated as exploratory and were not intended to support broad statistical generalisation. All behavioural measures concerned reported rather than directly observed practices. The study therefore captures participants' accounts and perceptions of their behaviour, not independently verified AI-use routines.

Interview transcripts were analysed using a hybrid deductive-inductive thematic procedure. Initial deductive codes were derived from the research questions and survey domains, including verification, privacy, disclosure, environmental consideration, time pressure and institutional guidance. Additional codes were developed inductively when recurrent ideas emerged from the transcripts. Codes were grouped into broader themes through iterative comparison across participants. Survey and interview findings were then integrated using a theme-by-source matrix that identified convergence, complementarity and divergence between the two datasets.

4 Findings

This section presents the main findings on students' awareness and reported practices concerning Green AI and Responsible AI in academic work.

4.1 AI Use as a Routine Academic Practice

AI use was a regular part of academic work for the majority of respondents. Of the 57 survey participants, 40.4% reported using AI tools daily and 45.6% reported using them several times a week. A further 8.8% used AI tools once a week, while only a very small proportion reported using them a few times a month or rarely. Taken together, more than 85% of respondents used AI tools at least several times a week, indicating that AI was not an occasional or experimental resource but part of students' routine study practices.

ChatGPT was the most widely used tool, selected by 80.7% of respondents. Other tools were also used, including Google Gemini by 47.4%, GitHub Copilot or similar AI code assistants by 35.1%, Grammarly or similar writing tools by 15.8%, and Microsoft Copilot

by 7.0%. A relatively large group, 54.4%, selected "other" tools, suggesting that students' AI use extended beyond the most visible commercial platforms.

Students used AI across a broad range of academic tasks. The most common use case was coding, debugging or code generation, reported by 78.9% of respondents. This was followed by summarising readings or research papers, reported by 77.2%. Writing essays, reports or assignments and using AI for exam preparation or concept explanation were each reported by 68.4% of respondents. Translation or grammar correction was reported by 57.9%, while 35.1% used AI for data analysis or statistics. These results show that AI tools were embedded in both technical and non-technical academic activities.

The interview data reinforced this pattern of routinisation. Students described AI as part of their normal study workflow rather than as a special or exceptional tool. One participant stated, "It's part of my workflow... I use it every day" (Int 2), while another described it as habitual: "It's like a habit... like brushing" (Int 5). This routine use is important because responsible AI behaviour is shaped through repeated small decisions, such as whether to verify an answer, disclose AI assistance, avoid sharing data or choose a lower-impact tool.

4.2 Awareness of Green AI and Responsible AI

The survey results indicate that students had a generally positive level of awareness of Green AI. In response to the Green AI awareness item, 43.9% of respondents selected "strongly agree" and 31.6% selected "agree," while 17.5% selected a neutral response. This suggests that approximately three-quarters of the respondents expressed agreement that they were aware of Green AI concepts. However, the presence of a notable neutral group also indicates that this awareness was not uniform across the sample.

Knowledge of AI energy consumption was more mixed. In the knowledge-based item on AI training and usage, 63.2% of respondents selected the correct answer: that AI training consumes energy during model development, while AI usage consumes energy continuously. However, 33.3% selected the option that training and usage consume the same amount of energy, 22.8% indicated that usage always consumes more energy than training, and 1.8% selected the option that energy used during training does not affect environmental impact. Because respondents could select more than one option, the percentages exceed 100%. These results suggest partial environmental literacy: many students understood the basic distinction between training and usage-related energy consumption, but misconceptions remained.

Students also showed awareness that AI use has environmental consequences. In relation to the perceived environmental impact of AI use, the largest

response categories were “agree” and “strongly agree.” However, perceived personal responsibility was more evenly distributed, with agreement, disagreement and neutral responses all present. This indicates a distinction between recognising that AI has environmental impact and feeling personally responsible for reducing that impact.

Interview findings also reflected this uneven awareness. Some students clearly recognised that AI systems require substantial resources. One participant stated, “I know it’s bad... data centres use a lot of water and energy” (Int 6). Another explained, “It definitely consumes power... I don’t know exact numbers, but I’m aware” (Int 2). These comments suggest that students often had a general awareness of environmental impact, but not necessarily detailed or systematic knowledge.

Awareness of Responsible AI was stronger. In the Responsible AI awareness item, 61.4% of respondents selected “strongly agree” and 26.3% “agree,” while 8.8% selected a neutral response. This indicates that nearly nine in ten respondents expressed agreement that they were aware of Responsible AI concepts. Awareness of data privacy risks was similarly high: 43.9% agreed and 42.1% strongly agreed that they understood the privacy risks of entering sensitive information into commercial AI tools.

Knowledge-based and scenario-based items further confirmed that students could identify responsible options in structured situations. In a Responsible AI principles knowledge check, 86.0% of respondents correctly identified human oversight and verification as the key issue. In a privacy-related knowledge item, 89.5% correctly identified that a dataset containing participant names and demographic information should not be uploaded to a public AI chatbot without institutional approval. These findings suggest that students are able to recognise responsible AI choices when risks are made explicit. Taken together, the survey and interview evidence converged in showing relatively high conceptual recognition of Responsible AI risks. However, the interviews qualified the survey findings by indicating that awareness was often general rather than detailed, especially in relation to environmental impacts.

4.3 Responsible Practice and the Awareness-Practice Gap

The central finding is that awareness did not consistently translate into responsible academic behaviour. Students were often able to identify responsible choices in abstract or scenario-based questions, but their reported everyday practices were less consistent. This awareness-practice gap was visible in three main areas: verification of AI outputs, privacy protection and environmentally restrained use.

In the academic integrity scenario, 42.1% of respondents selected rewriting the AI-generated content in their own words, while 33.3% indicated that

they would add a brief note acknowledging AI use. Smaller groups indicated that they would submit the work without mentioning AI use or ask the lecturer, while 10.5% selected “not sure.” These results suggest that many students could identify responsible options, but uncertainty and variation remained around disclosure and academic transparency.

Interview data showed that disclosure was often conditional rather than routine. One participant stated, “If they don’t ask, I wouldn’t mention it” (Int 6). Another explained, “If I wrote it and improved it with AI, it’s still my work” (Int 5). These comments illustrate that students may interpret AI assistance differently depending on whether they see it as support, collaboration or authorship. Disclosure therefore appears to depend strongly on perceived expectations rather than on a stable norm. The two forms of evidence were complementary: the survey demonstrated variation in chosen disclosure responses, while the interviews explained that disclosure decisions depended on perceived authorship and lecturer expectations.

Verification practices also showed an awareness-practice gap. Students recognised that AI outputs can be inaccurate, fabricated or biased. One participant noted, “It makes up sources... they don’t even exist” (Int 6), while another observed, “Sometimes it just accepts my ideas even if they are wrong” (Int 2). Despite this awareness, verification was not always systematic. Students reported checking more carefully for important assignments, but using AI outputs with less scrutiny for lower-stakes work. As one participant explained, “If it’s important, I check everything... if not, I just wing it” (Int 5). Another stated, “For less relevant work, I don’t always check” (Int 6). The survey and interview data together suggest that verification was not absent but conditional. Students reported applying stronger checks to high-stakes work and lighter checks to routine or lower-stakes tasks.

Privacy behaviour followed a similar pattern. In a privacy scenario, 64.9% of respondents selected removing identifiers before using an AI tool, while 28.1% indicated that they would avoid uploading the data altogether. This shows that students could identify safer privacy practices when the issue was clearly presented. However, routine privacy-related practices were weaker. Survey findings on checking privacy policies before using new AI tools showed that a large proportion of respondents selected “never” or “rarely,” with only a small group indicating that they did this often or always. The datasets showed a discrepancy between students’ ability to identify a safe response in a structured scenario and their reported reluctance to examine privacy conditions during routine tool selection.

The interviews help explain this gap. Students were aware that AI tools may store or reuse inputs, but this awareness did not always prevent data sharing. One participant reflected, “It came to my mind... but I just ignored it” (Int 5). Another stated more directly,

“You’re basically giving away your data” (Interview 6). These comments suggest that privacy awareness exists, but students may treat privacy checks as additional effort rather than as a normal part of AI-supported academic work.

Environmental responsibility was also inconsistent. In the environmentally responsible use scenario, 80.7% of respondents selected the use of a built-in spell-checker or lightweight grammar tool for a simple spelling correction task. This indicates that students could identify a lower-impact option when the choice was made explicit. However, environmental considerations were much less visible in routine tool selection. Survey results on whether students considered privacy policies or energy efficiency when choosing AI tools showed that “rarely” and “never” were common responses, while “often” and “always” were selected by only a limited number of respondents. Interview data supported this interpretation. Students recognised environmental impact but did not usually treat it as a priority in everyday use. One participant stated, “It’s in the back of my mind, but I don’t think about it daily” (Int 2). Another commented, “When I need information, I don’t think about the environment” (Int 5). A third summarised the tension clearly: “I know it’s bad... but it’s convenient” (Int 6). These findings indicate that environmental responsibility is particularly vulnerable to being displaced by speed, convenience and task completion. The scenario result therefore indicates recognition of a comparatively lower-impact option, whereas the interviews suggest that environmental considerations rarely functioned as a spontaneous criterion in everyday use.

4.4 Contextual Influences on Responsible AI Behaviour

The findings show that the awareness-practice gap was not simply caused by lack of knowledge. Instead, students’ behaviour was shaped by contextual factors, including time pressure, perceived task importance, institutional ambiguity and tool design.

Time pressure was one of the strongest barriers. In the survey item on whether deadlines make responsible AI use more difficult, agreement-based responses were common, indicating that many students perceived time pressure as a constraint on responsible behaviour. Interview participants also described deadlines as increasing dependence on AI. One student stated, “Closer to deadlines, it’s more tempting to use AI” (Int 6). Another explained that responsible checking takes additional effort: “It costs too much time to check everything yourself” (Int 6). These findings suggest that responsible AI practices are often perceived as desirable but costly.

Task importance also influenced behaviour. Students reported that they were more likely to verify outputs and act carefully in high-stakes situations, while shortcuts were more common in routine or lower-stakes tasks. This indicates that responsibility was

conditional rather than habitual. Students did not necessarily apply a fixed responsible AI routine across all academic work; instead, they adjusted their behaviour according to perceived risk and importance.

Institutional and lecturer guidance also shaped behaviour. Survey responses on institutional clarity showed a high proportion of neutral responses, suggesting uncertainty about whether university rules on student AI use were clear. Responses on lecturer guidance showed a similar pattern, with many students selecting neutral or mixed responses. Interview findings confirmed this ambiguity. One participant stated, “No course directly covered this” (Int 6), while another said, “It’s something I just kind of knew” (Int 2). These comments suggest that students often rely on personal judgement rather than structured institutional guidance.

Tool design was another important influence. Survey results on whether AI tools make it easy to choose options that protect privacy or use less energy showed limited confidence in current tool support. Many respondents selected neutral, disagreement or “not sure” responses. Interview participants suggested that AI tools rarely make environmental or privacy consequences visible at the point of use. One participant argued, “If it showed I wasted 100 liters of water, I might think twice” (Int 6). Another stated, “Systems don’t make you conscious about it” (Int 2). These findings suggest that responsible behaviour is affected not only by student awareness but also by the design of AI interfaces.

The comparative analysis between bachelor’s and master’s students showed only small differences. Bachelor’s students reported slightly higher mean scores for Green AI awareness ($M = 4.36$) than master’s students ($M = 4.12$), and also higher awareness of AI energy consumption ($M = 4.46$ compared with $M = 4.07$). Responsible AI awareness was high in both groups, with only a minimal difference between bachelor’s students ($M = 4.59$) and master’s students ($M = 4.50$). Privacy behaviour was also similar, although bachelor’s students reported slightly higher data-avoidance behaviour ($M = 4.10$) than master’s students ($M = 3.96$). In contrast, master’s students reported stronger privacy policy checking behaviour ($M = 2.68$) than bachelor’s students ($M = 1.62$). Thus, academic level did not strongly differentiate awareness or practice. Because these comparisons were descriptive and based on small subgroups, the observed differences should be interpreted as exploratory patterns rather than statistically established differences between bachelor’s and master’s students.

Within this sample, students in the two AI-related programmes were frequent users of AI tools and generally reported recognising ethical, privacy and environmental concerns. However, responsible behaviour remains inconsistent. Students are often able to identify responsible choices in structured scenarios, but everyday practice is shaped by convenience, time

pressure, unclear expectations and limited tool-level support. The central pattern observed in this exploratory sample was not a simple absence of awareness, but a gap between knowing what responsible AI use requires and consistently doing it in routine academic work.

5 Discussion

The findings challenge a simple information-deficit model of AI responsibility. Students' inconsistent responsible AI practices cannot be explained only by a lack of awareness. On the contrary, the results suggest that many students recognise key ethical, privacy and environmental concerns associated with AI use. However, this awareness competes with other factors in everyday academic work, including efficiency, convenience, time pressure, perceived task importance, confidence in AI outputs, institutional ambiguity and the design of AI tools. This supports behavioural perspectives that view action as shaped not only by knowledge and attitudes, but also by norms, habits and perceived behavioural control.

The awareness-practice gap should therefore be interpreted carefully. Inconsistent verification of AI outputs, for example, does not necessarily mean that students are unaware of AI risks or indifferent to academic quality. It may also reflect pragmatic judgement. Students may verify outputs more carefully when the task is high stakes, factually sensitive, unfamiliar or formally assessed, while applying lighter checks to low-stakes tasks such as brainstorming, language refinement or preliminary explanation. From this perspective, selective verification may be a form of risk-based decision-making rather than simple irresponsibility. Nevertheless, such pragmatism becomes problematic when students overestimate their ability to judge when verification is necessary, or when fluent AI outputs create a false sense of reliability. The issue is therefore not only whether students verify, but whether they possess sufficiently robust criteria for deciding when and how verification should occur.

A similar point applies to privacy behaviour. The finding that students do not routinely read privacy policies should not be interpreted only as a failure of individual responsibility. Privacy policies are often lengthy, legalistic and difficult to compare across tools. For students working under time constraints, reading such documents may appear unrealistic or disproportionate to the immediate academic task. This suggests that the gap between privacy awareness and privacy practice is partly structural. Students may understand that sensitive data should be protected, but lack accessible, actionable information about what different tools do with uploaded content. In this sense, privacy responsibility cannot be placed entirely on students; it also depends on institutional guidance, tool transparency and the availability of safer alternatives.

The Green AI findings raise a related issue. Environmental responsibility is particularly difficult to operationalise because the environmental impact of AI use is remote, distributed and invisible. Unlike academic integrity violations, which are directly connected to assessment rules, the energy and water costs of AI use are not visible in the interface or in the assignment process. Students may therefore understand that AI systems have environmental consequences while still failing to consider these consequences during ordinary tool selection. This does not necessarily indicate lack of concern. Rather, it reflects weak feedback mechanisms and limited perceived control. When students cannot see, estimate or compare environmental costs, sustainability becomes an abstract concern rather than a practical decision criterion.

These findings suggest that Responsible AI and Green AI should be understood as sociotechnical practices rather than only individual choices. Students make decisions within environments shaped by course expectations, institutional policies, assessment pressure and commercial platform design. When university guidance is ambiguous, students must infer what counts as acceptable assistance, collaboration, authorship or misconduct. When AI interfaces prioritise speed and ease of use, they encourage immediate task completion rather than reflection. When tools do not display privacy risks, uncertainty levels or environmental signals, responsibility is delegated to users at the moment when convenience is most salient. The result is a mismatch between the expectations placed on students and the supports available to help them act responsibly.

This distinction between individual, institutional and platform responsibility is important. Students remain responsible for how they use AI in academic work, particularly in relation to verification, authorship, disclosure and data protection. However, their responsibility should not be treated as isolated from the systems in which AI use occurs. Universities have responsibility for establishing clear norms, providing practical guidance and designing assessment conditions that encourage transparent and critical AI use. Platform providers also have responsibility for making risks visible, offering privacy-preserving defaults and supporting users in making proportionate choices. Responsible AI behaviour is therefore not produced by awareness alone, but by the alignment of student competence, institutional expectations and tool design.

The comparison between bachelor's and master's students did not reveal a uniform academic-level pattern. Awareness scores were broadly similar, while the clearest descriptive difference concerned privacy-policy checking. This suggests that progression from bachelor's to master's study may not, by itself, produce consistently more responsible AI routines. However, the subgroup sizes were too small to support a firm conclusion, and programme-specific curriculum,

research experience or assessment practices may be more relevant than study level alone. Future studies should examine these factors using larger stratified samples.

For Responsible AI, the practical challenge is operationalisation. Principles such as transparency, accountability, privacy and human oversight are widely accepted, but they must be translated into repeatable academic routines. A responsible student workflow might include defining the task before prompting, avoiding the upload of sensitive data, checking factual claims against reliable sources, revising outputs in one's own words, documenting the role of AI assistance and confirming whether AI use is permitted for the assignment. However, these routines should not simply be assumed. They need to be taught, practised, assessed and reinforced across courses. Otherwise, responsible AI remains an abstract ideal rather than a stable academic habit.

For Green AI, students need practical heuristics rather than only general environmental warnings. Since students cannot realistically calculate the precise environmental cost of each interaction, guidance should focus on proportionate and actionable choices. Examples include using built-in spellcheckers or lightweight grammar tools for simple tasks, avoiding unnecessary regeneration of responses, batching related prompts, selecting smaller or institutionally recommended tools where appropriate, and reserving large generative systems for tasks where they provide meaningful value. Such heuristics would make sustainability more usable in everyday academic decision-making.

The results also imply that universities should avoid treating AI policy primarily as a disciplinary instrument. Strict prohibition may drive AI use underground and increase non-disclosure, while vague permissiveness may leave students uncertain about acceptable practice. A more productive approach is to distinguish between acceptable, conditional and unacceptable uses of AI. For example, brainstorming, language refinement or debugging may be permitted under specified conditions, while undisclosed generation of substantive assessed content may be prohibited. Clear disclosure templates and course-level examples can reduce uncertainty and normalise transparency.

This study suggests that AI literacy should be extended to include both responsibility literacy and sustainability literacy. For students in AI-related programmes, this is especially important because their habits as users may shape their future assumptions as developers, analysts or decision-makers. If AI is experienced primarily as a tool for speed and convenience, students may reproduce systems that externalise ethical and environmental costs. Conversely, if verification, transparency, privacy protection and sustainability are embedded into academic routines, responsible AI can become part of students' professional identity. The central implication

is therefore that responsible AI use should not be framed merely as individual compliance, but as a shared academic and sociotechnical competence supported by education, policy and design.

6 Limitations and Future Research

The study is exploratory and context-specific. The survey sample of 57 students and six interviews provide rich insight into patterns at one university, but they do not support broad statistical generalisation. The study relies on retrospective self-reports rather than direct observation, system logs or diary records. This introduces social desirability, recall and interpretation biases. Participants may overstate practices perceived as responsible, such as verification or privacy protection, and may differ in how they understand response categories such as 'often' or 'rarely.' Accordingly, the study cannot determine how frequently these behaviours occurred in actual AI interactions or establish the exact magnitude of the awareness-practice gap. The findings should be interpreted as evidence of perceived and reported practice.

Future research should use larger and more diverse samples, compare disciplines, and examine how policy changes affect behaviour over time. Observational or diary-based methods could capture actual AI-use routines more accurately than retrospective self-reporting. Experimental work could test whether interface prompts, disclosure templates, energy feedback or privacy warnings change student behaviour. Longitudinal research is also needed to understand whether responsible AI habits developed during university persist into professional practice.

7 Conclusion

This study examined how students in AI-related programmes understand and practise Green AI and Responsible AI in their academic use of AI tools. Based on survey and interview data, the findings suggest that participants in these two programmes were frequent AI users and generally recognised ethical, privacy and environmental concerns. However, this awareness does not consistently translate into everyday academic practice, particularly in relation to verification, disclosure, privacy protection and environmentally restrained tool use.

The principal pattern identified in this exploratory, self-reported dataset is a possible awareness-practice gap: students often know what responsible and sustainable AI use requires, but their reported behaviour appears to be associated with time pressure, convenience, task importance, unclear institutional expectations and limited tool-level support. This suggests that responsible AI use cannot be explained

only through individual knowledge or intention; it must also be understood as a sociotechnical practice shaped by educational and platform environments.

The paper offers a context-specific contribution to higher education AI research by connecting Responsible AI and Green AI through the lens of routine academic practice, and by showing that AI literacy should include situated behavioural competence rather than only conceptual awareness. Responsible AI use in higher education depends on making responsible behaviour ordinary, visible and institutionally supported. These implications should be interpreted as propositions for further investigation rather than as general conclusions about university students. Future efforts should therefore focus on embedding verification, transparency, privacy protection and sustainability into teaching, assessment, institutional guidance and AI tool design.

References

- Ajzen, I. (1991). The Theory of Planned Behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179-211. doi: 10.1016/0749-5978(91)90020-T
- Almassaad, A., Alajlan, H., & Alebaikan, R. (2024). Student Perceptions of Generative Artificial Intelligence: Investigating Utilization, Benefits, and Challenges in Higher Education. *Systems*, 12(10), 385. doi: 10.3390/systems12100385
- Athanassopoulos, S., Tzavara, A., Aravantinos, S., Lavidas, K., Komis, V., & Papadakis, S. (2026). Teacher Education Students' Practices, Benefits, and Challenges in the Use of Generative AI Tools in Higher Education. *Education Sciences*, 16(2), 228. doi:10.3390/educsci16020228
- Blake, J. (1999). Overcoming the 'Value-Action Gap' in Environmental Policy: Tensions Between National Policy and Local Experience. *Local Environment*, 4(3), 257-278. doi: 10.1080/13549839908725599
- Carrington, M. J., Neville, B. A., & Whitwell, G. J. (2014). Lost in Translation: Exploring the Ethical Consumer Intention-Behavior Gap. *Journal of Business Research*, 67(1), 2759-2767. doi: 10.1016/j.jbusres.2012.09.022
- Chan, C. K. Y., & Hu, W. (2023). Students' Voices on Generative AI: Perceptions, Benefits, and Challenges in Higher Education. *International Journal of Educational Technology in Higher Education*, 20(1), 43. doi: 10.1186/s41239-023-00411-8
- Dalvi-Esfahani, M., Alaedini, Z., Nilashi, M., Samad, S., Asadi, S., & Mohammadi, M. (2020). Students' Green Information Technology Behavior: Beliefs and Personality Traits. *Journal of Cleaner Production*, 257, 120406. doi: 10.1016/j.jclepro.2020.120406
- Deng, R., Jiang, M., Yu, X., Lu, Y., & Liu, S. (2025). Does ChatGPT Enhance Student Learning? A Systematic Review and Meta-analysis of Experimental Studies. *Computers & Education*, 227, 105224. doi:10.1016/j.compedu.2024.105224
- Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., ... Wright, R. (2023). Opinion Paper: "So What if ChatGPT Wrote it?" Multidisciplinary Perspectives on Opportunities, Challenges and Implications of Generative Conversational AI for Research, Practice and Policy. *International Journal of Information Management*, 71, 102642. doi: 10.1016/j.ijinfomgt.2023.102642
- European Commission. Directorate General for Education, Youth, Sport and Culture. (2022). *Ethical guidelines on the use of artificial intelligence (AI) and data in teaching and learning for educators*. Publications Office. <https://data.europa.eu/doi/10.2766/153756>
- European Commission. (2024). *Living guidelines on the responsible use of generative AI in research*. Directorate-General for Research and Innovation. https://research-and-innovation.ec.europa.eu/document/download/2b6cf7e5-36ac-41cb-aab5-0d32050143dc_en
- European Commission. Directorate General for Communications Networks, Content and Technology & High Level Expert Group on Artificial Intelligence. (2019). *Ethics guidelines for trustworthy AI*. Publications Office. <https://data.europa.eu/doi/10.2759/346720>
- Floridi, L., & Cows, J. (2019). A Unified Framework of Five Principles for AI in Society. *Machine learning and the city: Applications in Architecture and Urban Design*, 535-545. doi: 10.1162/99608f92.8cd550d1
- Floridi, L., Cows, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People - An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations. *Minds and Machines*, 28(4), 689-707. doi: 10.1007/s11023-018-9482-5
- Fošner, A. (2024). University Students' Attitudes and Perceptions Towards AI Tools: Implications for Sustainable Educational Practices. *Sustainability*, 16(19), 8668. doi: 10.3390/su16198668

- Francis, N.J., Jones, S., & Smith, D.P. (2025). Generative AI in Higher Education: Balancing Innovation and Integrity. *British Journal of Biomedical Science*, 81:14048. doi:10.3389/bjbs.2024.14048
- Holmes, W., & Miao, F. (2023). *Guidance for Generative AI In Education and Research*. UNESCO Publishing.
- Jobin, A., Ienca, M., & Vayena, E. (2019). The Global Landscape of AI Ethics Guidelines. *Nature Machine Intelligence*, 1(9), 389-399. doi: 10.1038/s42256-019-0088-2
- Kasneci, E., Seßler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., ... & Kasneci, G. (2023). ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Education. *Learning and Individual Differences*, 103, 102274. doi: 10.1016/j.lindif.2023.102274
- 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. doi: 10.1080/13504620220145401
- Mittelstadt, B. (2019). Principles Alone Cannot Guarantee Ethical AI. *Nature Machine Intelligence*, 1(11), 501-507. doi: 10.1038/s42256-019-0114-4
- Morley, J., Kinsey, L., Elhalal, A., Garcia, F., Ziosi, M., & Floridi, L. (2023). Operationalising AI Ethics: Barriers, Enablers and Next Steps. *AI & Society*, 38(1), 411-423. doi: 10.1007/s00146-021-01308-8
- Nikolic, S., Wentworth, I., Sheridan, L., Moss, S., Duursma, E., Jones, R. A., ... & Middleton, R. (2024). A Systematic Literature Review of Attitudes, Intentions and Behaviours of Teaching Academics Pertaining to AI and Generative AI (GenAI) in Higher Education: An Analysis of GenAI Adoption Using the UTAUT Framework. *Australasian Journal of Educational Technology*, 40(6), 56-75. doi: 10.14742/ajet.9643
- OECD (2026). *OECD Due Diligence Guidance for Responsible AI*, OECD Publishing, Paris, doi: 10.1787/41671712-en.
- Raman, R., Pattnaik, D., Lathabai, H. H., Kumar, C., Govindan, K., & Nedungadi, P. (2024). Green and Sustainable AI Research: An Integrated Thematic and Topic Modeling Analysis. *Journal of Big Data*, 11(1), 1-28. doi:10.1186/s40537-024-00920-x
- Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54-63. doi:10.1145/3381831
- Stracke, C. M., Nahar, N., Punzo, V., Massaro, S., Pappa, D., Di Grassi, A., ... & Burgos, D. (2026). Artificial Intelligence Policies for Higher Education: Manifesto for Critical Considerations and a Roadmap. *MAP Education and Humanities*, 6, 61-73. doi: 10.53880/2744-2373.2025.6.61
- Stern, P. C. (2000). New Environmental Theories: Toward a Coherent Theory of Environmentally Significant Behaviour. *Journal of Social Issues*, 56(3), 407-424. doi: 10.1111/0022-4537.00175
- Tabbakh, A., Al Amin, L., Islam, M., Mahmud, G. I., Chowdhury, I. K., & Mukta, M. S. H. (2024). Towards Sustainable AI: A Comprehensive Framework for Green AI. *Discover Sustainability*, 5(1), 408. doi:10.1007/s43621-024-00641-4
- Verdecchia, R., Sallou, J., & Cruz, L. (2023). A Systematic Review of Green AI. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 13(4), e1507. doi: 10.1002/widm.1507
- Verplanken, B., & Wood, W. (2006). Interventions to break and create consumer habits. *Journal of Public Policy & Marketing*, 25(1), 90-103. doi: 10.1509/jppm.25.1.90
- Wang, H., Dang, A., Wu, Z., & Mac, S. (2024). Generative AI in Higher Education: Seeing ChatGPT Through Universities' Policies, Resources, and Guidelines. *Computers and Education: Artificial Intelligence*, 7, 100326. doi: 10.1016/j.caeai.2024.100326
- Xiao, L., Pyng, H. S., Ayub, A. F. M., Zhu, Z., Gao, J., & Qing, Z. (2025). University Students' Usage of Generative Artificial Intelligence for Sustainability: A Cross-Sectional Survey from China. *Sustainability*, 17(8), 3541. doi:10.3390/su17083541