Journal of ExoTechnology and Education
Vol. 2, No. 1, 2026

Online ISSN: 3079-7950

Prompting Change: Exploring Undergraduate Perceptions after AI Literacy Training

Author
Dr. Terence Govender

Affiliation
Indiana University Columbus
Center for Teaching and Learning

ORCID link
https://orcid.org/0000-0001-9145-040X

Correspondence to
Dr. Terence Govender

E-mail
tgovende@iu.edu

Dates
Received: 12 January 2026
Revised: 18 April 2026
Accepted: 19 April 2026
Published: 1 May 2026

How to cite
Govender, T. (2026) 'Prompting change: Exploring undergraduate perceptions after AI literacy training,' *jeXed* 2(1). 23 pages.
https://doi.org/10.36615/9953hz86

Article includes
☒ Peer review
☐ Supplementary material

Data availability
☒ Open data set
☒ All data included
☐ On request from author/s
☐ Not available
☐ Not applicable

Editor/s
Dr. Irene Lubbe

Funding
No funding was received for this article.

Conflict of interest
None.

Society with whom we Collaborate
International Conference on Information Communication Technologies in Education

Website: https://icicte.eu/

© 2026. The author/s. Published under a Creative Commons Attribution License.

Abstract

As generative artificial intelligence (GenAI) tools become increasingly integrated into daily life, society has been steadily adapting to changes in both work and education, often where guidelines and social contracts have yet to be established. Questions surrounding ethical usage are still being debated, and increased literacy in GenAI is necessary to enable voices, especially from younger generations, in these discussions. This paper presents an exploratory pilot case study on the effects of advancing AI literacy among a small group of university undergraduate students ($N = 6$) through a personalised training intervention. Drawing on pre- and post-test survey data and transcripts from training on Adobe Firefly, Google NotebookLM, and Google Gemini, the study examines changes in students' knowledge and ethical considerations. While participants initially expressed cautious or critical views (e.g., fearing an overreliance on GenAI or academic dishonesty), quantitative results after training indicate greater familiarity in using GenAI as a learning tool rather than a shortcut, with one student in the post-training survey stating how they now believe their learning 'will be enhanced far beyond what [they] were doing before.' Additionally, qualitative analysis reveals a persistent discrepancy between the perceived and actual understanding of GenAI and academic policies, further supporting the need for these discussions. Therefore, this work positions AI literacy as a critical component of higher education onboarding.

Interdisciplinary Implications

This article reports on a mixed-methods pilot intervention examining how undergraduate students' perceptions of generative AI shift after structured AI literacy training, with implications for curriculum design, the ethical integration of emerging technologies into academic practice, and faculty professional development. It draws on and contributes to higher education pedagogy, educational psychology (specifically constructivism and expectancy-value theory), computer and information science (GenAI and retrieval-augmented generation tools), and academic ethics.

Keywords
AI literacy, generative AI, academic integrity and ethics, higher education

1. Introduction

The increased availability of generative artificial intelligence (GenAI) tools since late 2022 has presented educators with a complex challenge. On the one hand, these GenAI tools, such as ChatGPT (https://chatgpt.com/), offer new ways to support student learning and promote equity. However, they also raise concerns about academic integrity and the possible reduction of critical thinking skills. Educators initially responded to this disruption in a variety of ways, ranging from banning GenAI to adopting it. However, as time went by, a consensus began to emerge that simply policing GenAI use is insufficient (Leslie and Perini, 2024). Instead, there is a growing recognition that students need to be AI-literate in order to use these technologies effectively and ethically (Yang *et al*., 2025).

AI literacy has been defined by prior research as a 'set of competencies that enables individuals to critically evaluate AI technologies; communicate and collaborate effectively with AI; and use AI as a tool online, at home, and in the workplace' (Long and Magerko, 2020, p. 2). In higher education (HE), these competencies include understanding how GenAI systems such as large language models (LLMs) work, knowing how to verify AI-generated outputs, and being aware of the ethical considerations and limitations of GenAI. In educational settings, AI literacy entails students learning not only how to use tools like chatbots or image generators, but also when and why to use them to enhance their learning journey. Recent studies show that prompt crafting and iterative dialogue with GenAI are valuable skills (also termed 'prompt engineering') because many GenAI systems respond differently depending on how tasks, roles, constraints, and examples are specified (Federiakin *et al*., 2024; Qian, 2025). For example, Knoth *et al*. (2024) found that higher AI literacy in students is associated with more effective prompting strategies and higher-quality outputs from GenAI, leading them to recommend integrating AI literacy into curricula.

Although the ecosystem in HE has changed considerably due to GenAI, new tools and features are continually released, which have created an additional challenge for those trying to stay up to date. Therefore, there are impactful GenAI tools and updates for education that may go unnoticed due to the influx of information. For example, students can benefit from utilising Google's NotebookLM (https://notebooklm.google/), which allows users to upload their own materials or notes and interact with a personalised AI assistant that can summarise readings, answer questions about the uploaded sources, and even generate podcast-style audio explanations of the content (Rebelo, 2025). Further, GenAI is not limited to text and audio. For example, tools like Adobe Firefly (https://firefly.adobe.com/), introduced in 2023, enable users to create images (and other media, including video) from text prompts. Adobe has integrated Firefly into its products with built-in ethical guardrails, such as restrictions on access to copyright training data, so that students can more safely use AI-generated images in projects (Adobe, 2024). Recently, Google's Gemini team has taken a step further with their Veo 3 model, which can create videos with native audio (Google, n.d.). These are just two companies from many that are pushing the boundaries of GenAI technology.

Unfortunately, GenAI also exacerbates the spread of misinformation and 'AI slop,' a term used to describe low-quality, superficial, or factually incorrect AI-generated outputs. Since these tools are trained on extensive datasets scraped from the internet, they inherently propagate inaccuracies and biases found in their training data. Additionally, users have been abusing these tools to create realistic but fake media. Students who are unfamiliar with evaluating information critically may inadvertently trust these GenAI outputs. A related concern is the 'liar's dividend,' a phenomenon where the mere existence of realistic AI-generated content creates doubt about the authenticity of genuine information (Schiff, Schiff and Bueno, 2025). Therefore, developing AI literacy among students becomes essential not only for identifying misinformation but also for strengthening their ability to navigate a media landscape where scepticism can be weaponised to challenge legitimate sources of information.

Employer expectations are also driving the need for AI literacy among graduates. Surveys of industry indicate that AI-related skills are now highly valued in the job market. A 2024 international study by Microsoft and LinkedIn (called the Work Trend Index) reports that over two-thirds of business leaders would not hire a candidate who lacks AI literacy, and nearly three-quarters said they would prefer a less-experienced applicant with AI literacy over a more-experienced one without it (Microsoft and LinkedIn, 2024). This trend is reflected in the rapid growth of job ads seeking AI-literate candidates, as well as wage premiums for those skills.

HE institutions are therefore increasingly expected to integrate AI literacy into their teaching. Within this context, this study addresses the key questions:

1. What are students' perspectives on and knowledge of GenAI prior to AI literacy training?
2. How does a targeted educational intervention change students' AI literacy and their attitudes toward using GenAI in learning?
3. Which design features of the training do students identify as most helpful and through what mechanisms?

The case study involved a workshop-style, one-on-one training session for students. By analysing qualitative data (participant transcripts) before, during, and after the training, as well as quantitative data (pre- and post-tests), this paper contributes a further understanding of the transformation potential that AI literacy has on HE. The aim is to help students learn more effectively and address institutional priorities related to academic integrity and graduate preparedness for the workplace in an AI-ubiquitous society.

2. Literature Review

The interaction between GenAI and HE has been explored from multiple angles in recent research, reflecting the fast-moving developments since 2022. Three interrelated topics are found in the literature: First, concerns around academic integrity and critical thinking in the age of GenAI; second, the potential benefits of GenAI for enhancing learning; and third, the imperative of AI literacy education as a response to the first two topics. This review synthesises these topics to provide context for the study.

2.1 Academic Integrity and Critical Thinking Concerns

Easily accessible GenAI tools have triggered debates over academic honesty and the nature of student work. Early commentary warned of ‘AI-giarism,’ wherein students could present AI-generated essays as their own, thus cheating plagiarism detectors (Chan, 2023; Cotton, Cotton and Shipway, 2023). Educators evaluating AI-generated essays have noted that GenAI writing is generally technically correct in terms of language use, often featuring advanced vocabulary and complex sentences, but tends to exhibit inadequate depth of analysis, incomplete development of ideas, and a lack of personal voice or authentic insight (Alexander, Savvidou and Alexander, 2023), raising flags about a new form of minimal-effort academics. Students are also aware of this issue. For example, a survey of 277 US HE students reports that a majority students believed using GenAI to complete assignments unquestionably constitutes academic misconduct (Lund *et al*., 2025). However, many students viewed smaller AI-assisted tasks such as using GenAI for ideas or proofreading, as falling into a grey area.

In terms of minimal-effort thinking (also known as cognitive offloading) versus critical thinking, students understand that an overreliance on GenAI tools without exercising independent thought could impact their ability to learn. A survey of 737 students found concerns about the devaluation of university education (Blahopoulou and Ortiz-Bonnin, 2025). This is why many scholars argue that, rather than prohibition, the solution lies in teaching students how to use GenAI appropriately (Apata *et al*., 2025).

2.2 GenAI’s Potential Benefits for Learning

Alongside the concerns, research also documents significant advantages of GenAI when used constructively. A survey of 23,218 students from over 100 countries reports that they used GenAI to support their learning through summarisation, aiding in ideation, and research (Ravšelj *et al*., 2025). These tools can also adapt to the user’s queries, providing iterative explanations until the student grasps a concept, functioning as an on-demand tutor (Yarlagadda, 2025).

GenAI can help bridge accessibility and equity gaps by providing automatic captioning and translation, text simplification, and multimodal content generation (Gabriel, 2024). However, maximising these benefits requires that students know how to prompt and cross-check GenAI outputs properly. Without guidance, these highly impactful benefits can turn into pitfalls, e.g., if a student unquestioningly trusts a flawed AI-generated explanation (Abuzar, Mahmudulhassan and Muthoifin, 2025). Therefore, researchers emphasise using GenAI tools from a ‘productive use’ perspective, which means using GenAI during the process of finding the answer rather than as a final answer generator (Bai and Wang, 2025). In this way, GenAI has the potential to complement human learning.

2.3 The Imperative of AI Literacy Education

AI literacy in HE should extend beyond technical skills. Rapanta *et al*. (2025) posit that both students and educators should understand GenAI’s limitations (e.g., potential biases, the concept of hallucination in LLMs), ethical considerations (privacy, consent, intellectual honesty), and effective strategies for using GenAI tools (how to craft prompts or refine GenAI outputs). Furthermore, they also comment on the challenge of creating a universal framework, as numerous factors need to be accounted for, and promote the idea of 'situated literacies' (meaning context-specific applications of skills).

Since AI literacy should be updated and reinforced periodically, as GenAI tools themselves evolve their capabilities, time should be invested in workshops, curriculum integration, and resource development. Ithaka S+R's multi-institution study found that many instructors have started embedding basic GenAI skills into student activities, but also that faculty desires more top-down support for formalising AI literacy (Baytas and Ruediger, 2025). Notably, instructors in the report advocate for clear policies and teaching modules on acceptable GenAI use so that students across courses receive consistent messaging.

A survey by Blahopoulou and Ortiz-Bonnin has found that there is a 'need for universities to actively involve students in shaping policies on artificial intelligence while offering targeted training to promote its responsible and ethical use' (Blahopoulou and Ortiz-Bonnin, 2025, p. 19741). Furthermore, this active involvement allows students to have agency in shaping guidelines, which could translate into a responsibility for upholding them. Accordingly, this study is situated at the intersection of these discussions, aiming to demonstrate how an educational intervention can help to realise GenAI's benefits (as a supportive learning tool) while mitigating its risks (misuse and overreliance). This research adds to a small but growing body of empirical evidence on how AI literacy programmes can shift student attitudes and behaviours.

2.4 Conceptual Framework

This study is guided by a constructivist view of learning, in which students actively build understanding through interaction, reflection, and contextualised engagement rather than passively receiving information. In the context of GenAI, this perspective is especially relevant because a meaningful use of these tools depends not only on technical access but also on how students interpret outputs and question limitations, while connecting new information to their existing knowledge and academic practices. In this sense, AI literacy is not treated here as a purely technical competency, but as a situated literacy shaped by context, purpose, and ethical judgement. This aligns with recent arguments that AI literacy in HE should extend beyond knowing what GenAI tools are, toward understanding when, why, and under what conditions they should be used responsibly.

This framing informed the design of the personalised one-on-one training sessions used in the study. Rather than delivering a standardised lecture about AI tools, the intervention emphasised dialogue, demonstration, guided experimentation, and critical reflection. Participants were encouraged to test tools for authentic academic purposes, reflect on the quality and trustworthiness of outputs, and discuss the ethical boundaries of use in relation to their own courses and disciplinary needs. The intention was not simply to increase tool familiarity, but to support students in constructing a more nuanced understanding of GenAI as both useful and limited. This also reflects the paper's broader concern by helping students to move from regarding GenAI as either a shortcut or a threat toward understanding it as a tool whose value depends on informed and reflective use.

In addition to this pedagogical framing, the organisation of the questionnaire was informed by an expectancy-value perspective (EVT) (Wigfield and Eccles, 2000). The survey categories of familiarity and usage, expectancy and academic usefulness, value and interest, concerns and ethical considerations, and future intentions were selected to capture how students perceived the relevance, utility, and risks of GenAI before and after the intervention. These dimensions are useful because students are more likely to engage productively with a tool when they regard it as valuable, feel capable of using it, and understand its limitations. Accordingly, the study interprets shifts in participants' responses not simply as attitude change, but as indications of evolving meaning-making around GenAI in academic contexts. Together, constructivism and an expectancy-value lens provide the conceptual basis for both the intervention design and the interpretation of findings in this pilot study.

3. Methodology

This section details the experiment design, including the participants and data collection methods before, during, and after the AI literacy training.

3.1 Participants

Six undergraduate students were recruited to participate in an AI literacy training session through a process of selective convenience sampling ($N = 6$). These participants were approached through an existing university network and invited to voluntarily participate in a one-on-one AI literacy training session. 10 invitations were sent out, and six responded. No incentives were offered or given for participation in the study. Selection criteria required that participants be 1) currently enrolled undergraduate students at the university; 2) in their second or third year of study; and 3) willing to be recorded and to share survey data for research purposes. Selective sampling allowed participants to be drawn from various disciplines (two from business, one from English literature, one from criminal justice, one from psychology, and one exploratory) in order to capture a range of perspectives. It also allowed the study to include an equal mix of genders (three male and three female students). This was intended to capture a broader range of perspectives in this exploratory pilot study, not to achieve statistical representativeness. Notably, their prior exposure to GenAI varied. Initially, two of the participants have disclosed never using tools like ChatGPT before, while three have used GenAI tools minimally for ideation and experimented with them for school or personal use. Only one participant stated using GenAI tools to a moderate degree. This variety was intentional, as the aim was to record how the training would affect both novice and experienced users.

To encourage an honest discussion, participants were assured that the study was not an evaluation of their abilities, but rather an exploration of learning with GenAI, and how they foresee the effects of its use. To comply with ethical standards, ethical clearance was obtained for working with human participants. Each student was informed about the entire experiment prior to signing consent, which included recording their session and sharing survey responses for research purposes, with the understanding that any identifying details would be removed. The context around the impetus for the study was explained as a push to better understand how to incorporate GenAI into teaching and learning. Thus, this pilot training served both research and practical development purposes, potentially informing broader institutional initiatives.

3.2 Experiment and Data Collection

Each participant engaged in a one-on-one training session conducted via an online meeting platform and followed a consistent three-part structure: 1) A pre-test survey; 2) an AI literacy training session; and 3) a post-test survey and reflection. This format was chosen to allow personalised interaction and candid discussion, and students could keep their cameras off to increase the sense of anonymity. The questionnaire used in this study was developed based on insights and constructs from prior research (Chan and Hu, 2023; Chan and Zhou, 2023; Baidoo-Anu *et al*., 2024; Gasaymeh, Beirat and Abu Qbeita, 2024; cf. Table 1). These studies informed both the theoretical framework and practical question formulation, ensuring a rigorous exploration of students' perceptions, familiarity, concerns, and intended use of GenAI. The sessions lasted approximately 60 to 75 minutes.

Step 1: Pre-Test Survey

Before the personalised training began, each participant completed a short online questionnaire to gauge their baseline attitudes, knowledge, and usage of GenAI (taking around 10 minutes). This survey included Likert-scale items that were divided into categories (cf. Table 1) and an optional open-ended question, which asked, 'In one or two sentences, how do you currently feel about using GenAI tools like ChatGPT for learning or studying?'

Step 2: AI Literacy Training Session

A semi-structured AI literacy training session, which lasted around 40-50 minutes included conversation-starting topics and hands-on activities (10-15 minutes), allowing for flexibility to follow participants' questions or areas of interest. The initial guiding questions included:

- What is your current understanding of GenAI, and how have you used it before?
- How do you think this technology will affect critical thinking?
- What are your concerns, especially related to how it will affect your education and future job?
- Do you know the policy around unacceptable usage at the university?

Table 1
Questionnaire categories and questions

Topic	Possible Answer
Familiarity and Usage (F)	F1: I am familiar with the concept of GenAI tools (e.g., ChatGPT, Claude, and Gemini). F2: I have used GenAI tools before this training. F3: I understand how these tools generate responses. F4: I feel confident navigating a GenAI interface.
Expectancy and Academic Usefulness (E)	E1: I believe GenAI can support my learning. E2: GenAI tools can help me complete academic tasks more efficiently. E3: I believe using GenAI could improve my academic performance. E4: I feel GenAI can be a helpful study partner.
Value and Interest (V)	V1: I am curious to explore how GenAI can assist with my studies. V2: I am interested in learning more about how GenAI works. V3: I think it is important to learn how to use GenAI tools ethically.
Concerns and Ethical Considerations (C)	C1: I am concerned that using GenAI might reduce my original thinking. C2: I am unsure about whether using GenAI is considered academic dishonesty. C3: I worry that an overreliance on GenAI tools may hurt my learning in the long run.
Future Intentions (I)	I1: I would like to use GenAI tools more in my academic work. I2: I would recommend these tools to my peers. I3: I believe that GenAI will play a major role in the future of education.

Using these topics, explanations were pitched to each participant's level, from basic analogies for how GenAI generates outputs through to the mechanics of how these models are trained. This included discussing concepts such as bias in GenAI (e.g., how training data can skew outputs) and hallucination (GenAI producing false information). The origin of some models' training data was also discussed (for example, the mention that early models were trained on internet forums, which introduced certain biases). Although the intervention was personalised, it followed a common instructional backbone across all six sessions. Each session included the same broad sequence: 1) Completion of the pre-test survey; 2) discussion of participants' prior understanding and use of GenAI; 3) explanation of core concepts such as statistical generation, bias, hallucination, and data privacy; 4) demonstration of institutionally available tools, especially Google Gemini and NotebookLM; 5) guided hands-on experimentation with at least one tool, including iterative prompt refinement; and 6) a discussion of academic integrity and acceptable use before the completion of the post-test survey. Personalisation occurred primarily in the depth, pacing, examples, and specific tools emphasised, which were adapted to every participant's prior experience, disciplinary context, and questions. This design was intended to make the training responsive and pedagogically relevant, but it also means that some facilitator effects and session-level variability were inherent to the intervention.

After going through foundational knowledge, at least two of the tools that the participants had access to via their institution (which uses enterprise licences for data protection) were demonstrated (15-20 minutes). The list at the time of this research was Microsoft Copilot, Google Gemini, Google NotebookLM, and Adobe Firefly. One of the main learning objectives of the demonstration was to explain how to create a prompt and how it can be refined iteratively, therefore progressing from a vague query to a more specific one using a structured approach. For text output, a simple prompting framework called CORE (Context, Output, Role, Example) was used to illustrate how providing context and role to the GenAI can yield better results. This was related back to the foundational knowledge of the models using statistics, where narrowing the scope increases accuracy by adding weight to the correct selection of words that need to be output (Park and Choo, 2025). For instance, rather than prompting 'give me a campaign e-mail template,' the participants were told that they could adjust the prompt with more information, for example:

> *[ROLE]*
>
> *Act as an expert e-mail marketing strategist who specialises in crafting compelling copy for mobile app launches. Your tone should be enthusiastic, persuasive, and user-focused.*
>
> *[CONTEXT]*
>
> *I am the marketing manager for a new language-learning mobile app called LinguaLeap. We are about to launch next week. We have a list of 5,000 potential users who have pre-registered because they are interested in the app. The app's key feature is its use of AI-powered conversational practice with virtual tutors. The goal of this e-mail is to announce the official launch date and persuade these pre-registered users to be among the first to download and use the app.*

[OUTPUT]
Generate an e-mail with the following components:
Subject Line: A short, catchy subject line (under 10 words) that builds excitement and includes the app's name.
E-mail Body: A concise e-mail (around 150 words) that thanks them for pre-registering, announces the launch date, highlights the main benefit (AI conversation practice), and includes a clear call-to-action to download the app. The tone should be celebratory and encouraging.
Call-to-Action Button Text: Suggest text for the main download button.

[EXAMPLE]
Here is an example of the style and tone I am looking for. Please do not copy this directly, but use it as a guide for the voice.
Subject: Your Language Journey Begins!
Hi [User Name],
The moment we've all been waiting for is almost here! Get ready to start talking your way to fluency....

This was demonstrated live, showing how the response improved as the prompt became more detailed (cf. Figure 1 below).

For the NotebookLM demonstration, a sample set of notes and a PDF article were uploaded in advance. The participants were then shown how the tool could generate a succinct summary and even create a conversational audio podcast discussing the material. One participant was surprised when the GenAI produced a natural-sounding 20-minute dialogue explaining a complex text, noting, 'It's weird, but it's cool' to hear the AI-generated voices break down the content. This demonstration also aimed to expand students' awareness of GenAI capabilities beyond just text generation, highlighting summarisation, retrieval augmented generation (RAG), and multi-modal outputs. The concept of RAG was highly intriguing to the participants, as they could also see its value in reducing hallucinations, since the tool utilised the sources provided by the user rather than relying on the billions of data points in the training model and on the internet.

Figure 1
Sample interaction with Google Gemini during training

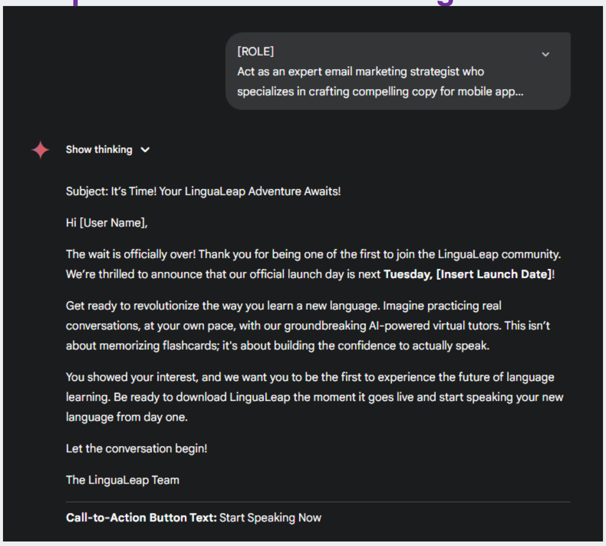

(Source: Author's screenshot of Google Gemini output)

Finally, for other media generations, emphasis was placed on understanding how to set the scene and the use of adjectives. However, for this study, none of the participants was interested in creating images or videos, where NotebookLM's RAG and podcast features generated the most attention.

Each student was then invited to experiment with a GenAI tool themselves (around 15 minutes). These activities were guided and adjusted to their field of study or interest. During this time, they were also coached on refining prompts. GenAI responses were discussed and critically evaluated, including fact-checking and identifying any gaps or biases that may have existed. Participants were encouraged to ask follow-up questions to the GenAI tool or request revisions (reinforcing the idea that GenAI output is an iterative dialogue).

Variability was primarily driven by student interest rather than facilitator deviation. In practice, the same core topics were addressed in all sessions, but participants differed in their prior AI exposure, concerns, and interests, which influenced the examples discussed, the amount of scaffolding required, and the tools that received the greatest attention. For example, some sessions focused more heavily on NotebookLM as a study aid, whereas others spent more time on prompting strategies, ethical concerns, or discipline-specific use cases. This variation reflects the exploratory and personalised nature of the intervention, but it also reduces strict replicability and should be taken into account when interpreting the findings.

In the latter part of the session, academic integrity and acceptable use were discussed. Participants were asked about current university policies, especially regarding what happens if a student is caught using GenAI tools in an unacceptable manner. At the time of this study, if a faculty member prohibits the use of GenAI tools and a student violates this prohibition, it is treated as cheating, with consequences such as redoing or failing the assignment. However, the punishment is not as severe as plagiarism, since plagiarism is defined as stealing someone else's work, whereas unacceptable GenAI usage is considered unauthorised assistance. This distinction was actually new information to some participants. One participant admitted, 'I thought it's plagiarism [similar to] taking stuff off the internet, but it's a little better that it's not [considered plagiarism], it's just cheating.' This open dialogue helped to clarify the parameters of acceptable use and encouraged students to articulate their experiences with GenAI tools more openly (even if they said they never used it at the start of the experiment).

Scenarios were used to elicit follow-up discussions, such as using GenAI to generate an essay that a student submits as their own would be cheating (unless explicitly permitted by the instructor), whereas using GenAI to brainstorm ideas or improve grammar might fall within acceptable use, depending on the course policies. Students were informed that many faculty were now including acceptable use statements in syllabi about GenAI. The importance of transparency was stressed from both ends. Students should ask their faculty how they are using GenAI in the course and also what would be acceptable for the students to use during their course.

Throughout the training, the tone was conversational and interactive. Students asked questions or voiced their thoughts, which were then addressed. For example, one participant asked about bias, which led to a discussion of how early image recognition models would return a positive skin cancer diagnosis if there was a ruler in the picture due to biased training data, and conversely, how over-correction can lead to strange outputs (like a GenAI bizarrely generating an African soldier image when asked to display an image of a Nazi, due to an overzealous attempt at diversity). These real examples underscore why one should not take GenAI outputs as a neutral truth. The participants were often prompted to reflect, such as, ‘Given what you just saw the AI did, how would you use this in studying for your next exam?’ and ‘Do you think using AI in that way would help you learn, or would it be doing too much of the work for you?’ This Socratic approach helped students to articulate their own stance on GenAI.

Step 3: Post-training Survey and Reflection
Before the end of the session, each participant was requested to complete the post-training survey, which consisted of the same set of Likert-scale questions as the pre-survey (to measure any shifts in agreement) and the same optional open-ended question about their feelings regarding the use of GenAI for learning. This step generally lasted around 10-15 minutes. Since the session was one-on-one, participants were allowed to verbally provide their open-ended answers to the survey, which were transcribed during the recording. Most participants chose to just discuss their responses aloud. Finally, there was an informal wrap-up conversation with questions such as ‘How do you feel about all this now?’ and ‘Do you think you will try using these tools going forward?’ This served both as a debrief and an opportunity for extra qualitative feedback. All sessions were recorded using a built-in transcription feature to capture the dialogue, and were later manually corrected for accuracy. These transcripts, along with the survey data, constituted the dataset.

Because the researcher facilitated both the one-on-one AI literacy sessions and conducted the qualitative analysis, reflexivity was important throughout the study. The researcher approached the project from a position that views AI literacy as educationally valuable, while also recognising concerns related to overreliance, bias, and ethical misuse. This dual role may have influenced both the conversational direction of the sessions and the interpretation of the participants’ responses. To reduce this risk, the analysis attended to both positive and critical participant perspectives, and the findings are presented as exploratory rather than generalisable.

3.3 Data Analysis

A mixed-methods analysis was employed. For the quantitative data, the survey datasets were entered into IBM SPSS Statistics Version 26 (IBM Corp., 2019). Each participant’s before-and-after answers were analysed to identify patterns in descriptive statistics. Since the study involved a small sample size ($N = 6$) and ordinal data derived from Likert-scale survey items, the Wilcoxon Signed Ranks Test was used to assess the intervention, with effect size (r) calculated using $r = \frac{Z}{\sqrt{N}}$, where Z is the z-score and N is the number of participants. The r was interpreted according to Cohen’s (2013) guidelines ($|r| = .1$ is small, $.3$ is medium, $.5$ is large). Given the small N, inferential statistics are reported primarily as descriptive indicators of within-participant change and should not be interpreted as population estimates.

For the qualitative data, a thematic analysis was conducted on the transcripts following the six-phase approach of Braun and Clarke (2006). In Phase 1 (familiarisation), the researcher read through all six transcripts, noting initial impressions and recurring ideas. After reading through the transcripts multiple times, the process moved to Phase 2 (generating initial codes). Systematic coding was conducted using NVivo Version 14 (Lumivero, 2023), where the focus was placed on recurring meaningful statements about perceived usefulness, ethical uncertainty, institutional awareness, learning support, and concerns about overreliance or societal impact. In Phase 3 (searching for themes), codes were then collated into potential themes by grouping related codes (e.g., codes relating to scepticism and accuracy concerns were grouped under what became Theme 1). In Phase 4 (reviewing themes), these candidate themes were reviewed and refined to ensure that they captured coherent patterns of meaning. This included checking the themes against the coded extracts and the full dataset to ensure internal coherence and external distinctiveness. Two candidate themes were merged during this phase, and one was discarded as it lacked sufficient support from the participants. In Phase 5 (defining and naming themes), the five final themes were refined, and in Phase 6, the report was produced with illustrative quotes.

As this was a small exploratory pilot conducted by a single researcher, inter-rater reliability was not calculated. Instead, analytic rigour was supported through an iterative revisiting of the transcripts, comparison of themes across cases, and an integration of the qualitative findings with the quantitative survey results. The absence of a second coder is acknowledged as a limitation (cf. Section 6.1), and future studies should incorporate independent coding and inter-rater agreement checks.

4. Results

Quantitative data extracted from the surveys have been summarised into two tables. The first table (Table 2) reports the corresponding Wilcoxon Signed Ranks Test results for each question, with effect size (r) and interpretation. Table 3 (found in Appendix A) presents the results from the pre-test and post-test as paired samples, categorised according to the details in Table 1. The qualitative results are a brief overview of the selectively coded themes.

The thematic analysis of the interview transcripts yielded five overarching themes that reflected the participants' perceptions, experiences, and concerns regarding GenAI in academic contexts. A thematic map is provided in Figure 2.

Theme 1: Perceptions of GenAI as a Supportive but Imperfect Tool

Students viewed GenAI tools as valuable for academic work (e.g., brainstorming and summarising) and supporting more equitable and accessible learning environments, but recognised their limitations in accuracy and originality. As one participant explained, GenAI is 'extremely good at recognising patterns [but] not really great at just making something from [prompts with little detail].' Another participant noted that AI-generated essays 'look pretty' but, after a few minutes' reading, reveal 'glaring problems' that require human correction.

Table 2
Wilcoxon Signed Ranks Test and effect size

	Z	$\|r\|$	Effect size (r)	N
Post_F1 - Pre_F1	0.000^{a}	0.000	Negligible	6
Post_F2 - Pre_F2	0.000^{a}	0.000	Negligible	6
Post_F3 - Pre_F3	-1.947^{b}	0.795	Large	6
Post_F4 - Pre_F4	-1.807^{b}	0.738	Large	6
Post_E1 - Pre_E1	-1.000^{b}	0.408	Medium	6
Post_E2 - Pre_E2	-1.633^{b}	0.667	Large	6
Post_E3 - Pre_E3	-1.890^{b}	0.772	Large	6
Post_E4 - Pre_E4	-1.342^{b}	0.548	Large	6
Post_V1 - Pre_V1	-0.577^{b}	0.236	Small	6
Post_V2 - Pre_V2	-1.000^{c}	0.408	Medium	6
Post_V3 - Pre_V3	0.000^{a}	0.000	Negligible	6

	Z	$\lvert r \rvert$	Effect size (r)	N
Post_C1 - Pre_C1	-0.447^{c}	0.183	Small	6
Post_C2 - Pre_C2	-1.633^{c}	0.667	Large	6
Post_C3 - Pre_C3	-1.656^{c}	0.676	Large	6
Post_I1 - Pre_I1	-1.633^{b}	0.667	Large	6
Post_I2 - Pre_I2	-1.633^{b}	0.667	Large	6
Post_I3 - Pre_I3	-1.414^{b}	0.577	Large	6

a. The sum of negative ranks equals the sum of positive ranks
b. Based on negative ranks
c. Based on positive ranks

Theme 2: Awareness Gaps in Institutional AI Tools

Students were unfamiliar with institutionally provided and approved GenAI tools but showed interest when being introduced to them. However, most students were familiar with other popular GenAI tools such as ChatGPT. Upon learning about the institution-approved tools, one participant stated, ‘I didn’t know that we had access to this through [the university], and I’ll definitely be using some of these in the future.’

Theme 3: Ethical Boundaries and Academic Integrity

Students expressed uncertainty about the acceptable use of GenAI and desired more explicit guidance from faculty, with some commenting on the inconsistent expectations across courses. One participant reflected, ‘You can’t force someone to critically think about something. If they’re just letting ChatGPT write everything, they were never going to develop the skills anyway.’ Others voiced similar concerns that unclear guidance could lead to misuse and undermine skill development.

Theme 4: GenAI as a Learning Companion

Students preferred to use GenAI to support, rather than replace, their own learning. Exposure to novel ways to use GenAI tools, such as generating podcasts, creating interactive flashcards, and summarising course materials, was met with enthusiasm. One student commented, ‘It’s not meant to replace studying; it’s just nice to have something to soundboard with.’ Another student also mentioned, ‘[GenAI] helps me get started, but I don’t trust it blindly.’ Discussions around hallucinations in GenAI outputs and bias present in the training model reinforced these reflections.

Theme 5: Concerns about Overreliance and Societal Impact

Participants expressed concerns about their personal overreliance on GenAI and its potential societal consequences, including job displacement and the reinforcement of systemic biases. ‘Top executives are exactly the kind of people who would accept a bad product if it means paying less,’ warned one student. Others stressed that GenAI ‘is just a tool. It’s the humans that make it a bad tool,’ highlighting the role of human agency in ethical adoption. One participant also raised concerns about the erasure of human creativity, stating, ‘AI is going to replace artists, and there’s no purpose for being an artist.’ They emphasised the need for a ‘line’ where AI remains a ‘backing voice, but not the main...replacement.’

The qualitative findings reveal a need for deeper conversations with students about GenAI. Even these brief conversations helped participants to understand the balance between recognising the benefits of GenAI and its limitations and risks.

5. Discussion

Baseline knowledge and attitudes toward GenAI (RQ1) and shifts in these attitudes following the AI-literacy session (RQ2) are entwined, since participants frequently described knowledge gains as the mechanism for attitudinal change. These two research questions are therefore discussed together, organised below according to the theoretical constructs introduced in the conceptual framework. Given the small N, these quantitative results should be treated as exploratory and interpreted cautiously. They are included to characterise the direction and possible magnitude of change in this cohort, not to claim generalisable effects.

Figure 2
Thematic map of qualitative data

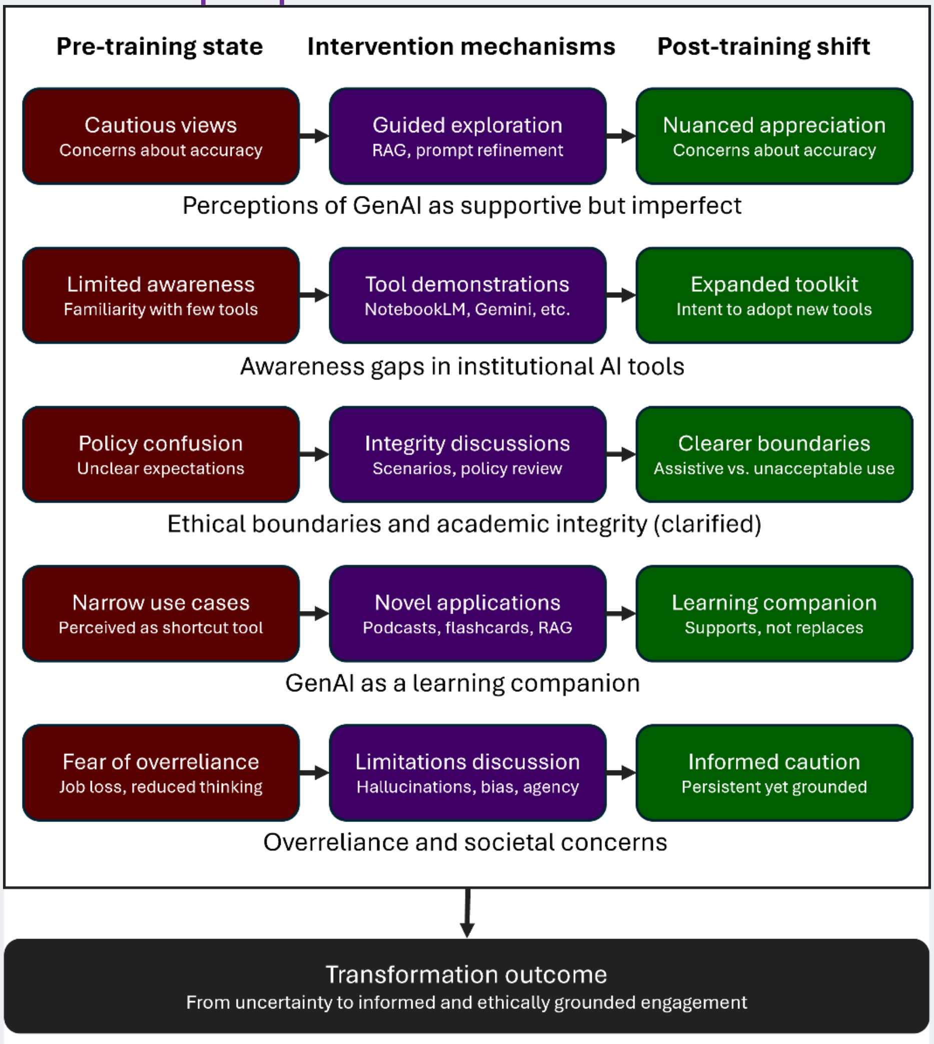

(Source: Personal archive)

5.1 Expectancy and Perceived Value: Shifting Appraisals of GenAI's Academic Utility

From an expectancy-value perspective, the intervention targeted the interrelated beliefs of students' confidence in their ability to use GenAI tools (expectancy) and their appraisal of those tools' relevance and usefulness for academic work (perceived value). The quantitative patterns observed in this pilot are consistent with the predictions of this framework. The Wilcoxon results indicate a shift in how students understand GenAI, with Familiarity items F3 (understanding how GenAI generates responses) and F4 (confidence in navigating a GenAI interface) showing large effects ($|r| = 0.795\ and\ 0.738$ respectively), suggesting that hands-on, guided practice increased participants' sense of competence. Meanwhile, items in the Expectancy and Academic Usefulness category (E1-E4) indicated medium-to-large effects ($|r| = 0.408 - 0.772$), indicating that participants' perceptions of GenAI's academic utility also shifted upward after the session.

Importantly, Future Intentions items (I1-I3) followed a similar trajectory, with large effects ($|r| = 0.577 - 0.667$). This convergence of increased expectancy and increased perceived value, followed by stronger behavioural intention, is consistent with the central prediction of EVT, where individuals are more likely to engage with a task when they feel both capable of performing it and regard it as worthwhile (Wigfield and Eccles, 2000). Qualitatively, this pattern was reflected in the participants' descriptions of how regarding NotebookLM include source links, hearing a generated study podcast, or watching a custom Gem act like a quiz coach helped to reframe GenAI as a study tool rather than a shortcut. One participant remarked, 'I didn't think I'd use it for studying, but now I can see it actually helps me go over material when I don't have time to read everything.' Another said, 'It saves time on busywork so I can focus on understanding the topic.' Students often commented on the time that could be saved on low-level tasks (such as searching for specific information or formatting a document) that could be spent on higher-order work (such as analysis and synthesis).

These findings extend the cross-sectional work of Chan and Zhou (2023), who have reported a strong positive correlation between perceived value and intention to use GenAI among 405 students using an EVT-based instrument. Whereas their design captured these relations at a single time point, the present study's pre-post design offers preliminary evidence that the expectancy-value-intention relation can be shifted through a brief, targeted intervention. This is a modest but meaningful extension: It suggests that rather than being fixed dispositions, students' expectancies and value beliefs about GenAI are responsive to pedagogical experiences, even within a single session. This finding also aligns with Wang *et al*. (2023) who have found that supportive environments play a significant role in improving expectancy-value beliefs around AI literacy, and with Bai and Wang (2025) who have emphasised the mediating role of interaction quality in shaping students' learning outcomes with GenAI.

Conversely, the null changes observed on F1 and F2 (general familiarity with, and prior use of GenAI) are also interpretable through an expectancy-value lens. These items had high baseline means ($4.500\ and\ 3.833$ respectively), indicating that most participants already considered themselves familiar with GenAI prior to the intervention. This ceiling effect is consistent with EVT's prediction that expectancy beliefs, which are already near their maximum, leave limited room for further change (Wigfield, 1994). In practical terms, these students did not need the intervention to know that GenAI exists. However, they still needed to understand how it operates and why it might be worth using thoughtfully, which is precisely where the largest shifts were observed.

5.2 Perceived Cost and Ethical Reasoning: A Constructivist Reappraisal

The Concerns and Ethical Considerations items (C1-C3) offer insight into what EVT refers to as perceived costs. These costs are the barriers, risks, and anxieties that may discourage engagement even when value and expectancy are high. Item C2 (uncertainty about whether GenAI constitutes academic dishonesty) and C3 (worry about overreliance) both indicated large downward effects ($|r| = 0.676$), while C1 (concern about reduced original thinking) indicated only a small effect ($|r| = 0.183$). From a purely quantitative standpoint, the reduction in C2 and C3 suggests that the intervention lowered perceived cost which, according to EVT, should reduce barriers to engagement and support the increased intentions observed in the I-category items.

However, the qualitative data reveal that this cost reduction was not simply a matter of students becoming less worried. Instead, it reflected a constructivist reappraisal of what ethical GenAI use means. Pre-session ambiguity around 'Is this dishonesty?' became a point of discussion during the training, where a distinction was made between assistive use with transparency versus unacceptable use. Through dialogue and reflection, participants moved from binary thinking ('Is GenAI cheating or not?') toward conditional, context-dependent reasoning. As one student noted, 'It's not cheating if I'm using it to learn, as long as I'm honest about how I used it.' Another observed, 'AI doesn't decide if it's ethical, we do.' These statements suggest that participants actively constructed a more nuanced ethical framework through social interaction and critical reflection, which is characteristic of constructivist learning.

This finding refines the observations of Chan and Hu (2023), as well as Lund *et al*. (2025), who documented widespread student confusion about whether GenAI use constitutes academic misconduct. Whereas those studies identified the confusion, the present study offers preliminary evidence that even brief, constructivist-oriented discussions, in which ethical reasoning is embedded within authentic workflows rather than delivered as a standalone lecture, can help students develop more differentiated positions. Notably, the small effect on C1 (concern about reduced original thinking) suggests that not all cost perceptions are equally amenable to short-term change. Some concerns may be more deeply held or may reflect legitimate caution rather than unfamiliarity. Two participants remained wary of overreliance even after the session. One put it, 'It's useful, but I don't want it to do the whole thing for me. I still want to learn it myself.' That stance aligns with the session's goal, which was to inform participants about their agency. Critically, from a constructivist perspective, this persistent concern should not be interpreted as a failure of the intervention. Rather, some students constructed a more sophisticated worry, not less worry, precisely because they now understood the tools' capabilities. As one student reflected, 'I'm still careful with it, but now I know what to be careful about.' This informed caution represents a qualitative change in the nature of the concern, even where the quantitative score did not shift substantially.

5.3 Value and Interest: Stability in Intrinsic Motivation

The Value and Interest items (V1-V3) indicated comparatively little movement: V1 (curiosity about GenAI for studying) disclosed a small effect ($|r| = 0.236$), V2 (interest in learning how GenAI works) indicated a medium effect ($|r| = 0.408$), while V3 (importance of learning ethical use) displayed no change ($|r| = 0.000$), with a near-ceiling baseline mean of 4.833. These patterns suggest that students' intrinsic interest in GenAI and their valuing of ethical engagement were already well established before the intervention.

This is consistent with broader survey findings. Ravšelj *et al*. (2025), in a study of over 23,000 students from more than 100 countries, found that students generally reported high curiosity and interest in GenAI tools, while Blahopoulou and Ortiz-Bonnin (2025) observed that even students who had not used GenAI still expressed interest in its potential. The stability of V3 in particular, with its near-ceiling mean, aligns with prior findings that students broadly recognise the importance of ethical AI use, even when they are unsure about what ethical use looks like in practice (as reflected in the pre-training C2 scores). This distinction between motivating students toward GenAI and equipping motivated students to use it well has implications for how institutions design AI literacy programmes, as it suggests that programmes should prioritise guided practice and critical reflection over persuasion.

5.4 Constructivist Pedagogy and Workshop Design Mechanisms

For RQ3, the three workshop design features that appear to drive the observed changes are:

1. Students' interactions with RAG tools, where they could observe source-level anchoring, increased perceived transparency, and control of outputs. One participant stated, 'Once I saw where the answers came from, it felt less like magic and more like something I could control.' From a constructivist perspective, this feature made the tool's reasoning process visible, allowing students to evaluate and interrogate outputs rather than passively accepting them. This aligns with the constructivist emphasis on students actively questioning and verifying knowledge.
2. Academic integrity was not a separate discussion and was embedded in each workflow (e.g., when to verify, and what to disclose). A participant acknowledged, 'You don't have to trust AI completely…you can check.' This embedded approach reflects the constructivist principle that knowledge is best constructed in context. It also aligns with the concept of situated literacies, where AI literacy is understood as context-specific rather than a universal, decontextualised competency.
3. Just-in-time learning examples matched authentic academic tasks (outline, create flashcards, and generate practice questions) which, as one participant put it, helped them to understand how they 'can use this tonight.' This design feature exemplifies the constructivist emphasis on situating learning in authentic, personally meaningful contexts, and it may explain why the largest effects were observed in the Expectancy (E) and Intention (I) categories. Participants could directly envision how GenAI would fit into their existing academic practices, thereby increasing both perceived utility value and expectancy.

These three mechanisms provide some guidance on what constructivist AI literacy training looks like in practice, contributing to the growing but still limited evidence base on pedagogical design for GenAI education. Importantly, these design features also connect back to the expectancy-value framework by increasing transparency (which builds expectancy), demonstrating practical relevance (which builds perceived value), and reducing uncertainty about ethical boundaries (which lowers perceived cost).

6. Conclusion

This study set out to examine whether a focused intervention could 'prompt change' in how undergraduates perceive and use GenAI. After engaging with current GenAI tools and reflecting on their use, students shifted from uncertainty to enthusiasm and a commitment to the ethical use of GenAI. They discovered that these tools, when approached with critical thinking, can be powerful for their learning, helping to personalise their study process. At the same time, they gained a better understanding of GenAI's limitations, therefore becoming more AI-literate students who can leverage the technology's strengths for their education. Just as digital literacy became essential in the last few decades, AI literacy is rapidly becoming indispensable. These tools are becoming ubiquitous in society and by supporting students in developing AI literacy now, the purposeful implementation of training not only enhances their current learning outcomes but also empowers them with skills and mindsets for lifelong learning, where an overreliance on these tools could lead to the loss of critical thinking.

Prompting change, on the one hand, refers to the idea that by teaching students how to prompt GenAI effectively, they can change the way in which they approach learning tasks (by making them more efficient and giving them more time to be creative). On the other hand, it signifies prompting a broader change in educational practice, in an effort to change the culture to one where GenAI is integrated into learning in a thoughtful way, rather than left unregulated. The experiences of the six students in this study offer a glimpse into how this can be achieved. Educators must be proactive in equipping students with AI literacy. This involves providing resources and training, establishing clear ethical guidelines, and encouraging exploration and responsible innovation. One participant aptly noted, ‘AI should be fully integrated into teaching and learning objectives…[focusing on] responsible and productive use.’

6.1 Limitations and Future Research

While this study yielded valuable insights, it has limitations. First, the sample size ($N = 6$) was small and not demographically diverse since all participants were from the same university and of traditional college age, that substantially limit both statistical inference and external validity. Nonparametric tests (e.g., Wilcoxon) have limited power to detect small-to-moderate changes. Therefore, null findings should not be interpreted as evidence of no effect. Conversely, apparent ‘large’ effects may be unstable and disproportionately influenced by one or two participants, and p-values are highly sensitive to single-case shifts. Accordingly, the quantitative results are best interpreted as exploratory within-participant patterns that motivate hypotheses and inform the design of future interventions, rather than as confirmatory evidence of population-level effects. The primary contribution of this pilot is therefore mechanistic and design-oriented, identifying plausible learning mechanisms and illustrating how students reason about GenAI after a targeted literacy intervention within a specific context, rather than estimating generalisable effect magnitudes. To reiterate, this study functioned as a pilot, and larger studies are necessary across varied contexts. Future research should involve a control group and a larger experimental group to more rigorously measure the impact of AI literacy interventions on actual academic performance or longer-term attitude changes.

Second, the data on actual behaviour change are indirect, as they comprise self-reported intentions and immediate post-training enthusiasm. How these students will integrate GenAI into their study practices over time is unknown. There is a risk of ‘workshop effect’ wear-off. Longitudinal research tracking students’ GenAI usage patterns over time (with tools that log usage or through periodic surveys/interviews) would be valuable. This longitudinal data would help answer questions such as, ‘Do students continue to apply what they learned from the AI literacy training?’ and ‘Does AI literacy training have any measurable effect on learning outcomes or efficiency (e.g., do students produce better essays, or do they save time on homework)?’

Another limitation is that the qualitative analysis is interpretive and potentially subject to researcher bias. As the trainer and principal investigator were the same person, there is a possibility of inadvertently leading students or interpreting their statements in an overly optimistic manner. In an effort to mitigate this, direct quotes were placed throughout the paper, but bias cannot be truly mitigated without cross-validation and interrater agreement. Future studies should consider having independent observers or coders, and possibly even different trainers to see if results are replicable – is there anything idiosyncratic about a facilitator's style that influenced outcomes?

While the core literacies (prompting, evaluating output, and ethical use) remain broadly applicable, future training needs to adapt to the continually updating technology and the emergence of new GenAI systems. Therefore, research should continue to evaluate AI literacy interventions for new modalities and tools, ensuring that pedagogical approaches keep pace with technological advances (e.g., how to teach students to critically evaluate AI-generated diagrams or deepfake videos in an academic context).

Finally, while this study focused on students, an important complementary angle is faculty development. For AI literacy to truly support student learning, instructors themselves must be literate and comfortable with GenAI, so that they can design assignments that incorporate GenAI constructively and evaluate student work fairly. This study did not address faculty perspectives, but some participants noted inconsistent policies among educators. Future research might explore how faculty can be trained to create an environment conducive to AI literacy, using AI-responsive assessments.

7. Acknowledgements

The author discloses that he has no actual or perceived conflicts of interest. The author discloses that he has not received any funding for this manuscript beyond resourcing for academic time at his university. The author has not used artificial intelligence in the ideation, design, or write-up of this research.

It also needs to be stated that in this article no reporting checklists are applicable. The writer of this article did not use AI or an LLM to write this article or correct it.

8. References

Abuzar, M., Mahmudulhassan[1] and Muthoifin.[1] (2025) 'University students' trust in AI: Examining reliance and strategies for critical engagement,' *International Journal of Interactive Mobile Technologies*, 19(07), pp. 70-82. https://doi.org/10.3991/ijim.v19i07.52875

Adobe. (2024) *Generative AI in higher education: Insights from campus thought leaders*. Adobe Blog. Available at: https://blog.adobe.com/en/publish/2024/06/28/generative-ai-in-higher-education-insights-from-campus-thought-leaders. (Accessed 28 June 2025).

Alexander, K., Savvidou, C. and Alexander, C. (2023) 'Who wrote this essay? Detecting AI-generated writing in second language education in higher education,' *Teaching English with Technology*, 2023(2), pp. 25-43. https://doi.org/10.56297/BUKA4060/XHLD5365

Apata, O.E., Ajamobe, J.O., Ajose, S.T., Oyewole, P.O. and Olaitan, G. (2025) 'The role of Artificial Intelligence in enhancing classroom learning: Ethical, practical, and pedagogical considerations,' *2025 ASEE -GSW Annual Conference Proceedings*. *2025 ASEE -GSW Annual Conference*, Arlington, TX, Texas: ASEE Conferences. 55084. 9 pages. https://doi.org/10.18260/1-2--55084

[1] No initials given.

Bai, Y. and Wang, S. (2025) ‘Impact of generative AI interaction and output quality on university students’ learning outcomes: A technology-mediated and motivation-driven approach,’ *Scientific Reports*, 15(1). 24054. 16 pages. https://doi.org/10.1038/s41598-025-08697-6

Baidoo-Anu, D., Asamoah, D., Amoako, I. and Mahama, I. (2024) ‘Exploring student perspectives on generative artificial intelligence in higher education learning,’ *Discover Education*, 3(1). 98. 21 pages. https://doi.org/10.1007/s44217-024-00173-z

Baytas, C. and Ruediger, D. (2025) *Making AI generative for higher education: Adoption and challenges among instructors and researchers*. Ithaka S+R. Available at: https://sr.ithaka.org/publications/making-ai-generative-for-higher-education/. (Accessed 28 November 2026).

Blahopoulou, J. and Ortiz-Bonnin, S. (2025) ‘Student perceptions of ChatGPT: benefits, costs, and attitudinal differences between users and non-users toward AI integration in higher education,’ *Education and Information Technologies* 30, pp. 19741-19764. https://doi.org/10.1007/s10639-025-13575-9

Braun, V. and Clarke, V. (2006) ‘Using thematic analysis in psychology,’ *Qualitative Research in Psychology*, 3(2), pp. 77-101. https://doi.org/10.1191/1478088706qp063oa

Chan, C.K.Y. (2023) ‘Is AI changing the rules of academic misconduct? An in-depth look at students’ perceptions of “AI-giarism”,’ *arXiv*, 2306.03358. https://doi.org/10.48550/ARXIV.2306.03358

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*, 20(1). 43. 18 pages. https://doi.org/10.1186/s41239-023-00411-8

Chan, C.K.Y. and Zhou, W. (2023) ‘Deconstructing student perceptions of Generative AI (GenAI) through an Expectancy Value Theory (EVT)-based instrument,' *arXiv*, 2305.01186. https://doi.org/10.48550/ARXIV.2305.01186

Cohen, J. (2013) *Statistical power analysis for the behavioral sciences*. Burlington: Elsevier Science.

Cotton, D.R.E., Cotton, P.A. and Shipway, J.R. (2023) ‘Chatting and cheating: Ensuring academic integrity in the era of ChatGPT,’ *Innovations in Education and Teaching International*, 61(2), pp. 228-239. https://doi.org/10.1080/14703297.2023.2190148

Federiakin, D. Molerov, D., Zlatkin-Troitschanskaia, O. and Maur, A. (2024) ‘Prompt engineering as a new 21st century skill,’ *Frontiers in Education*, 9. 1366434. 14 pages. https://doi.org/10.3389/feduc.2024

Gabriel, S. (2024) ‘Generative AI and educational (in)equity,’ *International Conference on AI Research*, 4(1), pp. 133-142. https://doi.org/10.34190/icair.4.1.3153

Gasaymeh, A.-M.M., Beirat, M.A. and Abu Qbeita, A.A. (2024) ‘University students’ insights of Generative Artificial Intelligence (AI) writing tools,’ *Education Sciences*, 14(10). 1062. 18 pages. https://doi.org/10.3390/educsci14101062

Google. (n.d) *Gemini AI video generator powered by Veo 3: Break the silence with Veo 3*. Available at: https://gemini.google/overview/video-generation/?hl=en (Accessed 30 June 2025).

IBM Corp. (2019) *IBM SPSS statistics for Windows*. Armonk: IBM Corp. Available at: https://www.ibm.com/products/spss-statistics. (Accessed 30 June 2025).

Knoth, N., Tolzin, A., Janson, A. and Leimeister, J.M. (2024) ‘AI literacy and its implications for prompt engineering strategies,’ *Computers and Education: Artificial Intelligence*, 6. 00225. 14 pages. https://doi.org/10.1016/j.caeai.2024.100225

Leslie, D. and Perini, A.M. (2024) ‘Future shock: Generative AI and the International AI Policy and governance crisis,’ *Harvard Data Science Review* [Preprint], (Special Issue 5). 92 pages. https://doi.org/10.1162/99608f92.88b4cc98

Long, D. and Magerko, B. (2020) ‘What is AI literacy? Competencies and design considerations,’ *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. CHI ‘20: CHI Conference on Human Factors in Computing Systems*, Honolulu HI USA: ACM, 16 pages. https://doi.org/10.1145/3313831.3376727

Lumivero (2023) *NVivo*. Lumivero. Available at: https://www.lumivero.com/. (Accessed 30 June 2025).

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, 1545-1565. https://doi.org/10.1007/s10805-025-09613-3

Microsoft and LinkedIn. (2024) *AI at work is here. Now comes the hard part*. Microsoft Work Lab. Available at: https://www.microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part. (Accessed 28 June 2025).

Park, J. and Choo, S. (2025) ‘Generative AI prompt engineering for educators: Practical strategies,’ *Journal of Special Education Technology*, 40(3), pp. 411-417. https://doi.org/10.1177/01626434241298954

Qian, Y. (2025) ‘Prompt engineering in education: A systematic review of approaches and educational applications,’ *Journal of Educational Computing Research*, 63(7-8), pp. 1782-1818. https://doi.org/10.1177/07356331251365189

Rapanta, C., Bhatt, I., Bozkurt, A., Chubb, L.A., Erb, C. Forsler, I., Gravett, K., Koole, M., Lintner, T., Örtegren, A., Petricini, T., Rodgers, B., Webster, J., Xu, X., Christensen, I.-M.F, Dohn, N.B., Christensen, L.L.W., Zeivots, S., and Jandrić, P. (2025) ‘Critical GenAI literacy: Postdigital configurations,’ *Postdigital Science and Education*, 7, pp. 1296-1333. https://doi.org/10.1007/s42438-025-00573-w

Ravšelj, D. Keržič, D., Tomaževič, N., Umek, L., Brezovar, N., Iahad, N.A., Abdulla, A.A., Akopyan, A., Segura, M.W.A., AlHumaid, J., Allam, M.F., Alló, M., Andoh, R.P.K., Andronic, O., Arthur, Y.D., Aydın, F., Badran, A., Balbontín-Alvarado, R., Saad, H.B., Bencsik, A., Benning, I., Besimi, A., Da Silva Bezerra, D., Buizza, C., Burro, R., Bwalya, A., Cachero, C., Castillo-Briceno, P., Castro, H., Chai, C.S., Charalambous, C., Chiu, T.K.F., Clipa, O., Colombari, R., Corral Escobedo, L.J.H., Costa, E., Crețulescu, R.G., Crispino, M., Cucari, N., Dalton, F., Kaya, M.D., Dumić-Čule, I., Dwidienawati, D., Ebardo, R., Egbenya, D.L., Faris, M.E., Fečko, M., Ferrinho, P., Florea, A., Fong, C.Y., Francis, Z., Ghilardi, A., González-Fernández, B., Hau, D., Hossain, Md.S., Hug, T., Inasius, F., Ismail, M.J., Jahić, H., Jessa, M.O., Kapanadze, M., Kar, S.K., Kateeb, E.T., Kaya, F., Khadri, H.O., Kikuchi, M., Kobets, V.M., Kostova, K.M., Krasmane, E., Lau, J., Law, W.H.C., Lazăr, F., Lazović-Pita, L., Lee, V.W.Y., Li, J., López-Aguilar, D.V., Luca, A., Luciano, R.G., Machin-Mastromatteo, J.D., Madi, M., Manguele, A.L., Manrique, R.F., Mapulanga, T., Marimon, F., Marinova, G.I., Mas-Machuca, M., Mejía-Rodríguez, O., Meletiou-Mavrotheris, M., Méndez-Prado, S.M., Meza-Cano, J.M., Mirķe, E., Mishra, A., Mital, O., Mollica, C., Morariu, D.I., Mospan, N., Mukuka, A., Jiménez, S.G.N., Nikaj, I., Nisheva, M.M., Nisiforou, E., Njiku, J., Nomnian, S., Nuredini-Mehmedi, L., Nyamekye, E., Obadić, A., Okela, A.H., Olenik-Shemesh, D., Ostoj, I., Peralta-Rizzo, K.J., Peštek, A., Pilav-Velić, A., Pires, D.R.M., Rabin, E., Raccanello, D., Ramie, A., Rashid, Md.M.ur, Reuter, R.A.P., Reyes, V., Rodrigues, A.S., Rodway, P., Ručinská, S., Sadzaglishvili, S., Salem, A.A.M.S., Savić, G., Schepman, A., Shahpo, S.M., Snouber, A., Soler, E., Sonyel, B., Stefanova, E., Stone, A., Strzelecki, A., Tanaka, T., Cortes, C.T., Teira-Fachado, A., Tilga, H., Titko, J., Tolmach, M., Turmudi, D., Varela-Candamio, L., Vekiri, I., Vicentini, G., Woyo, E., Yorulmaz, Ö., Yunus, S.A.S., Zamfir, A.-M., Zhou, M. and Aristovnik, A. (2025) 'Higher education students' perceptions of ChatGPT: A global study of early reactions,' *PLOS ONE*, 20(2). 53 pages. https://doi.org/10.1371/journal.pone.0315011

Rebelo, M. (2025) *An inside look at Google's AI-powered NotebookLM*. Available at: https://zapier.com/blog/google-ai-notebook-notebooklm/. (Accessed 28 June 2025).

Schiff, K.J., Schiff, D.S. and Bueno, N.S. (2025) 'The liar's dividend: Can politicians claim misinformation to evade accountability?,' *American Political Science Review*, 119(1), pp. 71-90. https://doi.org/10.1017/S0003055423001454

Wang, F., King, R.B., Char, C.S. and Zhou, Y. (2023) 'University students' intentions to learn artificial intelligence: The roles of supportive environments and expectancy-value beliefs,' *International Journal of Educational Technology in Higher Education*, 20(1). 51. 21 pages. https://doi.org/10.1186/s41239-023-00417-2

Wigfield, A. (1994) 'Expectancy-value theory of achievement motivation: A developmental perspective,' *Educational Psychology Review*, 6(1), pp. 49-78. https://doi.org/10.1007/BF02209024

Wigfield, A. and Eccles, J.S. (2000) 'Expectancy-Value Theory of Achievement Motivation,' *Contemporary Educational Psychology*, 25(1), pp. 68-81. https://doi.org/10.1006/ceps.1999.1015

Yang, Y., Zhang, Y., Sun, D., He, W. and Wei, Y. (2025) 'Navigating the landscape of AI literacy education: insights from a decade of research (2014-2024),' *Humanities and Social Sciences Communications*, 12(1). 374. 12 pages. https://doi.org/10.1057/s41599-025-04583-8

Yarlagadda, K.C. (2025) 'AI in Education: Personalized learning and intelligent tutoring systems,' *European Journal of Computer Science and Information Technology*, 13(32), pp. 15-27. https://doi.org/10.37745/ejcsit.2013/vol13n321527

Appendix A

Table 3
Paired samples descriptive statistics

		Mean	Median	Mode	*N*	Std. Deviation	Std. Error Mean
Pair 1	Pre_F1	4.500	4.500	4[b]	6	0.548	0.224
	Post_F1	4.500	4.500	4[b]	6	0.548	0.224
Pair 2	Pre_F2	3.833[a]	4.000	4	6	1.472	0.601
	Post_F2	3.833[a]	4.000	4	6	1.472	0.601
Pair 3	Pre_F3	3.333	3.000	3	6	1.033	0.422
	Post_F3	4.667	5.000	5	6	0.516	0.211
Pair 4	Pre_F4	3.000	3.000	3	6	1.414	0.577
	Post_F4	4.333	4.000	4	6	0.516	0.211
Pair 5	Pre_E1	4.000	4.000	4	6	0.632	0.258
	Post_E1	4.167	4.000	4	6	0.753	0.307
Pair 6	Pre_E2	3.333	3.500	3[b]	6	1.366	0.558
	Post_E2	4.000	4.000	4	6	1.095	0.447
Pair 7	Pre_E3	3.167	3.000	3	6	1.329	0.543
	Post_E3	4.000	4.000	4	6	1.095	0.447
Pair 8	Pre_E4	3.500	4.000	4	6	1.225	0.500
	Post_E4	4.333	4.000	4	6	0.516	0.211
Pair 9	Pre_V1	4.000	4.000	3[b]	6	0.894	0.365
	Post_V1	4.167	4.000	4	6	0.753	0.307
Pair 10	Pre_V2	4.167	4.000	4	6	0.753	0.307
	Post_V2	3.833	4.000	4[b]	6	1.169	0.477
Pair 11	Pre_V3	4.833[a]	5.000	5	6	0.408	0.167
	Post_V3	4.833[a]	5.000	4[b]	6	0.408	0.167
Pair 12	Pre_C1	2.833	2.500	2	6	1.169	0.477
	Post_C1	2.667	2.500	5	6	1.366	0.558
Pair 13	Pre_C2	2.333	2.500	1[b]	6	1.211	0.494
	Post_C2	1.667	1.000	2[b]	6	1.211	0.494
Pair 14	Pre_C3	3.833	3.500	3	6	0.983	0.401
	Post_C3	2.833	3.000	1	6	1.722	0.703
Pair 15	Pre_I1	3.167	3.000	3	6	1.329	0.543
	Post_I1	3.833	4.000	1[b]	6	0.753	0.307
Pair 16	Pre_I2	3.167	3.000	3	6	1.329	0.543
	Post_I2	3.833	4.000	4	6	0.753	0.307
Pair 17	Pre_I3	4.000	4.000	4	6	1.095	0.447
	Post_I3	4.333	4.500	5	6	0.816	0.333

a. Identical values between pre-test and post-test

b. Multiple modes exist, with the smallest shown