

DOI: https://doi.org/10.48009/4_iis_2026_132

Guided, not given: Designing and pilot testing an AI-powered micro-lesson for teaching AI ethics

Joseph Rene Corbeil, *The University of Texas Rio Grande Valley*, rene.corbeil@utrgv.edu

Maria Elena Corbeil *The University of Texas Rio Grande Valley*, mariaelena.corbeil@utrgv.edu

Abstract

This article describes the design, development, and pilot testing of an AI-powered micro-lesson on ethical AI use in education. The lesson used a chatbot named Ava to guide graduate Educational Technology students through three case studies focused on transparency, authorship, accuracy, bias, and verification. Rather than providing direct answers, Ava used open-ended questions, tiered hints, and feedback to support student reasoning. The pilot test compared the AI-powered micro-lesson with a non-chatbot version of the same lesson. Although the AI-powered version demonstrated a larger descriptive gain score, the difference was not statistically significant. Qualitative reflections indicated that students perceived Ava as engaging, reflective, and supportive of guided reasoning. Findings suggest that chatbot-guided micro-lessons may support interaction, reflection, and perceived engagement with ethical reasoning when intentionally designed as learning guides.

Keywords: AI ethics, educational technology, chatbot, microlearning, instructional design, artificial intelligence in education

Introduction

Generative AI tools are rapidly becoming part of students' academic lives (Stephenson & Armstrong, 2026), creating both opportunities and challenges for educators. These tools can support brainstorming, writing, revision, feedback, and study. However, they can also encourage students to bypass important parts of the learning process by generating answers, summaries, or polished written products without requiring them to fully engage with the content. This creates an instructional challenge: How can educators help students use AI in ways that support learning rather than replace the thinking students are expected to do? One promising response is to design AI-supported learning experiences that position the chatbot as a guide rather than an answer provider. Instead of using AI to deliver direct responses, educators can design chatbots to ask questions, provide hints, redirect learners to relevant information, and encourage students to explain their reasoning. In this way, the chatbot becomes

a form of instructional scaffold, helping learners work through a problem or scenario, while preserving the cognitive effort needed for meaningful learning.

This issue is especially important in the area of AI ethics. Students need more than general warnings about academic dishonesty or overreliance on AI. They need opportunities to apply ethical principles to realistic academic situations, such as using AI to draft or revise written work, deciding when and how to disclose AI assistance, preserving their own voice and authorship, and verifying AI-generated information for accuracy and bias. Case-based learning is well suited to this need because it allows students to examine specific choices, consider consequences, and practice ethical decision-making in context rather than simply recall general rules about responsible AI use.

To address this need, we developed an AI-powered micro-lesson on *AI Ethics in Education*. The lesson used a chatbot named Ava to guide students through three short case studies focused on common forms of AI use in academic work. Ava was intentionally designed not to provide direct answers. Instead, the chatbot asked open-ended questions, offered tiered hints, and guided students toward appropriate conclusions through dialogue. The goal was to create a brief, but complete instructional experience that combined content presentation, learner interaction, multimedia, and assessment, while modeling responsible AI use.

This article reports a design-and-development pilot study that describes the creation of the AI-powered micro-lesson and examines preliminary evidence of student learning and perceptions. The article first reviews literature related to AI chatbots, microlearning, and case-based learning. It then describes the design process used to develop the chatbot-guided lesson, followed by the methodology and findings from a pilot comparison of the AI-powered micro-lesson and a non-chatbot version of the same lesson.

Review of Literature

AI Chatbots as Scaffolds for Learning

Educational chatbots are software tools designed to interact with learners through text or voice, often within websites, mobile applications, learning management systems, or other digital learning environments. In educational settings, they can function as scaffolds for learning when they guide learners through the thinking process rather than simply provide correct answers. In this role, chatbots may support learners as they work through a task by asking questions, offering feedback and hints, and directing attention to relevant information. Kumar (2021) examined the use of educational chatbots as instructional supports that provide personalized, conversational interaction. In a quasi-experimental study with 60 undergraduate students enrolled in an instructional design course, Kumar compared a chatbot-supported group with a traditional classroom control group. Students in the chatbot-supported group demonstrated higher learning performance and teamwork outcomes. However, the chatbot was not associated with significant differences in affective or motivational outcomes, including perception of learning, need for cognition, motivation, and creative self-efficacy.

Broader reviews of educational chatbot research provide additional context for understanding chatbot design patterns before the widespread adoption of generative AI tools. Kuhail et al. (2023), for example, conducted a systematic review of 36 educational chatbot studies identified from literature published between 2011 and 2021. They found that educational chatbots were most often designed as teaching agents, which provide instructional guidance, or peer agents, which interact with learners more like study partners or collaborators. Many of the chatbots followed scripted conversation paths, meaning that learners were guided through preplanned prompts and responses, while some adapted their responses based on learner needs. Although the reviewed studies often reported improved learning, usefulness, or learner satisfaction, Kuhail et al. also noted recurring design limitations, including insufficient dataset training and limited attention to usability principles.

Lee et al. (2025) provide a more recent example of scaffolded chatbot design. In their study of second-year students enrolled in a university multimedia course, they examined a self-regulated learning intervention that used chatbot-guided feedback during students' reflection activities. Students in the experimental group worked through these reflections with the chatbot, while students in the comparison group completed reflection sheets and received teacher feedback. Results indicated that, compared with conventional instruction, the chatbot-supported intervention was associated with higher learning performance and higher levels of motivation, self-regulated learning, reflection, and metacognitive tendency. In discussing these findings, Lee et al. noted that students reported the chatbot provided information, guidance, and support for reflecting on and thinking through problems.

Generative AI chatbots raise a related, but more complex instructional issue because they can support open-ended dialogue, while also shaping how much students contribute to the interaction. Song et al. (2025) found that students who interacted with a generative AI chatbot demonstrated stronger creative problem-solving performance than students who interacted with peers. The chatbot supported students by providing knowledge-based responses, elaborating on complex ideas, organizing information, and helping students work through a research proposal task. However, students produced less dialogue when interacting with the chatbot than when interacting with peers.

Other studies point to the importance of metacognitive feedback in chatbot-supported learning. Yin et al. (2024) compared educational chatbot interactions that used metacognitive feedback with interactions that used neutral feedback. Their findings indicated that metacognitive feedback reduced some negative emotions, including distress, nervousness, and jitteriness, and was associated with higher interest, perceived competence, and perceived value. The authors also emphasized that feedback could help learners understand their progress, evaluate their responses, and adjust their learning strategies. This is relevant to scaffolded chatbot design because students may need support that helps them monitor their own understanding, not just confirmation that an answer is correct or incorrect.

At the same time, recent findings suggest the need for caution in assuming that AI-supported scaffolding will improve all outcomes. In a study of first-semester university students from business informatics and social work-related programs, Melanou and Beege (2026) found that

students using generative AI as a tutor showed gains in knowledge, critical thinking, and reflective use, while motivation remained stable. However, differences between the full-scaffolding, light-scaffolding, and control conditions were not statistically significant. These findings suggest that short-term learning gains in AI-supported settings may depend less on scaffolding intensity alone and more on how learners engage with the AI during the learning process.

The distinction between guided learning and answer-giving is especially important in educational uses of generative AI. Sørhaug (2026), in a thematic analysis of Norwegian K-12 students' use of educational chatbots for writing support, found that students requested several types of support, including information, structural guidance, examples, content creation, feedback, and follow-up clarification. However, the chatbot often generated content on students' behalf, even when students appeared to seek scaffolding-oriented support. Sørhaug noted that effective chatbot scaffolding requires not only stronger student prompting strategies, but also system designs that better prioritize pedagogical support over extensive text creation. This finding is important because it helps demonstrate that a chatbot does not automatically provide instructional guidance simply because it can respond conversationally. Without clear instructional boundaries, the interaction may shift from helping students think through the task to producing the work for them.

The literature points to the value of designing chatbots as guided learning tools rather than answer providers. When their role is deliberately limited and pedagogically structured, chatbots can ask open-ended questions, provide progressive hints, offer feedback, prompt reflection, and encourage learners to justify their reasoning. This approach preserves the learner's cognitive responsibility, while still providing support. For AI ethics instruction, this form of scaffolding is particularly appropriate because students need opportunities to reason through complex situations, consider alternatives, and apply principles such as transparency, authorship, accuracy, bias, and verification in context.

Microlearning and Case-Based Learning for AI Ethics

Brief instructional interventions are often used when learners need focused guidance on a specific concept, decision, or practice, such as ethical decision making related to AI use. In these contexts, microlearning provides a useful structure because it focuses instruction on a small, clearly defined learning goal that can be supported in a short period. In a systematic review of microlearning studies, Monib et al. (2025), described microlearning as “targeted, action-oriented, bite-sized content” designed “to achieve specific objectives within a short period” (p. 1). Their framework further emphasizes the importance of a single objective, concise content, and manageable cognitive load. These instructional design considerations are especially important when a brief learning activity addresses ethical decision making, where the objective, learning activity, and assessment need to remain closely aligned.

Case-based learning is also appropriate for AI ethics instruction because ethical decision-making often requires students to apply principles in context rather than simply recall definitions. Barry et al. (2025) described case-based learning as a way for students to apply theoretical principles to real-world scenarios and noted its value for fostering decision-making and self-reflection in

situations involving uncertainty and competing interests. In academic work, students may understand general rules about honesty, disclosure, or accuracy, but still struggle to decide what responsible AI use looks like in a realistic situation. Cases give students an opportunity to examine specific choices, consider consequences, and explain their reasoning. Barry et al. also noted that although case-based learning is supported in ethics education, the complexity of AI requires careful attention to how cases are integrated into instruction.

Microlearning, case-based learning, and responsible AI use provide complementary support for brief AI ethics instruction. Microlearning supports a focused instructional structure, case-based learning gives learners realistic situations in which to apply ethical principles, and responsible AI use provides the conceptual foundation for evaluating AI-supported academic work. These approaches support instruction that asks learners to reason through ethical situations rather than simply receive rules or definitions.

Purpose of the Study and Research Questions

The purpose of this study is to describe the design, development, and pilot testing of an AI-powered micro-lesson that used a chatbot to guide students through case-based learning without providing direct answers. The article presents a replicable process for creating a chatbot-guided micro-lesson, including the development of realistic cases, the design of the chatbot's base prompt, the use of tiered hints, the creation of brief pre- and post-test assessments, and the refinement of the chatbot based on testing and student feedback.

In addition to describing the design process, this article reports preliminary evidence from a pilot implementation with graduate Educational Technology students. Quantitative data were used to compare learning gains between students who completed the AI-powered micro-lesson and students who completed a non-chatbot version of the same lesson. Qualitative data from student reflections were used to examine how learners experienced the chatbot, including how it shaped their understanding, thinking process, and engagement with AI ethics content.

The study was guided by the following research questions:

1. How do student learning gains compare between the AI-powered micro-lesson and the non-chatbot micro-lesson?
2. How did students describe the AI chatbot's influence on their understanding, thinking process, and engagement with AI ethics content?
3. What strengths and limitations did students identify in using a chatbot-guided micro-lesson to support learning?

The following section describes how the AI-powered micro-lesson was developed, including the selection of the chatbot platform, design of the case-based lesson structure, creation of the base prompt, and refinement of the chatbot's instructional behavior.

Developing the AI-Powered Micro-lesson

Platform Selection and Chatbot Configuration

In Corbeil and Corbeil (2024), we described the process used to evaluate and select an AI chatbot platform for online course integration. After reviewing several available tools, we selected *AI Coaches* (<https://aicoaches.io>) because it aligned with our instructional and practical requirements, including an intuitive design interface, ease of use, compatibility with online course environments, flexibility for course-specific customization, and a cost-effective model that supported implementation across the program.

The chatbot, named Ava, was configured using the platform's widget editor. This editor enabled customization of both the chatbot's appearance and instructional behavior, including its visual presentation, personality, coaching style, and response patterns. Rather than functioning as a general-purpose question-answering tool, Ava was designed as a conversational learning agent that supported students through prompting, clarification, feedback, and guided reflection.




Figure 1. Setting Up the AI Chatbot

Case-Based Curriculum Design

The AI Ethics micro-lesson was designed as a brief, interactive learning experience that introduced students to responsible AI use in academic work. Rather than presenting AI ethics as a set of abstract rules, the lesson used three realistic case studies to help students recognize ethical issues they are likely to encounter when using AI tools for writing, revising, and

summarizing course-related materials.

The curriculum focused on five core principles: transparency, authorship, accuracy, academic integrity, critical evaluation, and responsible use of AI-generated content. Ava guided students through three scenarios: *The Shortcut Paper*, *The Invisible Editor*, and *The Biased Bot*. Each case addressed a different ethical challenge.

The *Shortcut Paper* focused on academic honesty and disclosure when students use AI to generate or support written work. The *Invisible Editor* addressed authorship, voice, and the risks of allowing AI to substantially rewrite a student's ideas. The *Biased Bot* emphasized the need to verify AI-generated summaries, sources, and claims because AI tools can omit important information, introduce bias, or generate inaccurate content.

These case studies were selected because they represent common and practical ways students are already using AI in educational settings. Rather than focusing on extreme examples of misconduct, the cases were grounded in everyday academic tasks: writing a paper, revising a draft, and summarizing a source. This made the lesson more relatable and allowed students to examine ethical decision-making in situations they could realistically face in their own coursework. The cases were also sequenced from more obvious to more nuanced ethical concerns. For example, in *The Shortcut Paper*, Sofia's behavior illustrates academic misconduct and nondisclosure; Miguel's case, in *The Invisible Editor*, introduces the more complex issue of how much AI assistance is acceptable before the work no longer reflects the student's own thinking; and Ananya's case in *The Biased Bot*, shifts attention from student behavior to the trustworthiness of AI output. These cases created a compact curriculum that addressed ethical AI use from multiple angles, while remaining manageable within a brief micro-lesson.

Base Prompt and Behavioral Guardrails

The chatbot's base prompt was designed to function as both an instructional script and a set of behavioral guardrails. Ava was instructed to teach the lesson through guided dialogue using case studies, open-ended questions, and progressive hints. The goal was to prevent Ava from simply supplying answers, while still providing enough support to help students make progress. Table 1 summarizes the major prompt design elements.

Table 1. Key Prompt Design Elements Used to Configure Ava

Prompt Design Element	Instructional Purpose	Example Prompt Language
Role definition	Establishes Ava as an instructional guide	"You are Ava, a friendly, engaging EdTech instructor chatbot..."

Case-based structure	Keeps the lesson focused on applied ethical reasoning	“Use three short case studies to teach core AI ethics concepts.”
Open-ended questioning	Requires learner participation and reflection	“After each case, ask open-ended questions.”
No-answer guardrail	Prevents direct answer-giving	“Do not give them the correct answer right away.”
Tiered hints	Preserves productive struggle, while offering support	“Offer a small hint... then give another nudge.”
Redirection	Keeps students focused on the case	“Guide them back to the case details when needed.”
Adaptive support	Adjusts pacing and difficulty	“Track performance and adjust tone, pace, and follow-up difficulty.”

A central feature of the prompt was tiered scaffolding. When a student’s response was incorrect, incomplete, or off-topic, Ava was instructed to first affirm the student’s participation, then provide a small hint, analogy, or follow-up question. If the student continued to struggle, Ava provided additional nudges that gradually narrowed the focus. This structure allowed Ava to support learners without removing the cognitive work required for understanding.

For example, in the case about Sofia using AI to write her paper, Ava asked: “What do you think Sofia should have done instead?” If students struggled, Ava used hints such as: “Think about honesty in academic work” and “What’s one way to give credit when using ideas, text, or tools created by someone, or something, else?” These prompts were intended to move students toward the concept of disclosure without immediately naming the answer.

Iterative Prompt Refinement

The base prompt for Ava was refined through four major iterations. Each iteration improved the chatbot’s ability to function as a guided instructional coach rather than a tool that simply provided answers.

In the first iteration, the prompt established Ava’s basic role, tone, and lesson sequence. Ava was introduced as a friendly EdTech instructor chatbot who would guide students through a five-minute AI Ethics micro-lesson using three short case studies. This version included the basic instructional flow: welcome students, present each case, ask reflective questions, provide feedback, summarize key takeaways, and invite students to complete the quiz.

In the second iteration, the prompt was strengthened to better prevent answer-giving. This version added clearer instructions for how Ava should respond when students provided incorrect, incomplete, or off-topic answers. Instead of giving the correct answer immediately, Ava was instructed to affirm the student's participation, provide a small hint or follow-up question, and continue narrowing the focus until the student arrived at a correct or mostly correct answer.

In the third iteration, the lesson was converted into a more complete scripted dialogue. Each case study was separated into its own message, and tiered hints were added for each question. This helped control the pacing of the lesson and ensured that Ava paused for student responses before moving forward.

In the fourth iteration, adaptive branching was added. Ava was instructed to track learner performance and adjust tone, pace, and follow-up difficulty based on how many first-attempt responses were correct. High-performing students received more challenging "what-if" extensions, mid-performing students received prompts that deepened their reasoning, and students needing more support received simpler, step-by-step guidance.

Initial testing showed that some interactions became circular when learners provided repeated incorrect or partially correct responses. The final refinements helped reduce this issue by strengthening redirection strategies, clarifying acceptable hints, and improving the chatbot's ability to move students toward the intended concept without directly providing the answer.

Assessment Development

The pre- and post-tests were designed to provide a brief, focused measure of students' understanding of the key concepts addressed in the AI Ethics micro-lesson. Because the lesson was intended to be completed in less than five minutes, the assessment also needed to be concise, while still addressing the essential learning outcomes. To accomplish this, we developed five multiple-option questions aligned with the three case studies and the lesson's major takeaways.

Each item included a question stem, multiple correct responses, and plausible distractors. This format was selected because ethical AI use often requires students to recognize several appropriate practices at once, such as disclosing AI assistance, preserving their own voice, verifying accuracy, and evaluating outputs for bias. The multiple-option format also increased the difficulty of the assessment because students had to select all correct responses, while avoiding incorrect ones.

To support content alignment, each item was developed by identifying the core concept from the micro-lesson, writing a scenario-based stem connected to that concept, creating two or more correct responses, and developing plausible distractors that reflected common misconceptions about AI use. The same five-question assessment was administered as both the pre-test and post-test to measure changes in student understanding before and after the micro-lesson. Table 2 presents the assessment blueprint for the AI Ethics Micro-Lesson.


Table 2. Assessment Blueprint for the AI Ethics Micro-Lesson

Item	Lesson Component	Core Concept Assessed	Distractor Focus
1	The Shortcut Paper	Transparency, disclosure, authorship	Avoiding detection, nondisclosure
2	The Invisible Editor	Voice, ownership, accuracy	Treating AI as a source, overreliance
3	The Biased Bot	Bias, hallucination, verification	Plagiarism avoidance, irrelevant concerns
4	Lesson takeaways	Transparency, voice, critical thinking	AI as primary author, neutrality assumption
5	Overall responsible use	Ethical partnership with AI	Larger model eliminates bias, unnecessary formalization

Putting It All Together

After the chatbot and assessments were developed, the micro-lesson was assembled as a web-based learning module using Google Sites. The lesson began with a 50-second introductory video in which Ava introduced the purpose of the micro-lesson and emphasized the importance of practicing ethical AI use in educational settings. The video, chatbot interaction, and quiz created a concise instructional sequence that included content presentation, multimedia, learner interaction, and assessment within a brief learning experience.

AI ETHICS IN EDUCATION



AI ETHICS IN EDUCATION: WHAT ARE YOUR RESPONSIBILITIES?


Welcome to this micro-lesson on AI Ethics in Education. AI tools like ChatGPT, Grammarly, and Quillbot can support writing, research, and productivity—but without clear ethical boundaries, misuse can blur the line between help and dishonesty.

In this lesson, you'll review 3 short case studies and learn **5 key tips** for ethically using AI in your academic work.

INSTRUCTIONS

After watching the video introduction above, start a conversation with the AI Ethics chatbot below to explore some of the key do's and don'ts of using AI responsibly in your academic work. Once you've completed the chat, you'll take a short quiz to check your understanding.

IMPORTANT: If the conversation seems to repeat or gets stuck, simply tell Ava to move on to the next section (for example, "go to the second case," "go to the third case," or "go to the summary") and continue from there.



AI Ethics in Education Micro-Lesson

Ava

Hi and welcome to a quick micro-lesson on AI Ethics in Education! I'm Ava, your friendly EdTech instructor. What's your name?

Type a response

TEST YOUR KNOWLEDGE

This short quiz will assess your understanding of key principles related to ethical AI use in academic settings. You'll encounter real-world scenarios and questions designed to test your ability to make responsible choices when using AI tools in your coursework and research. Take your time, think critically, and reflect on what it means to use AI with integrity.

Take the Quiz

Note: You will earn a **Certificate of Achievement** for scoring **80%** or above on this quiz!

Figure 2. The Complete AI Ethics Micro-Lesson

Methodology

Research Design

This pilot study used a quasi-experimental, pretest-posttest comparison design to examine learning gains associated with two versions of an AI Ethics micro-lesson. The first version was an AI-powered chatbot micro-lesson in which students interacted with Ava to review the lesson

content through guided case-based dialogue. The second version was a non-chatbot micro-lesson that presented the same core content in a plain-text digital format.

The study was conducted during the Fall 2025 semester across two seven-week modules. The AI-powered chatbot micro-lesson was implemented during the first week of Fall 2025 Module 1, and the non-chatbot micro-lesson was implemented during the first week of Fall 2025 Module 2. Because students were not randomly assigned to conditions, the study should be understood as a pilot comparison rather than a causal test of instructional quality.

Participants and Setting

Participants were graduate students enrolled in Educational Technology courses in a fully online master's program. All students enrolled in the participating courses during the two Fall 2025 modules completed the micro-lesson as part of their Week 1 course activities.

The final analytic sample included 92 students. Of these, 50 students completed the AI-powered chatbot micro-lesson, and 42 students completed the non-chatbot micro-lesson. Participant demographic information, including age, gender, teaching experience, and prior experience with AI tools, was not collected as part of this pilot study. The study was reviewed by institution's Institutional Review Board (IRB) and determined to be exempt.

Instructional Conditions

The instructional intervention consisted of two versions of a brief micro-lesson on AI Ethics in Education. Both versions addressed the same core concepts: ethical transparency, responsible disclosure of AI use, preservation of student voice and authorship, critical evaluation of AI-generated content, bias, hallucinations, and verification of sources.

The AI-powered version was delivered through Ava, who guided students through three short case studies. Ava was designed to ask open-ended questions, provide tiered hints, and guide students toward appropriate ethical conclusions without directly providing answers.

The non-chatbot version presented the same instructional content in a plain-text digital format. The content, key takeaways, and assessment alignment were kept as similar as possible across both versions to support comparison between instructional formats.

Procedure

During Week 1 of Fall 2025 Module 1, students accessed the AI-powered chatbot micro-lesson through the learning management system course schedule. Students first completed a pretest, then completed the AI-powered micro-lesson, and then completed the posttest.

During Week 1 of Fall 2025 Module 2, students accessed the non-chatbot version of the same micro-lesson through the learning management system course schedule. These students followed the same sequence: pretest, micro-lesson, and posttest.

Both versions were designed to be completed in approximately five minutes, excluding the time required to complete the pretest and posttest. The study was conducted as part of normal course activities. Students completed the learning activity online and asynchronously.

Assessment Instrument

A five-item multiple-option assessment was used as both the pretest and posttest to measure changes in students' understanding of ethical AI use following the micro-lesson. Items were aligned with the lesson concepts and included plausible distractors reflecting common misconceptions about AI-generated content, bias, and inappropriate paraphrasing. The assessment was scored using partial credit, with deductions for incorrect responses and possible scores ranging from 0 to 100. Because the instrument was developed for a brief micro-lesson, it emphasized alignment with the instructional content and feasibility rather than comprehensive measurement; future revisions will include expert review and item analysis to strengthen validity evidence.

Qualitative Data Source

Qualitative data were collected from student reflections posted to the EdTech Blog. Students responded to a prompt asking them to reflect on their experience with the AI-powered micro-lesson and describe whether and how the chatbot shaped their understanding, thinking process, or engagement with the content. Because some students submitted more than one relevant reflection, the qualitative dataset included 70 usable reflection entries from students who completed the AI-powered micro-lesson.

Data Analysis

Descriptive statistics were calculated for each instructional condition, including pretest means, posttest means, standard deviations, and gain scores. Paired-samples *t* tests were conducted separately for each group to determine whether students demonstrated significant improvement from pretest to posttest within each instructional condition.

An independent-samples *t* test was conducted to compare gain scores between the AI-powered chatbot micro-lesson and the non-chatbot micro-lesson. Because the groups differed in pretest performance, an analysis of covariance was conducted as the primary between-group analysis. In the ANCOVA, *posttest score* served as the dependent variable, *instructional condition* served as the independent variable, and *pretest score* was entered as the covariate. Because the ANCOVA was used to adjust for pretest differences, findings should be interpreted with attention to the assumptions associated with this analysis.

Statistical significance was evaluated using an alpha level of .05. Effect sizes were reported where appropriate, including Cohen's *dz* for paired-samples *t* tests, Cohen's *d* and Hedges' *g* for the independent-samples *t* test comparing gain scores, and partial eta squared (partial η^2) for the ANCOVA.

Student reflections were analyzed using thematic analysis. The researchers reviewed responses to identify recurring patterns related to engagement, guided reasoning, reflection, perceived usefulness, and limitations. Initial codes were grouped into broader themes, and themes were

reviewed and refined through discussion until consensus was reached. Representative quotes were selected to illustrate each theme.

Results

Quantitative Findings

Quantitative analyses were conducted to compare learning outcomes between the AI-chatbot micro-lesson and the non-chatbot micro-lesson. The primary outcome was posttest performance, with pretest performance used to account for students' initial knowledge. Gain scores were calculated by subtracting each student's pretest score from the posttest score.

Descriptive Statistics

As shown in Table 3, students in both groups demonstrated improvement from pretest to posttest. The AI-chatbot micro-lesson group had a lower pretest mean ($M = 73.33$, $SD = 20.36$) than the non-chatbot micro-lesson group ($M = 81.40$, $SD = 12.37$). The AI-chatbot group demonstrated a larger descriptive gain score ($M = 12.12$, $SD = 19.01$) than the non-chatbot group ($M = 8.25$, $SD = 13.24$).

Table 3. Descriptive Statistics by Instructional Condition

Instructional condition	N	Pre-test		Post-test		Gain Score	
		Mean	SD	Mean	SD	Mean	SD
AI-powered micro-lesson	50	73.33	20.36	85.45	14.67	12.12	19.01
Non-chatbot micro-lesson	42	81.40	12.37	89.65	9.12	8.25	13.24

Within-Group Learning Gains

Paired-samples t tests were conducted to determine whether each group demonstrated significant improvement from pretest to posttest. Results indicated that students in the AI-chatbot micro-lesson group improved significantly from pretest to posttest, $t(49) = 4.51$, $p < .001$. Students in the non-chatbot micro-lesson group also demonstrated significant improvement, $t(41) = 4.04$, $p < .001$. These results indicate that both instructional formats produced statistically significant immediate learning gains.

Table 4. Paired-Samples t Tests for Pretest to Posttest Gains

Instructional Condition	Pretest Mean	Posttest Mean	Mean Gain	t	df	p	Cohen's d_z
-------------------------	--------------	---------------	-----------	-----	------	-----	---------------

AI-chatbot micro-lesson	73.33	85.45	12.12	4.51	49	<.001	0.64
Non-chatbot micro-lesson	81.40	89.65	8.25	4.04	41	<.001	0.62

Comparison of Gain Scores

An independent-samples *t* test was conducted to compare gain scores between the two groups. Although the AI-chatbot group demonstrated a larger mean gain than the non-chatbot group, the difference was not statistically significant, $t(87.18) = 1.15, p = .255$. The effect size was small, Cohen's $d = 0.23$, Hedges' $g = 0.23$. Thus, while the AI-chatbot micro-lesson produced a larger descriptive gain score, the difference between groups was not statistically significant.

Table 5. Independent-Samples *t* Test Comparing Gain Scores

Comparison	Mean Difference	<i>t</i>	<i>df</i>	<i>p</i>	Cohen's <i>d</i>	Hedges' <i>g</i>
AI-chatbot vs. Non-AI chatbot	3.87	1.15	87.18	.255	0.23	0.23

ANCOVA

Because the two groups differed in pretest performance, an analysis of covariance was conducted with posttest score as the dependent variable, instructional condition as the independent variable, and pretest score as the covariate. Pretest score was a significant predictor of posttest performance, $F(1, 89) = 17.56, p < .001$, indicating that students' initial knowledge was related to their posttest scores.

After controlling for pretest scores, instructional condition did not have a statistically significant effect on posttest performance, $F(1, 89) = 0.55, p = .460$, partial $\eta^2 = .006$. Pretest score was a significant predictor of posttest performance, $F(1, 89) = 17.56, p < .001$, partial $\eta^2 = .165$.

Table 6. ANCOVA Results for Posttest Scores by Instructional Condition

Source	Sum of Squares	<i>df</i>	Mean Square	<i>F</i>	<i>p</i>	Partial η^2
Instructional Condition	72.10	1	72.10	0.55	.460	.006
Pretest	2299.38	1	2299.38	17.56	<.001	.165
Error	11656.09	89	130.97			

Table 7. Adjusted Posttest Means Controlling for Pretest Scores

Instructional Condition	Adjusted Posttest M
AI-chatbot micro-lesson	86.53
Non-chatbot micro-lesson	88.36

Qualitative Findings: Student Perceptions of the Learning Experience

Qualitative data were collected from student reflections posted to the EdTech Blog. Five themes emerged from the responses.

Theme 1: Guided Thinking Rather Than Answer-Giving. Many students noted that Ava did not simply provide answers. Instead, the chatbot asked questions, offered prompts, and encouraged them to reason through the ethical scenarios. One student explained that Ava “didn’t just give me answers right away,” but instead “guided me by asking questions and making me think through my responses.” Another student described the chatbot as a “thinking partner rather than an answer provider,” noting that it prompted reflection and required justification of responses.

Theme 2: Reflection and Productive Struggle. Students frequently reported that the chatbot made them slow down, explain their thinking, and reconsider incomplete responses. One student wrote that the experience “made learning harder in the moment, but in a good way” because it required them to “actually understand” the content rather than simply recognize a concept. Another noted that the chatbot “pushed me to articulate my reasoning rather than simply absorb information.”

Theme 3: Engagement and Interaction. Many students described the micro-lesson as more engaging than reading text, watching a video, or completing a static activity. One student stated that the lesson “felt less like completing a task and more like working through a problem.” Another wrote that the chatbot made the lesson “more interactive and reflective than a traditional lesson” because it responded to the learner’s input and prompted continued thinking.

Theme 4: Modeling Responsible AI Use. Several students connected the experience to broader questions about responsible AI use. One student observed that the chatbot modeled how AI can “support learning when it promotes thinking rather than passive response.” Another explained that Ava helped them see AI as “a tutor and not as an author” of their work. Similarly, one student noted that the experience showed how AI can be used “as a tool to guide learners’ thinking rather than supplying answers.”

Theme 5: Limitations of Chatbot Guidance. Although most responses were positive, a few students identified limitations. One student described Ava as “clunky and basic,” noting that if the chatbot “doesn’t get just what she is expecting she gets hung up.” Another student, writing

from the perspective of an experienced tutor, observed that Ava used guided reasoning, but lacked the “human touch and the nuanced judgment that a human tutor provides.”

Overall, the student reflections suggested that the AI-powered micro-lesson supported active engagement, reflection, and ethical reasoning. Students generally perceived Ava as a guide that helped them think through the cases rather than as a tool that supplied answers. The feedback also identified areas for refinement, particularly the need to improve redirection, reduce circular responses, and provide appropriate support without becoming overly scripted.

Limitations

Because this was a pilot study, the findings should be interpreted cautiously. Students were not randomly assigned to conditions; instead, instructional format was determined by module enrollment. The two groups also participated during different seven-week modules, which may have introduced differences unrelated to the micro-lesson format. As a result, instructional condition was confounded with module timing, limiting causal interpretation.

In addition, participants were graduate-level Educational Technology students who may have had stronger academic skills, higher motivation, and greater familiarity with AI tools and online learning than other student populations. Their generally high pretest and posttest performance may have limited the ability to detect larger differences between instructional formats. The assessment was also brief, consisting of only five items, which may have limited measurement sensitivity.

The study measured immediate learning gains only and did not include a delayed posttest to examine retention. The study also did not analyze chatbot transcripts, which could have provided more direct evidence of how students interacted with Ava during the lesson. Finally, qualitative reflections may have reflected the perspectives of students who were more engaged or more willing to comment on the experience.

Discussion and Future Directions

This pilot study examined the design and preliminary effectiveness of an AI-powered micro-lesson intended to guide students through ethical reasoning without providing direct answers. Both the AI-powered and non-chatbot versions produced significant immediate learning gains, suggesting that a brief, focused micro-lesson can support students’ understanding of AI ethics. However, the quantitative results did not show a statistically significant advantage for the AI-powered version. Although the AI-chatbot group demonstrated a larger descriptive gain score, the gain-score comparison was not significant, and instructional condition did not significantly predict posttest performance after controlling for pretest scores. These findings indicate that the chatbot format did not produce stronger immediate achievement outcomes in this pilot sample. The qualitative findings suggest that the chatbot shaped the learning experience in ways not fully captured by posttest scores. Students described Ava as interactive, conversational, and reflective. Many noted that the chatbot encouraged them to slow down, explain their reasoning, and think through ethical scenarios instead of simply receiving information. In this sense, the chatbot’s

perceived value appeared to be more process-oriented than achievement-based, with students emphasizing engagement, reflection, productive struggle, and reasoning through ethical scenarios.

These findings should be interpreted cautiously because this was a pilot study with limited evidence of comparative effectiveness. Rather than suggesting that AI chatbots improve learning outcomes by default, the results point to several design considerations that may be useful for future development and research. Ava was intentionally designed as a learning guide rather than an answer provider. The chatbot asked open-ended questions, offered hints, and redirected students to the case scenarios instead of supplying direct responses. In this study, that design helped position the chatbot as a scaffold for student thinking, although additional research is needed to determine whether this type of support contributes to longer-term learning, transfer, or performance in other instructional contexts.

The case-based structure also appeared important to the learning experience. The scenarios gave students realistic situations in which to apply principles such as transparency, authorship, accuracy, bias, and verification. Within a brief micro-lesson, these cases provided a focused context for students to consider how responsible AI use applies to academic decisions they may encounter in their own coursework. Similarly, the use of tiered hints provided a way to support students when their responses were incorrect or incomplete without immediately giving them the correct answer. These features may be useful to examine more systematically in future studies, particularly in relation to students' reasoning processes and ability to transfer ethical AI principles to new situations.

The findings also reinforce the importance of designing and testing AI-powered micro-lessons as complete instructional experiences. Ava's usefulness depended not only on the chatbot prompt but also on the surrounding lesson design, including the case scenarios, multimedia content, learner interaction, assessment, prompt language, tiered hints, pacing, and behavioral guardrails. Student feedback highlighted areas for refinement, particularly when chatbot interactions became too scripted, circular, vague, or limited in their ability to interpret unexpected student responses. These issues suggest that prompt testing should include expected, unexpected, incomplete, and incorrect student responses so that designers can identify where the chatbot may need clearer guidance, better question wording, or more flexible feedback patterns.

Building on this pilot study, future research should examine the AI-powered micro-lesson with larger samples and stronger comparative designs, including quasi-experimental or experimental studies with random assignment where possible. Future studies should also include delayed posttests, transfer tasks, or scenario-based assessments to examine whether students retain AI ethics concepts and apply them in new academic or professional situations. Analysis of chatbot transcripts could further clarify how students interact with Ava and which prompts, hints, and feedback patterns support ethical reasoning. Additional comparisons across AI-only, text-only, and AI-plus-peer-discussion conditions, as well as different learner populations, could help determine when and how chatbot guidance is most useful. In summary, this pilot study contributes a practical model for designing AI-powered micro-lessons that guide students through ethical reasoning rather than provide direct answers.

Although the chatbot version did not produce significantly stronger immediate achievement outcomes than the non-chatbot version, students perceived Ava as interactive, reflective, and supportive of guided reasoning. The study suggests that the value of AI-powered instruction depends less on the presence of a chatbot itself and more on how the chatbot is integrated into a focused, case-based, and carefully tested learning experience.

References

- Barry, B. A., Sharp, R. R., & McGowan, M. L. (2025). Teaching the bioethics of information technologies and artificial intelligence in healthcare: Case-based learning for identifying and addressing ethical issues. *International Journal of Ethics Education*, 10. <https://doi.org/10.1007/s40889-025-00218-3>
- Corbeil, J. R. & Corbeil, M. E. (2024). Developing AI chatbots to enhance student learning experiences in online courses. *Issues In Information Systems*, 25(4), 78-90. 10.48009/4_iis_2024_106.
- Kuhail, M. A., Alturki, N., Alamlawi, S., & Alhejori, K. (2023). Interacting with educational chatbots: A systematic review. *Education and Information Technologies*, 28, 973–1018. <https://doi.org/10.1007/s10639-022-11177-3>
- Kumar, J. A. (2021). Educational chatbots for project-based learning: Investigating learning outcomes for a team-based design course. *International Journal of Educational Technology in Higher Education*, 18, Article 65. <https://doi.org/10.1186/s41239-021-00302-w>
- Lee, Y. F., Hwang, G. J., & Chen, P. Y. (2025). Technology-based interactive guidance to promote learning performance and self-regulation: A chatbot-assisted self-regulated learning approach. *Educational Technology Research and Development*, 73, 2279–2304. <https://doi.org/10.1007/s11423-025-10478-x>
- Melanou, C., & Beege, M. (2026). Scaffolding generative AI as a tutor: A quasi-experimental study of learning outcomes and motivational, cognitive and metacognitive processes. *Education Sciences*, 16(4), 651. <https://doi.org/10.3390/educsci16040651>
- Monib, W. K., Qazi, A., & Apong, R. A. (2025). Microlearning beyond boundaries: A systematic review and a novel framework for improving learning outcomes. *Heliyon*, 11(2), Article e41413. <https://doi.org/10.1016/j.heliyon.2024.e41413>
- Song, Y., Huang, L., Zheng, L., Fan, M., & Liu, Z. (2025). Interactions with generative AI chatbots: Unveiling dialogic dynamics, students' perceptions, and practical competencies in creative problem-solving. *International Journal of Educational Technology in Higher Education*, 22, Article 12. <https://doi.org/10.1186/s41239-025-00508-2>

Sørhaug, J.O. (2026). Prompting for scaffolding: A thematic analysis of K-12 students' use of educational chatbots for writing support. *Journal of Writing Research*, 17(3), 453-481. <https://doi.org/10.17239/jowr-2026.17.03.04>

Stephenson, R., & Armstrong, C. (2026). *Student generative artificial intelligence survey 2026* (HEPI Report No. 199). Higher Education Policy Institute. <https://www.hepi.ac.uk/reports/student-generative-ai-survey-2026>

Yin, J., Goh, T. T., & Hu, Y. (2024). Interactions with educational chatbots: The impact of induced emotions and students' learning motivation. *International Journal of Educational Technology in Higher Education*, 21, Article 47. <https://doi.org/10.1186/s41239-024-00480-3>