



AI Guide

ONLINE COURSE CONTENT

MODULE 1: AI CHATBOTS IN EDUCATIONAL CONTEXTS

2025-1-LT01-KA220-HED-000354784

Artificial Intelligence for Green Understanding, Insights & Decision Empowerment



Co-funded by
the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the National Agency of Erasmus+ Programme in Lithuania – Education Exchanges Support Foundation (ŠMPF). Neither the European Union nor the ŠMPF can be held responsible for them.

Table of Contents

<i>AI chatbots in educational contexts</i>	2
A. General Information	3
B. Content for STEM HE Teachers	3
Highlights.....	3
Subtopic Content	3
Green Decision Context	6
Illustrative Classroom Scenarios	7
Real World Examples.....	7
Example Projects	7
Recommended Readings.....	7
Videos/Interactive Resources	7
References	8
C. Assessment and Practice Activities	8
Suggested 60-Minute Sequence	8



AI chatbots in educational contexts

A. General Information

Module No	Module 1
Subtopic	AI chatbots in educational contexts
AI Literacy Component	Effective prompting; understanding of chatbot capabilities and limitations
Decision Making Skills	Framing tasks clearly; judging how reliable the results are; solving problems through repeated refinement
Learning Objectives for STEM HE Teachers	Use prompting strategies and explain the capabilities and limitations of AI chatbots.
Learning Objectives for STEM Students	Use AI responses cautiously and justify the reliability of chatbot outputs
Suggested AI Tools	ChatGPT; Claude for Education; Gemini for Education; Perplexity AI; Poe
Estimated Duration	60 minutes

B. Content for STEM HE Teachers

Highlights

After this subtopic, teachers should emphasise:

- AI chatbots are a form of generative AI built on large language models
- Chatbots generate responses by predicting likely text, not by retrieving verified facts
- Effective use of chatbots depends on clear, iterative prompting
- Chatbot outputs must be critically verified, especially on environmental or scientific claims
- Chatbots meaningfully support, rather than replace, STEM and green-decision teaching when used with a Socratic, guided-questioning approach

Subtopic Content

Introduction

Building on an understanding of AI types, this subtopic focuses on the form of AI that teachers and students are most likely to interact with directly: conversational AI chatbots such as ChatGPT, Gemini, and Claude. These tools sit within the generative AI category introduced in the previous subtopic and are increasingly used in higher education for brainstorming, explanation, feedback, and content drafting, a trend documented in a growing body of research reviewed later in this subtopic (Kasneci et al., 2023). Because chatbots are so accessible and so fluent, they raise AI literacy and decision-making questions in a particularly immediate way: how they work, how to use them well, and how to judge when to trust what they say.

What Are AI Chatbots?

Large language models, the technology many chatbots are built on, can be described, drawing on IBM's definition, as AI systems that understand and generate human language by processing large amounts of text data in order to carry out a wide range of language-based tasks (IBM, n.d.). An AI chatbot, in turn, is a conversational interface built on top of such a model that allows a user to type a request in natural language and receive a natural-language response. Unlike earlier customer-service chatbots, which followed scripted decision trees with a fixed set of possible replies, modern AI chatbots can respond to an almost unlimited range of prompts, hold a multi-turn conversation, and adapt their tone or format to the user's instructions. This flexibility is what makes them useful across so many educational tasks, from explaining a concept in simpler terms to drafting a lesson plan outline, but it is also what makes their limitations less obvious than those of older, more constrained tools. This places AI chatbots squarely within the generative AI category, a family of systems defined by producing new outputs rather than retrieving or classifying existing information (Banh & Strobel, 2023).



How Large Language Models Generate Responses

AI chatbots are built on large language models trained on enormous quantities of text drawn from books, websites, and other written material. Training starts with a massive amount of data. That data can amount to billions or trillions of words from books, articles, websites, code, and other text sources (IBM, n.d). During training, the model learns statistical relationships between words and phrases. Essentially, given the words that came before, it learns which word is likely to come next. Much of this capability rests on the transformer architecture introduced in 2017, whose self-attention mechanism allows a model to weigh the relevance of every other word in a passage when generating each new word, rather than processing text strictly in sequence (Vaswani et al., 2017). When given a prompt, a base chatbot with no connected search, retrieval, or tool-use capability does not search a database or retrieve a stored, pre-written answer. Instead, it generates its response one piece at a time, selecting its next word, or token, by sampling from a probability distribution over likely continuations, shaped by patterns learned during training and by the specific wording of the prompt itself.

Watch: [“Introduction to Large Language Models” \(Google Cloud Tech\)](#) for a short explainer on the technology described above.

This generative, prediction-based process explains several everyday observations about chatbots. It explains why responses can sound confident and fluent while still being incomplete, outdated, or factually incorrect, a phenomenon often called “hallucination,” because the model is optimised to produce plausible-sounding text, not to verify facts against a trusted source. It also explains why chatbots can behave inconsistently: asking the same question twice, or with slightly different wording, can produce noticeably different answers, since each response is generated fresh rather than looked up.

The Rise of Conversational AI in Higher Education

The rapid adoption of AI chatbots in higher education has been driven partly by their low barrier to entry, no specialised training is required to type a question and receive a useful-sounding answer, and partly by the range of opportunities documented in the research literature. A widely cited 2023 review of large language models in education catalogued benefits including personalised support, automated feedback generation, and assistance with brainstorming and drafting, alongside risks such as the perpetuation of societal biases, threats to academic integrity, and an over-reliance on AI-generated explanations without independent verification (Kasneci et al., 2023). At the same time, this same accessibility has raised concerns about unequal access, since not all students have the same familiarity or comfort with these tools. STEM higher education teachers are therefore expected not only to use chatbots effectively themselves, but to explicitly teach students how to use them responsibly. A systematic review of empirical studies on AI chatbots in science education adds a further caveat: most existing research measures whether chatbots hold students' interest or build isolated skills, while rarely examining whether they help students see themselves as capable, recognised participants in scientific practice, a gap worth bearing in mind when judging a chatbot's educational value (Cheung, 2026).

Chatbots versus Search Engines

Students often assume a chatbot works like a search engine, retrieving and summarising existing web pages, but the underlying process is fundamentally different. A search engine returns links to existing documents that a human author wrote and that can, in principle, be traced back to their original source. A chatbot generates new text word by word based on learned patterns, and while that text is often accurate, it is not a citation or retrieval of any specific document, even when the chatbot describes it in those terms. Some chatbot products now combine a language model with a live search function, allowing them to cite real, retrievable sources for at least part of their answer, and teachers should help students distinguish between this search-grounded mode and a chatbot's default mode of generating an answer purely from patterns learned during training, since the reliability of the two can differ substantially.

Because these differences matter for how much a given answer can be trusted, it helps to name the underlying product categories explicitly rather than treating all “chatbots” as one thing. A base chatbot answers purely from patterns learned during training, with no access to the internet or any external document. A search-grounded chatbot adds a live web search step, so it can cite current, retrievable sources for at least part of its answer. A retrieval-augmented chatbot, such as the NotebookLM workflow described elsewhere in this module, is instead grounded in a specific set of documents the user supplies, and its



answers should be checked against those documents rather than treated as general knowledge. A tool-enabled or agentic chatbot can additionally call external functions, such as a calculator, code runner, or a specific database, to complete part of a task. Knowing which of these categories a given chatbot product falls into and being aware that a single product can offer more than one mode, is itself part of judging how much to trust its output.

Effective Prompting Strategies

Effective use of chatbots depends heavily on prompting: framing a request with clear context, a specific goal, and, where useful, an example of the desired output or format. A vague prompt such as “tell me about air pollution” will produce a generic, unfocused response, whereas a more specific prompt such as “explain how PM2.5 particulate pollution affects respiratory health in children, in three short paragraphs suitable for a first-year STEM course” gives the model the context it needs to produce a targeted, usable answer.

Because a first response is rarely perfect, refining a prompt based on the chatbot’s output, asking it to shorten an answer, add an example, adopt a different tone, or reconsider a specific claim, is itself a decision-making and problem-solving skill that mirrors good practice more broadly: state the goal, evaluate the first attempt, and iterate. Teachers can model this iterative process explicitly in class, showing students that a chatbot’s first answer is a starting draft to be interrogated and improved, not a final authority to be accepted as-is.

Practice: [Learn Prompting](#)’s interactive tutorials to try refining prompts hands-on.

Capabilities of AI Chatbots in STEM and Sustainability Teaching

Within STEM and green-decision teaching contexts, chatbots can be genuinely useful for tasks such as explaining a concept using multiple analogies until one resonates with a particular student, translating technical scientific language into plain terms, drafting first versions of case studies or discussion scenarios (such as the air-pollution scenario explored elsewhere in this module), and helping non-native speakers rephrase their ideas more clearly. Because their responses can be regenerated instantly and adjusted through follow-up prompts, chatbots are also well suited to exploratory brainstorming, generating multiple possible framings of an environmental problem for students to critique and compare, rather than presenting a single fixed answer.

Watch: [Meet Khanmigo: The Student Tutor AI Being Tested in School Districts](#) (60 Minutes, CBS) to see a Socratic tutoring chatbot deployed at scale.

Limitations: Hallucination, Bias, and Overconfidence

Equally important is understanding chatbot limitations. During what is commonly termed a hallucination, a model generates information that is false or misleading while sounding entirely plausible, and large language models can similarly reflect and amplify biases present in their training data, producing outputs that are unfair or offensive (IBM, n.d.). Outputs are probabilistic and should never be accepted uncritically, particularly on topics with real-world consequences such as environmental policy, health guidance, or scientific claims. Most consumer-facing chatbots have no readily visible mechanism for signalling how confident it actually is in a given statement; a fabricated statistic can be delivered in exactly the same fluent, assured tone as an accurate one. This overconfident tone, combined with the risk of hallucination described above, is the single most important limitation for teachers to convey to students: fluency is not the same as accuracy. Large-scale benchmarking of generative foundation models confirms that hallucination is not a rare glitch but a persistent, measurable property of current large language models, alongside related failure modes such as sycophancy, where a model shifts its stated answer to match a user’s opinion rather than the evidence (Huang et al., 2025).

Evaluating and Verifying Chatbot Outputs

Teachers need to model how to judge the reliability of a chatbot’s response, rather than treating fluency as a proxy for correctness. Practical verification habits include the following. Cross-check any specific fact, statistic, or citation against an independent, credible source before relying on it. Ask the chatbot to explain its reasoning or show its sources, while remembering that it may still generate a plausible-sounding but incorrect justification. Be especially cautious with numerical claims, dates, and named sources, which large language models are particularly prone to generating incorrectly. Building this verification habit is a direct extension of the decision-making skills introduced earlier in this module: framing the task clearly, judging how reliable a result is, and refining the approach through repeated iteration.



Data Protection and Institutional-Use Boundaries

Because chatbot providers may store, and in some cases use, conversation content to improve their underlying models, STEM higher education teachers and students should observe clear data-protection boundaries whenever they use these tools. Personal data, student records, unpublished research data, and confidential institutional documents should never be uploaded to a consumer AI account. Wherever an institutionally approved account or configuration is available, such as an enterprise or education-tier subscription with stronger data-handling guarantees, it should be used in preference to a personal consumer account, and any chatbot use involving student work or institutional data should follow the relevant institutional AI-use policy. These precautions reflect a broader pattern documented across generative AI applications: privacy exposure and cybersecurity risk are consistently identified as major concerns wherever such systems are deployed, education included, precisely because conversational input is rarely treated with the same caution as a formal data-sharing request (Uddin et al., 2026).

A Short History of Chatbot Development

Chatbots as a concept are not new. ELIZA, created in 1966, was among the first programs able to sustain a text conversation, using simple pattern-matching and scripted responses to simulate a Rogerian psychotherapist (Weizenbaum, 1966), and scripted customer-service chatbots following the same basic template became common on websites in the 2000s and 2010s. What changed dramatically in recent years is the underlying technology: the shift from scripted, rule-based dialogue to large language models built on the transformer architecture (Vaswani et al., 2017) produced a leap in fluency, flexibility, and general knowledge that earlier chatbots could not approach. This shift happened quickly enough that many educators, students, and institutions are still developing shared norms around appropriate use, which is precisely why explicit classroom conversations about prompting, verification, and academic integrity, rather than assuming students already know how to use these tools responsibly, are so important at this stage.

Watch: [The History of Chatbots: From ELIZA to ChatGPT](#) for the fuller backstory behind this shift.

Chatbots in the Broader AI Landscape

Placing chatbots back within the classification introduced in the previous subtopic reinforces how these ideas connect across the module. A chatbot is a generative AI system, meaning its core function is producing new content rather than forecasting a numeric outcome or sorting data into categories. This is why a chatbot is not the right tool for tasks like predicting next month's energy demand or classifying a water sample as safe or unsafe, even though it can fluently discuss those topics in general terms. Those tasks call for predictive or classification systems built and validated on the relevant data, not a general-purpose language model. Recognising a chatbot as one specific type of AI helps set realistic expectations. It has a specific set of strengths, including generating explanations, drafting text, and brainstorming ideas. It also has a specific set of blind spots, including verified factual precision, numerical accuracy, and up-to-date information. This allows teachers to set realistic expectations for what it can contribute to a STEM or green-decision lesson, and where a different type of tool, or a human expert, remains necessary.

Implications for Teachers and Students

For STEM higher education teachers, this means developing and explicitly teaching prompting strategies while openly discussing where chatbots are useful and where they fall short, rather than either banning their use outright or endorsing them uncritically. For students, it means learning to use chatbot responses cautiously, treating them as a starting point for further inquiry rather than a final answer, and being able to justify, in their own words, why a given output can or cannot be trusted. At an institutional level, universities are encouraged to pair the adoption of these tools with explicit staff and student training on their appropriate, ethical use, rather than leaving norms to develop informally (UNESCO, 2023). These skills, framing tasks clearly, judging reliability, and refining outputs through iteration, prepare both teachers and students for the broader landscape of AI tools explored in the next subtopic.

Green Decision Context

AI chatbots can support green decision-making teaching by:

- guiding students through structured reasoning about environmental trade-offs, such as comparing the lifecycle carbon costs of different energy sources



- helping students interpret regional emissions, energy, or climate datasets
- drafting first versions of environmental policy briefs or case studies for students to critique
- supporting multi-source, cited research on sustainability topics

However:

- a chatbot's fluent tone is not evidence of factual accuracy, especially on environmental statistics
- outputs should be checked against primary sources such as IPCC reports or peer-reviewed literature
- teachers should favour Socratic or guided-questioning configurations, where the tool asks questions rather than supplies conclusions, to build genuine student reasoning.

This guided-questioning approach echoes recommendations from the broader science education literature: a systematic review of environmental socioscientific issues found that argumentation and decision-making are the most consistently targeted skills in this field and recommends structured, dialogic teaching approaches, precisely the role a Socratic-mode chatbot is well suited to support rather than simply supplying conclusions (Kumar et al., 2024).

Illustrative Classroom Scenarios

The following two scenarios are illustrative examples adapted from the AIGUIDE WP2-A3 Best Practices Report to show how these tools could be used in a STEM higher education classroom. They are not documented case studies (AIGUIDE, 2026).

- **ChatGPT in a STEM HE Environmental Engineering Course:** student teams use ChatGPT's Study Mode to compare the lifecycle carbon costs of wind versus solar power for a given region, referencing IPCC data. Rather than handing over a ready-made answer, the tool's Socratic questioning, such as "What criteria matter most to the affected community?", builds evidence-based and ethical reasoning. (AIGUIDE, 2026).
- **Gemini for Education in an Environmental Science Course:** students upload government environmental impact reports into NotebookLM, which grounds its mind maps and study guides strictly in that content, then use Guided Learning and Gemini Canvas quizzes to test their understanding of renewable-energy trade-offs. (AIGUIDE, 2026).

Real World Examples

- [Khanmigo \(Khan Academy\) – AI Socratic Tutor for Personalised Learning at Scale](#): according to Khan Academy's own 2023–24 annual report, more than 221.2K individuals used Khanmigo that school year (Khan Academy, 2024). Khanmigo is deployed primarily in K–12 settings rather than higher education; it is included here as a transferable example of Socratic-tutoring design at scale, not as a verified higher-education implementation. The tool is built specifically to guide students toward their own problem-solving decisions rather than supplying direct answers, an instructive model for deploying chatbots responsibly at scale.

Example Projects

- [OREO – To chatbOt in higher Education or not to chatbOt?](#) (2023-2025): helps the academic community understand the potential and risks of AI chatbots in education through unit-based training and a board game, and builds an online learning community, closely mirroring this subtopic's own goals.
- [DialogEduShift – Transforming Higher Education Teaching and Evaluation Approaches in the Era of AI Chat Tools](#) (2023-2025): equips academic educators with the skills to responsibly integrate AI chat tools into teaching and assessment through a course, explainer videos, and classroom policy guidance.

Recommended Readings

- UNESCO (2023). [Guidance for Generative AI in Education and Research](#).
- Learn Prompting. [Prompt Engineering Guide](#) (free, open-source).

Videos/Interactive Resources

- [“Introduction to Large Language Models” – Google Cloud Tech](#): a short explainer on how large language models, the technology behind AI chatbots, actually work.



- [Learn Prompting](#) Interactive Prompting Tutorials: hands-on exercises for practicing and refining prompts.
- [The History of Chatbots: From ELIZA to ChatGPT](#).
- [Meet Khanmigo: The Student Tutor AI Being Tested in School Districts](#) – 60 Minutes (CBS): a televised look at the Socratic-tutoring chatbot referenced in the Real World Examples above.

References

- AIGUIDE. (2026). AIGUIDE WP2-A3 best practices report: AI tools for green decision-making in STEM higher education. Artificial Intelligence for Green Understanding, Insights & Decision Empowerment (2025-1-LT01-KA220-HED-000354784).
- Banh, L., & Strobel, G. (2023). Generative artificial intelligence. *Electronic Markets*, 33, Article 63. <https://doi.org/10.1007/s12525-023-00680-1>
- Cheung, K. K. C. (2026). Science Identity and AI Chatbots: A Systematic Review of Empirical Studies in Science Education. *Journal of Science Education and Technology*. <https://doi.org/10.1007/s10956-026-10295-8>
- Huang, Y., Gao, C., Wu, S., Wang, H., Wang, X., Zhou, Y., et al. (2025). On the Trustworthiness of Generative Foundation Models – Guideline, Assessment, and Perspective. <https://doi.org/10.48550/arXiv.2502.14296>
- IBM. (n.d.). [What are large language models \(LLMs\)?](#) IBM Think. Retrieved August 26, 2026.
- Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., et al. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Khan Academy. (2024). [Annual Report 2023–24: Khanmigo](#).
- Kumar, V., Choudhary, S. K., & Singh, R. (2024). Environmental socio-scientific issues as contexts in developing scientific literacy in science education: A systematic literature review. *Social Sciences & Humanities Open*, 9, 100765. <https://doi.org/10.1016/j.ssaho.2023.100765>
- Uddin, M., Arfeen, S. U., Alanazi, F., Hussain, S., Mazhar, T., & Rahman, M. A. (2026). A Critical Analysis of Generative AI: Challenges, Opportunities, and Future Research Directions. *Archives of Computational Methods in Engineering*, 33, 1763–1793. <https://doi.org/10.1007/s11831-025-10355-Z>
- UNESCO. (2023). [Guidance for Generative AI in Education and Research](#).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł. & Polosukhin, I. (2017). Attention Is All You Need. *Advances in Neural Information Processing Systems*, 30. <https://doi.org/10.48550/arXiv.1706.03762>
- Weizenbaum, J. (1966). ELIZA—A Computer Program for the Study of Natural Language Communication Between Man and Machine. *Communications of the ACM*, 9(1), 36–45. <https://doi.org/10.1145/365153.365168>

C. Assessment and Practice Activities

The following activities give teachers concrete, measurable ways to check whether STEM higher education teachers can apply prompting strategies and whether students can justify the reliability of chatbot outputs, the two learning objectives stated for this subtopic. They are designed to fit within the subtopic’s 60-minute estimated duration alongside the reading above. Activity (e) below is a dedicated teacher-performance task: teachers revise a weak prompt, justify their changes, and evaluate the resulting output using the Reliability Checklist.

Suggested 60-Minute Sequence

- 0–10 minutes: introduce chatbots and large language models using the “What Are AI Chatbots?” and “How Large Language Models Generate Responses” sections and the linked video.
- 10–20 minutes: guided reading and discussion of prompting strategies, capabilities and limitations.
- 20–35 minutes: Activity (a), the before-and-after prompt-improvement task.
- 35–50 minutes: Activity (b), the claim-verification exercise, scored with the reliability checklist in Activity (c).
- 50–60 minutes: formative quiz and self-reflection, Activity (d).
- Activity (e), the teacher prompt-revision and classroom implementation task, is completed separately



by teachers as professional preparation and is not part of the 60-minute in-class sequence.

(a) Before-and-After Prompt-Improvement Task

Give students a deliberately vague prompt related to this module's themes, for example "tell me about renewable energy", and have them record the chatbot's response. Using the prompting principles above, clear context, a specific goal, and a desired output format, students then rewrite the prompt and record the improved response. As a measurable outcome, students submit both prompts and responses side by side with two or three sentences explaining which specific changes improved the result, directly assessing the learning objective of applying prompting strategies effectively.

(b) Verifying a Chatbot Response

Provide students with a short chatbot-generated response on an environmental topic covered elsewhere in this module, such as the wind-versus-solar carbon comparison, containing a mix of accurate statements and at least one fabricated statistic, one misattributed source, and one overstated scientific claim planted by the teacher. Students identify which statements they trust, which they doubt, and why, then check each doubtful claim against an independent source such as an IPCC report or peer-reviewed literature. As a measurable outcome, students submit an annotated version of the response marking each claim as verified, corrected, or unverifiable, together with the source they used.

Example flawed response for direct classroom use: "Solar panels convert roughly 85% of sunlight into electricity, making them more efficient than wind turbines in almost all climates. According to a NASA (2019) study, replacing a single coal-fired power plant with solar panels prevents 40 million tonnes of CO₂ emissions every year. Because solar energy is now the cheapest and cleanest option everywhere, most environmental engineers agree that wind power will be phased out within the next decade."

Answer key: the 85% efficiency figure is fabricated (commercial solar panels typically convert roughly 15–22% of sunlight into electricity); the "NASA (2019)" attribution is a misattributed source (no such NASA study or figure is independently verifiable); and the claim that wind power will be phased out within a decade is an overstated, unsupported generalisation not reflected in current energy-transition literature.

(c) Reliability Checklist

Teachers can use or adapt the following short checklist, scored yes / no / unsure, to assess any chatbot response as part of Activity (b) or independently:

- Is each specific fact, statistic, or citation traceable to an independent, credible source?
- Does the response acknowledge uncertainty where the underlying evidence is genuinely uncertain, rather than stating estimates as fact?
- Are named sources, dates, and figures checked rather than assumed accurate?
- Would the response change materially if asked again or reworded, and if so, has that been accounted for?
- Are the cited sources relevant, current, and accurately represented, rather than fabricated or used out of context?
- Has the response been checked for bias or a one-sided framing of the environmental issue involved?

(d) Formative Quiz and Self-Reflection

To close the lesson, students complete a short formative quiz aligned with the subtopic's learning objectives, for example:

- In one sentence, explain why a chatbot's fluent tone is not evidence that its answer is accurate.
- Give one concrete change you could make to a vague prompt to get a more useful answer.
- Name one type of claim, numerical, source-based, or scientific, that should always be independently verified before use, and explain why.
- Reflect in two or three sentences: after this subtopic, what will you personally do differently the next time you use a chatbot for a STEM or environmental task?

Teachers can score the first three questions against a simple correct / partially correct / incorrect key and use the reflection question as a qualitative check that students can justify chatbot-output reliability in their own words, directly evidencing the student learning objective for this subtopic.



(e) Teacher Prompt-Revision and Classroom Implementation Task

This activity directly targets the teacher learning objective. Give each teacher a deliberately weak, vague prompt drawn from this subtopic's themes, for example "explain climate change." The teacher revises the prompt using the principles set out above, clear context, a specific goal, and a desired output format, and records the chatbot's response to the revised prompt. In two or three sentences, the teacher justifies which specific changes they made and why, then scores the resulting response against the Reliability Checklist in (c). Finally, the teacher writes a short classroom implementation note (three to five sentences) explaining how they would introduce this revised-prompting approach to their own students, including one anticipated obstacle and how they would address it.

As a measurable outcome, the submitted prompt revision, completed checklist, and implementation note together provide direct evidence that the teacher can apply prompting strategies and evaluate chatbot-output reliability, the two learning objectives stated for this subtopic.

