

STIAEM

Slow Tech Innovation

AI, Ethics and
New Teaching Methods

Few-Shot prompting strategies: using pre-solved exemplar models for output stability







Co-funded by the European Union

2025-1-IT02-KA220-SCH-000365213

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 European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them... Project reference code: 2025-1-IT02-KA220-SCH-000365213

STIAEM

Slow Tech Innovation
AI, Ethics and
New Teaching Methods

Co-funded by
the European Union

2025-1-IT02-KA220-SCH-000365213

eduOpen

CC BY-NC-SA 4.0 International

This license requires that reusers give credit to the creator. It allows reusers to distribute, remix, adapt, and build upon the material in any medium or format, for noncommercial purposes only. If others modify or adapt the material, they must license the modified material under identical terms.

Università di Foggia
Studi Umanistici, Lettere, Beni Culturali,

OER 4.2.1: Few-Shot prompting strategies: using pre-solved exemplar models for output stability
Reference Module: Module 4 – Prompt Engineering for Inclusive Education
Related Video-Pill: Video 4.2: Advanced Prompting & Chain-of-Thought
Authoring partner: Rhodes Evening High School - Greece
Topic / Thematic Focus: Few-shot prompting, pre-solved exemplars and output stability

Introduction and Learning Objectives

Generative Artificial Intelligence systems are highly responsive to the way tasks are presented. Few-shot prompting addresses the limitations of vague or ambiguous instructions by providing a language model with a small number of pre-solved examples, or exemplars, before presenting a new task. These examples demonstrate the expected relationship between an input and an appropriate output, helping the system recognise patterns and produce more consistent, structured and relevant responses.

The resource is mapped to the European Framework for the Digital Competence of Educators (DigCompEdu), with specific reference to Area 2: Digital Resources (selecting digital resources 2.1, modifying and creating digital resources 2.2), Area 3: Teaching and Learning (teaching 3.1, guidance 3.2) and Area 6: Facilitating Learners' Digital Competence (information and media literacy 6.1, digital content creation 6.3).

Upon completion of this resource, teachers will be able to:

- Understand the principles of few-shot prompting and the role of exemplars in shaping AI-generated outputs.
- Design and apply few-shot prompts for specific educational tasks.
- Select and adapt effective examples to improve output consistency and relevance.
- Evaluate AI responses for accuracy, bias, relevance and over-standardisation.
- Use AI as pedagogical scaffolding while safeguarding learner agency, cognitive effort and teacher oversight.

1. From Instructions to Exemplars: Understanding Few-Shot Prompting

Large language models can perform new tasks without being specifically retrained for each individual task. Instead, they can use the context provided in a prompt to infer what kind of response is expected. Brown et al. (2020) demonstrated that large language models could perform a wide range of tasks through zero-shot, one-shot and few-shot prompting, making the examples contained within the prompt an important source of contextual guidance.

A zero-shot prompt provides an instruction without an example. A one-shot prompt includes a single example demonstrating the relationship between the task and the desired output. A few-shot prompt includes several examples before the final task. These demonstrations provide additional contextual information and can help the model infer a more specific response pattern (Brown et al., 2020; Min et al., 2022).

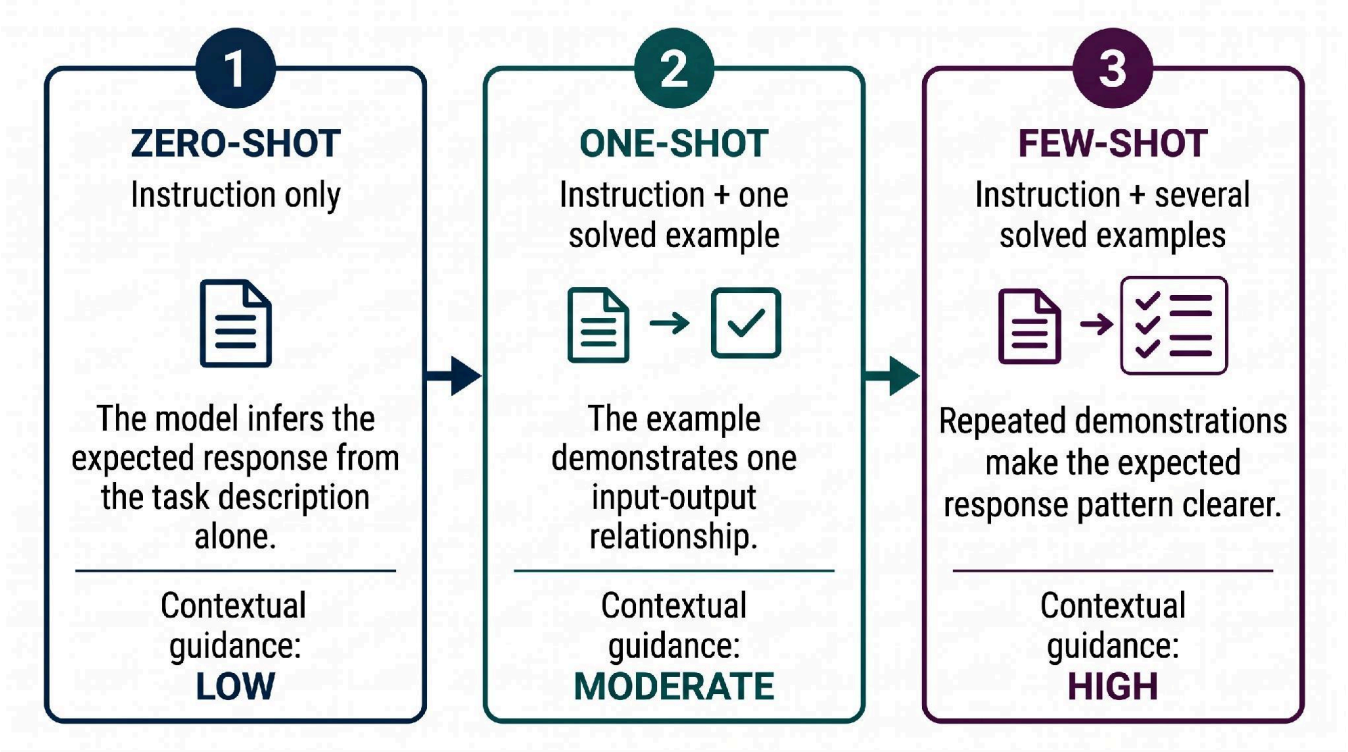


Figure 1. The progression from zero-shot to few-shot prompting. As relevant demonstrations are added to a prompt, the AI system receives a more explicit contextual pattern for interpreting the expected response.

The essential feature of few-shot prompting is therefore not simply the inclusion of more text. The examples demonstrate what should be recognised as an appropriate relationship between an input and an output. If a teacher wishes an AI system to provide feedback using a particular structure, for instance, pre-solved exemplars can show the model how previous student responses were analysed and how feedback was formulated.

For educators, this represents an important connection with DigCompEdu Area 2: Digital Resources. The examples used in a prompt are not neutral pieces of information. They become deliberately selected digital resources that shape the AI interaction. Teachers must therefore consider their accuracy, relevance and appropriateness before including them in an AI-supported task (Redecker, 2017).

2. Pre-Solved Exemplars and Output Stability

Pre-solved exemplars can help reduce ambiguity by making the expected form of a response visible. In a few-shot prompt, the AI system can identify recurring features in the demonstrations, such as the structure of an explanation, the categories used for classification, the tone of feedback or the level of detail expected in the final response (Brown et al., 2020; Min et al., 2022).

For example, a teacher may want an AI system to create structured classroom activities. Instead of writing only 'Create an activity about this historical topic', the teacher may provide two or three examples using the structure Topic → Learning Objective → Student Task → AI-Supported Activity → Reflection. The AI is then asked to create a new activity using the same pattern. The exemplars reduce uncertainty because the expected format is demonstrated rather than merely described.

However, the concept of output stability must be interpreted critically. A response can be stable in the sense that it repeatedly follows the expected pattern without necessarily being factually correct, pedagogically appropriate or ethically acceptable. This distinction is particularly important in education.

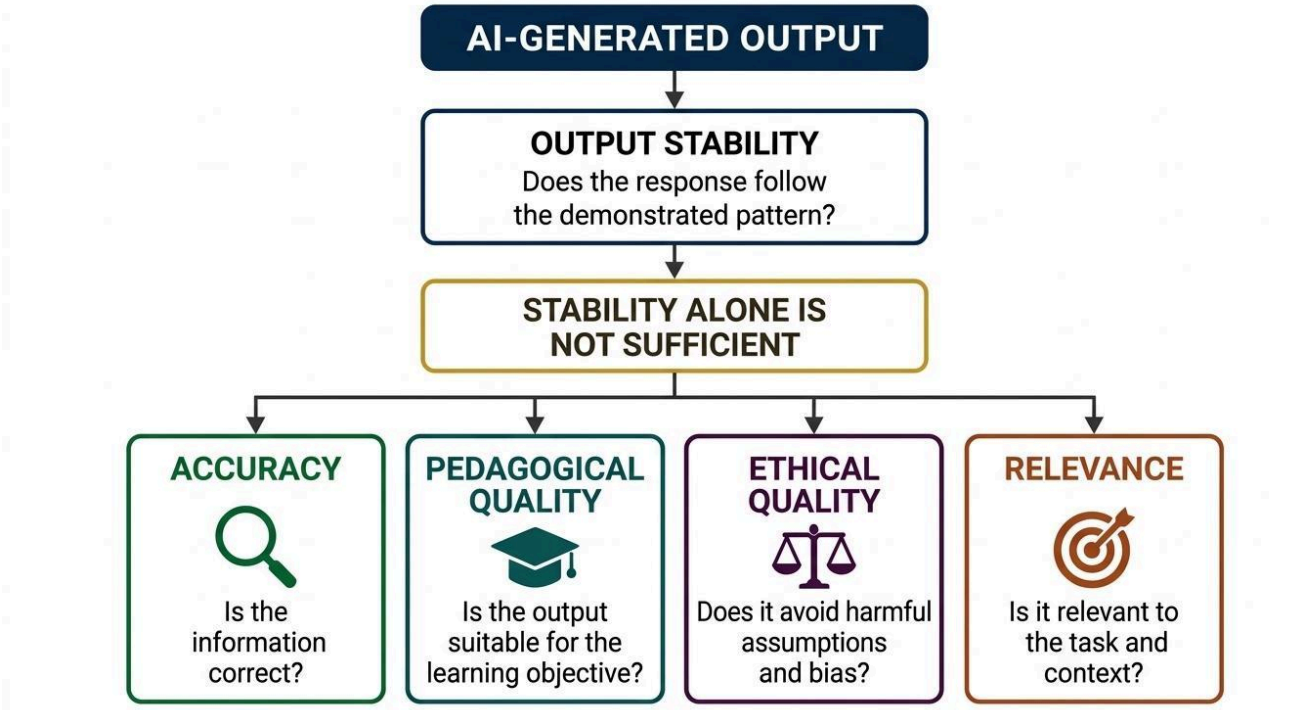


Figure 2. Output stability in relation to accuracy, pedagogical quality, ethical quality and relevance. A consistent response pattern does not guarantee that the generated content is correct, educationally appropriate, ethically responsible or relevant to the task.

3. Few-Shot Prompting and the Pedagogical Logic of Worked Examples

The use of pre-solved exemplars has strong connections with established pedagogical approaches. In education, worked examples provide learners with a completed solution that can be studied before they attempt similar problems independently. Research on example-based learning has shown that carefully designed examples can support the development of knowledge structures and reduce unnecessary search processes during initial learning (Sweller & Cooper, 1985; Renkl, 2014).

Few-shot prompting follows a related logic. Instead of presenting examples directly to a learner, the examples are also presented to an AI system as contextual demonstrations. Nevertheless, the pedagogical importance of the example remains central. The quality and representativeness of the model influence the subsequent process.

This creates an important opportunity for teachers. Rather than interacting with AI through repeated trial and error, educators can use examples to communicate a clear pedagogical intention. Few-shot prompting can therefore transform prompting into a more deliberate process of instructional design.

AI-supported examples should be embedded within activities that preserve meaningful cognitive effort. A useful learning sequence can therefore be: human analysis of examples; AI generation; comparison; error detection; revision; and independent application. This

sequence is closely connected to DigCompEdu Area 3: Teaching and Learning, because the teacher's task is not simply to introduce an AI tool but to design and manage a pedagogical process around it (Redecker, 2017).

It is also consistent with the current Universal Design for Learning perspective, which identifies learner agency as purposeful, reflective, resourceful, authentic and strategic. Examples and scaffolds should therefore support learners in developing their capacity to act rather than creating permanent dependence on external support (CAST, 2024).

4. Designing Effective Pre-Solved Exemplars

The effectiveness of few-shot prompting depends less on the number of examples than on their quality. Research on demonstrations used in in-context learning indicates that the characteristics and selection of examples can significantly influence model performance (Min et al., 2022). Teachers should therefore treat exemplars as carefully selected pedagogical resources rather than interchangeable pieces of prompt text.

An effective exemplar should first be clear. It should make the expected relationship between the task and the response understandable. Secondly, it should be relevant to the new task. Examples that are too distant from the intended activity may encourage the AI to identify superficial similarities rather than the desired pattern.

The exemplar should also be accurate and pedagogically legitimate. If an incorrect or misleading example is included in the prompt, the AI may reproduce the same weakness with greater consistency. Finally, the educator should be able to explain why the example was selected and what function it performs. This is particularly relevant to DigCompEdu Area 2: Digital Resources, which requires educators to select, create and manage digital resources in ways that are appropriate to educational objectives (Redecker, 2017).

5. Few-Shot Prompting Through the STlAEM Pedagogical Framework

The STIAEM pedagogical framework can be interpreted through the interconnected dimensions of Agency, Reflection, Transformation and Ethics (ARTE). Few-shot prompting can be understood through these four dimensions.

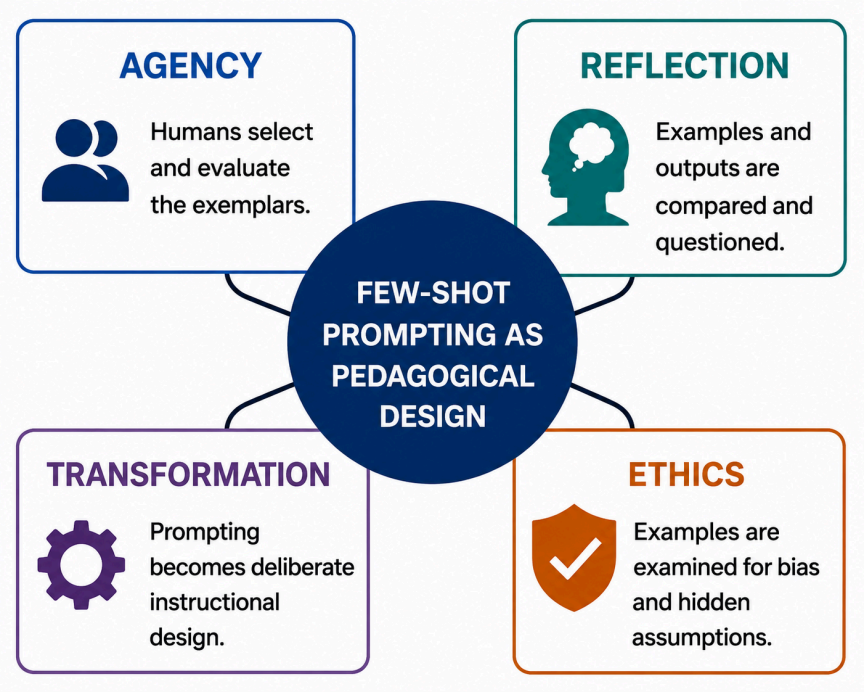


Figure 3. Few-shot prompting interpreted through the ARTE framework. The use of exemplars becomes a pedagogical practice when human agency, critical reflection, instructional transformation and ethical evaluation remain connected.

Agency

Agency requires teachers and learners to remain active decision-makers. In few-shot prompting, humans select the examples, define the task and evaluate the result. The AI may generate a response, but it does not independently determine what should count as a pedagogically appropriate example or outcome.

Reflection

Reflection requires users to examine the relationship between the examples and the generated output. Teachers and learners can consider which features of the examples were reproduced, which features were ignored, whether the AI introduced unexpected assumptions and whether the response remains accurate and educationally appropriate.

Transformation

Few-shot prompting can transform AI interaction from random experimentation into deliberate design. Teachers can create and refine reusable exemplar patterns that make their expectations more explicit and transparent.

Ethics

Ethical reflection requires attention to the examples themselves. Every exemplar may contain assumptions about language, culture, ability, gender or acceptable forms of expression. If examples are treated as neutral, their biases may be reproduced and normalised through subsequent AI outputs.

6. Cognitive and Ethical Risks of Example-Based Prompting

Few-shot prompting can create an impression of reliability because the outputs become more predictable. However, predictability can generate false confidence. A response may look trustworthy because it resembles the examples while still containing factual inaccuracies or inappropriate assumptions.

The first major risk is cognitive offloading. If learners repeatedly delegate both the production and evaluation of answers to AI, the technology may reduce the amount of independent processing required from the learner. For this reason, educators should consider carefully which stages of an activity benefit from AI support and which stages require independent cognitive effort.

The second risk is imitation without understanding. This is particularly important when stable output structures are mistaken for evidence of conceptual learning. Learners may reproduce the visible form of a response without understanding the principles that produced it. A third risk concerns bias replication. Examples may contain cultural assumptions, stereotypes or narrow definitions of successful performance. Because few-shot prompting encourages the model to follow the demonstrated pattern, poorly selected examples may cause these limitations to be repeated across subsequent outputs.

A fourth risk is over-standardisation. Consistent formats can be useful for classification, feedback or structured learning materials. However, not every educational activity should lead to a uniform response. Tasks involving creativity, divergent thinking or personal expression may require more open-ended forms of AI interaction.

These risks reinforce the importance of DigCompEdu Area 6, which includes supporting learners in using digital technologies critically and responsibly. They also connect with the STIAEM concern that technological efficiency should not become more important than human judgment and cognitive development.

7. A Responsible Few-Shot Prompting Cycle for Education

Few-shot prompting can be integrated into educational practice through a structured cycle. The cycle positions AI generation within a broader process of human planning, careful exemplar selection, critical evaluation, revision and independent cognitive consolidation.

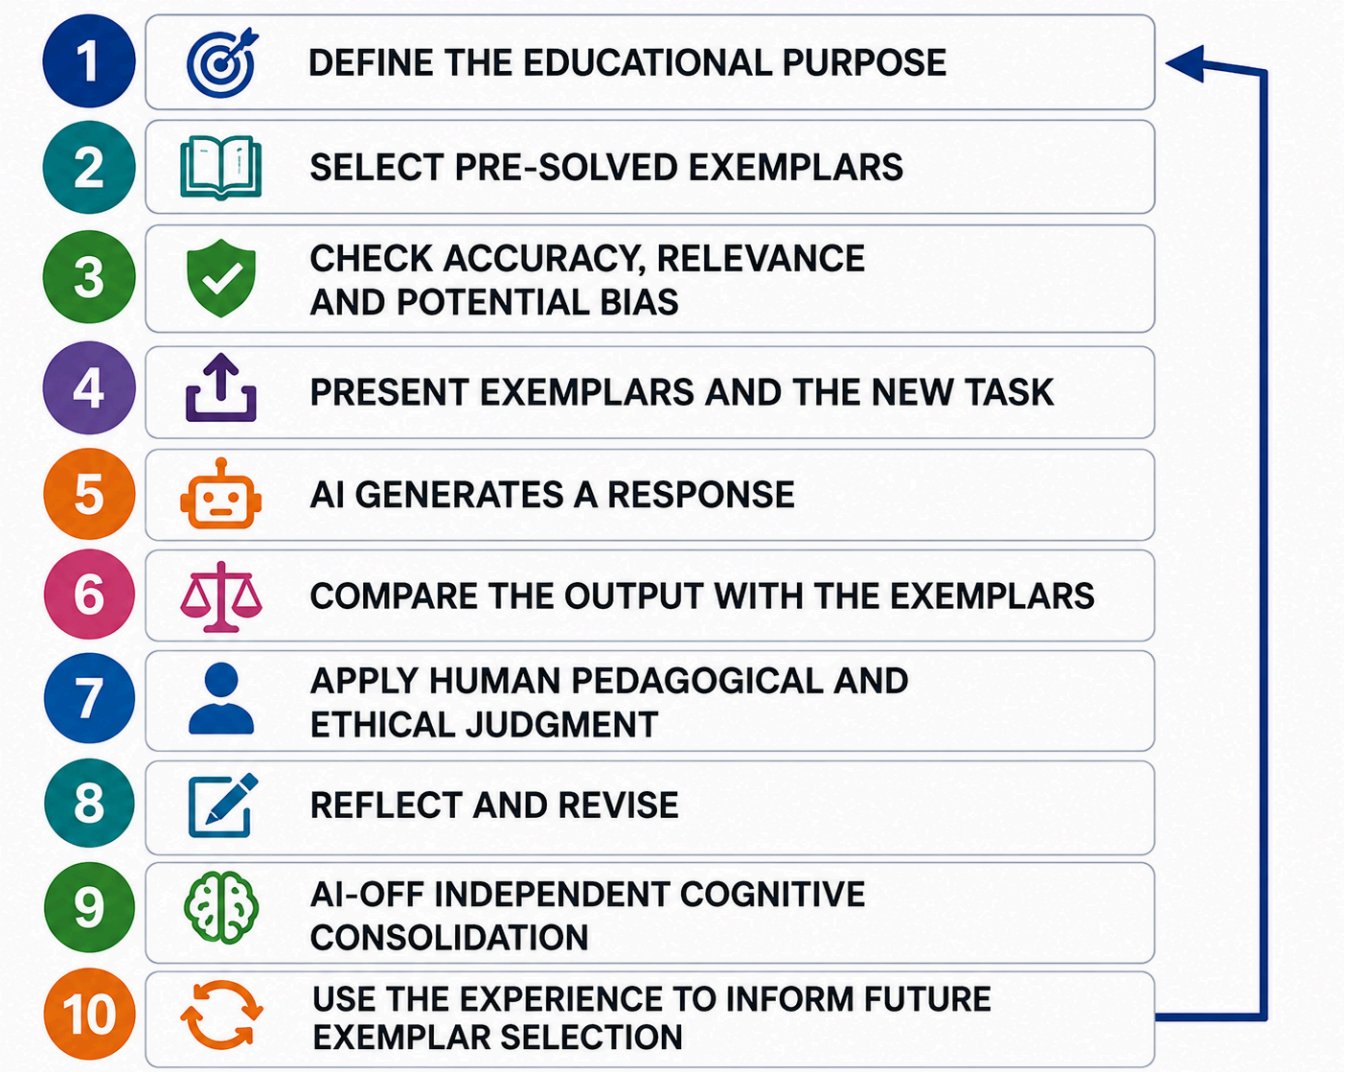


Figure 4. The responsible few-shot prompting cycle for educational practice. The cycle positions AI generation within a broader process of human planning, exemplar selection, critical evaluation, revision and independent cognitive consolidation.

1. Define the educational purpose
The teacher first identifies what the AI is expected to support. The purpose may involve feedback, classification, explanation, differentiation or another educational function.
2. Select pre-solved exemplars
A limited number of representative examples are selected to demonstrate the intended structure or quality.
3. Audit the exemplars
The teacher checks the examples for accuracy, relevance, bias, stereotypes and hidden assumptions.
4. Present the examples and the new task
The examples are clearly separated from the new task so that the AI receives a structured context.
5. Generate the output
The AI system produces a response using the contextual pattern established by the exemplars.

6. Compare rather than merely accept
Teachers or learners compare the generated output with the original examples and identify similarities, deviations and possible errors.
7. Apply human pedagogical judgment
The output is evaluated for factual accuracy, educational relevance and ethical appropriateness. Educational uses of generative AI should remain subject to human-centred ethical and pedagogical validation (UNESCO, 2023).
8. Reflect and revise
The teacher or learner considers what should be retained, corrected, rejected or improved before the output is used.
9. Include an AI-off moment
Where learning is the primary objective, learners should also complete a stage without AI support. They may independently explain, solve, revise or justify their reasoning.
10. Use the experience to inform future exemplar selection
Reflection on the quality of the generated output informs the future selection and design of exemplars, allowing the process to become more deliberate over time.

8. Conclusion

Few-shot prompting demonstrates how pre-solved exemplars can influence the behaviour of generative AI systems and contribute to greater consistency across similar tasks. By providing examples, teachers can make expectations more explicit and reduce ambiguity in AI-assisted interactions. Research on in-context learning demonstrates the importance of demonstrations in shaping model behaviour, while educational research on worked examples provides a useful pedagogical perspective on the value of well-designed models and exemplars (Brown et al., 2020; Min et al., 2022; Renkl, 2014).

However, consistency must never be confused with correctness. A stable AI response may still contain factual errors, reproduce biased assumptions or be unsuitable for a particular learning context. The quality of the output is therefore closely connected to the quality of the examples selected by the human user.

Within the STIAEM framework, few-shot prompting should be understood as a form of deliberate technological mediation. The examples are not neutral technical instructions; they are pedagogical choices. Teachers remain responsible for selecting exemplars, reflecting on their implications, transforming them into meaningful learning designs and evaluating their ethical consequences.

The ultimate educational question is therefore not simply whether AI can imitate an example successfully. It is whether the interaction helps teachers and learners develop greater understanding, critical awareness and capacity for independent judgment.

Conclusions and Metacognitive Reflection Prompts

Few-shot prompting shows that examples can become powerful mechanisms for structuring AI behaviour. Nevertheless, educational technology should not be evaluated solely according to its ability to generate predictable outputs. The central question is whether consistency supports genuine learning or merely creates more efficient forms of imitation.

The STIAEM approach requires educators to remain active mediators of the interaction. Pre-solved exemplars should therefore function as visible and critically examinable forms of scaffolding, while teacher and learner agency remain central.

- In which educational tasks would stable AI output genuinely support learning, and in which tasks could it unnecessarily standardise students' thinking or expression?
- How can a teacher determine whether an exemplar represents genuine quality or merely reproduces a familiar pattern?
- What types of bias, exclusion or hidden assumptions might be embedded in the examples selected for a few-shot prompt?
- How could students compare AI-generated responses with pre-solved exemplars in order to identify errors, unexpected deviations or limitations?
- At which point in an AI-supported learning activity should students deliberately work without AI in order to demonstrate independent understanding and preserve cognitive effort?
- Which DigCompEdu competences would you develop when using few-shot prompting as part of your own teaching practice, and how would you make this development visible to your learners?

Technical Metadata Sheet

Metadata Field	Specification / Parameter Value
Identifier	OER 4.2.1
Title	Few-Shot Prompting Strategies: Using Pre-Solved Exemplar Models for Output Stability
Creator	Rhodes Evening High School - Greece
Description	The resource presents few-shot prompting: the technique of giving the machine two or three pre-solved examples before the real task in order to stabilise output and align it with the teacher's expectations. It explains how to choose and order exemplars and when zero-shot or one-shot is enough. The critical focus is that the stability of form does not guarantee the correctness of content.
Keywords	few-shot prompting; exemplars; output stability; zero-shot; one-shot; prompt engineering; in-context learning; verification
Language	EN (English)
Affiliation	Rhodes Evening High School - Greece
Contributor	STIAEM Transnational Scientific Board (peer review, Rhodes validation protocol)
Media Format	Text; Infographics
Genre	Training Manual; Educational Guide
URL	
Citation	Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., ... Amodei, D. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33, 1877–1901. CAST. (2024). Universal Design for Learning Guidelines version 3.0. https://udlguidelines.cast.org/ Min, S., Lyu, X., Holtzman, A., Artetxe, M., Lewis, M., Hajishirzi, H., & Zettlemoyer, L. (2022). Rethinking the role of demonstrations: What makes in-context learning work? In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (pp. 11048–11064). Association for Computational Linguistics. https://aclanthology.org/2022.emnlp-main.759/ Redecker, C. (2017). European framework for the digital competence of educators: DigCompEdu. Publications Office of the European Union. https://doi.org/10.2760/159770 Renkl, A. (2014). Toward an instructionally oriented theory of example-based learning. Cognitive Science, 38(1), 1–37. https://doi.org/10.1111/cogs.12086 Sweller, J., & Cooper, G. A. (1985). The use of worked examples as a substitute for problem solving in learning algebra. Cognition and Instruction, 2(1), 59–89. https://doi.org/10.1207/s1532690xci0201_3 UNESCO. (2023). Guidance for generative AI in education and research. UNESCO.
Educational Level	Secondary School
Associated Module / MOOC Video Lesson	Module 4 - Prompt Engineering for Inclusive Education / Video 4.2 - Advanced Prompting and Chain-of-Thought
Learning Objectives	explain what distinguishes few-shot prompting from a simple request of a task, and why it stabilises the output; build two or three mutually coherent worked examples and use them as a model for the new task; evaluate critically the limits of the technique, knowing that stability of form does not guarantee correctness of content.
Rights	Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0)
Date Available	14/09/2026
Date Modified	

References

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., ... Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems, 33*, 1877–1901.

CAST. (2024). *Universal Design for Learning Guidelines version 3.0.* https://udlguidelines.cast.org/

Min, S., Lyu, X., Holtzman, A., Artetxe, M., Lewis, M., Hajishirzi, H., & Zettlemoyer, L. (2022). Rethinking the role of demonstrations: What makes in-context learning work? In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing* (pp. 11048–11064). Association for Computational Linguistics. https://aclanthology.org/2022.emnlp-main.759/

Redecker, C. (2017). *European framework for the digital competence of educators: DigCompEdu.* Publications Office of the European Union. https://doi.org/10.2760/159770

Renkl, A. (2014). Toward an instructionally oriented theory of example-based learning. *Cognitive Science, 38*(1), 1–37. https://doi.org/10.1111/cogs.12086

Sweller, J., & Cooper, G. A. (1985). The use of worked examples as a substitute for problem solving in learning algebra. *Cognition and Instruction, 2*(1), 59–89. https://doi.org/10.1207/s1532690xci0201_3

UNESCO. (2023). *Guidance for generative AI in education and research.* UNESCO.