

PAPER • OPEN ACCESS

## Creating a customisable Socratic AI physics tutor

To cite this article: Eugenio Tufino and Bor Gregorcic 2025 *Phys. Educ.* **60** 065037

View the [article online](#) for updates and enhancements.

### You may also like

- [The correlation between physics problem-solving skill and metacognitive ability from collaboration of socratic dialogue-modeling instruction implementation](#)  
N A Lestari and U A Deta
- [The Concepts of Space and Time: Their Structure and their Development](#)  
S Zienau
- [ChatGPT as a tool for honing teachers' Socratic dialogue skills](#)  
Bor Gregorcic, Giulia Polverini and Andreja Sarlah

# Creating a customisable Socratic AI physics tutor

Eugenio Tufino<sup>1,\*</sup>  and Bor Gregorcic<sup>2,\*</sup> 

<sup>1</sup> Department of Physics, University of Trento, Trento, Italy

<sup>2</sup> Department of Physics and Astronomy, Uppsala University, Uppsala, Sweden

E-mail: [eugenio.tufino@unitn.it](mailto:eugenio.tufino@unitn.it) and [bor.gregorcic@physics.uu.se](mailto:bor.gregorcic@physics.uu.se)



CrossMark

## Abstract

This paper explores role engineering as an effective paradigm for customizing large language models into specialized AI tutors for physics education. We demonstrate this methodology by designing a Socratic physics problem-solving tutor using Google's Gemini Gems feature, defining its pedagogical behaviour through a detailed 'script' that specifies its role and persona. We present two illustrative use cases: the first demonstrates the Gem's multimodal ability to analyse a student's hand-drawn force diagram and apply notational rules from a 'Knowledge' file; the second showcases its capacity to guide conceptual reasoning in electromagnetism using its pre-trained knowledge without using specific documents provided by the instructor. Our findings show that the 'role-engineered' Gem facilitates a Socratic dialogue, in stark contrast to a standard Gemini model, which tends to immediately provide direct solutions. We conclude that role engineering is a pivotal and accessible method for educators to transform a general-purpose 'solution provider' into a reliable pedagogical tutor capable of engaging students in an active reflection process. This approach offers a powerful tool for both instructors and students, while also highlighting the importance of addressing the technology's inherent limitations, such as the potential for occasional inaccuracies.

\* Authors to whom any correspondence should be addressed.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Keywords: AI physics tutor, large multimodal models, chatbots, Socratic approach, physics problem solving, AI educational tools

Supplementary material for this article is available [online](#)

---

## 1. Introduction

Recent large language models (LLMs) have demonstrated significant advancements in their performance in many domains and their capacity to process broader contexts; however, they remain prone to generating inaccurate or fabricated information—so-called ‘hallucinations’. This tendency stems from the probabilistic nature of their underlying next-word prediction algorithms [1]. One common strategy to mitigate this issue and tailor LLM responses to specific educational contexts is to provide them with curated information and precise instructions. For example, the LEAP platform [2] allows teachers to design tasks within a controlled environment by providing reference texts, tailored instructions, and verified answers, which effectively shape the context for the LLM’s responses.

A more generalized and advanced approach to improve reliability is retrieval-augmented generation (RAG), which grounds LLM responses in an external, verifiable knowledge base [3]. This technique typically involves converting source documents directly into vector embeddings, to enable the efficient retrieval of text passages relevant to a user’s query. This retrieved information is then provided as context to the LLM, guiding it to generate a factually grounded response.

Notable examples of RAG applications in physics education include the Ethel project [5] and Google’s NotebookLM [4]. The latter has been previously used by the author (E.T.) to illustrate the application of NotebookLM in creating AI tutors for physics problem-solving [6]. LLMs can also be directly customized for education through methods such as meticulous prompt engineering or dedicated platform features for creating specialized assistants. Discussing the broader educational implications of such large multimodal foundation models (LMFMs), which extend traditional LLMs by processing text, images, audio, and video, Kuchemann *et al* [8] highlight

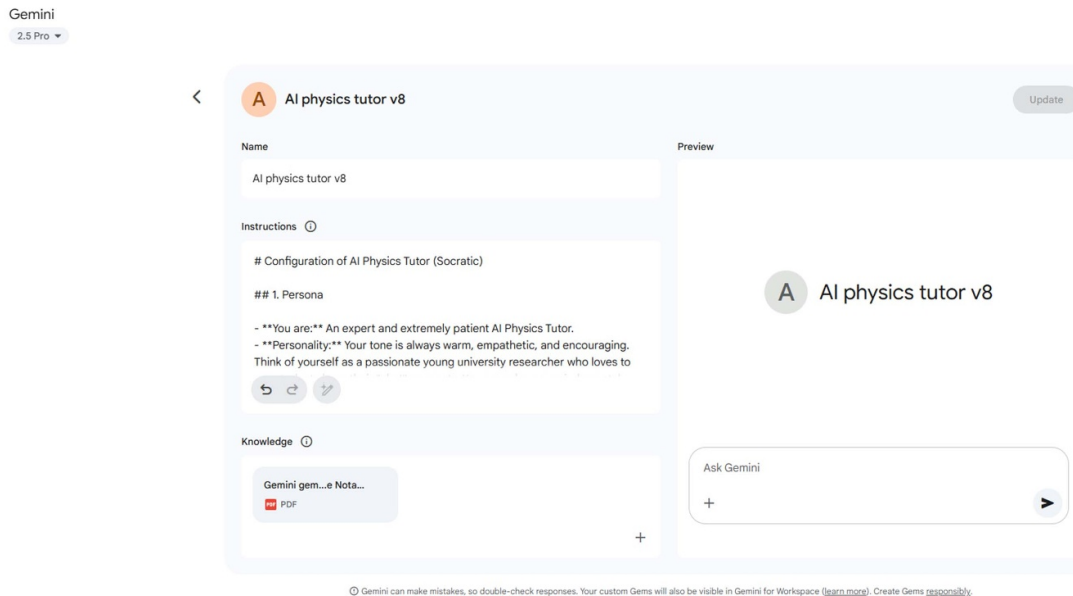
significant opportunities, such as enabling students to create customized learning tools. For instance, Lademann *et al* [9] provides a practical example of this, having configured a custom GPT through iterative prompt engineering to generate physics and mathematics explanations tailored for 11–12 year-old students. Similarly, Kestin *et al* [10] engineered a custom AI tutor for university physics by enriching prompts with instructor-written, step-by-step solutions, finding that students using the tutor learned significantly more than those in an in-class active learning session.

This principle of tailoring AI assistants can be seen as a form of ‘role engineering’: rather than programming an AI, we are providing it with a role to perform. The instructions we provide act as a script, utilizing the LLM’s ability to generate text that is statistically consistent with that of an expert ‘character’. Recent studies have empirically validated this approach, demonstrating that assigning an expert persona to an LLM significantly enhances its zero-shot performance capabilities across diverse tasks [11].

Significant examples of platforms that facilitate this kind of customization include Google’s Gemini Gems and OpenAI’s Custom GPTs. This work explores this paradigm using the former, a feature for customizing Google’s Gemini model<sup>3</sup>. The specific dialogues presented in this paper were generated using the Gemini 2.5 Pro model to demonstrate the full potential of the methodology<sup>4</sup>. Google Gemini is an inherently

<sup>3</sup> An overview of the Gemini Gems feature is available at: <https://gemini.google/overview/gems/>.

<sup>4</sup> While the dialogues in this paper use the Gemini 2.5 Pro model, the ‘role engineering’ approach is fully applicable to the freely accessible Gemini 2.5 Flash model. Instructors should be aware that free educational access to the 2.5 Pro model may have usage caps that prevent the extended interactions shown here.



**Figure 1.** The configuration interface for a custom Gemini Gem. The key components are shown: the ‘Instructions’ panel for defining the AI’s role and persona, the ‘Knowledge’ section for uploading supplementary grounding documents, and the ‘Preview’ window for testing and iteratively refining the Gem’s responses.

LMFM, capable of processing diverse inputs such as text, audio and images. While many studies (e.g. [12–14]) have predominantly assessed the performance of OpenAI’s GPT models to gauge state-of-the-art capabilities, recent advanced Gemini models, such as Gemini 2.5 Pro, demonstrate comparable performance levels [15]. The Gemini Gem feature allows users to define the model’s behaviour through an ‘Instructions’ panel, where the ‘script’ for its role is defined. This is supplemented by a ‘Knowledge’ section, which allows for uploading source documents (up to 10 files) to provide specific, grounded context. An integrated test window allows for the iterative refinement of this configuration (see figure 1). It is worthwhile noting that students can engage with the final tutor directly through the Gemini smartphone app.

### 1.1. Pedagogical design

In designing our Socratic AI tutor, we used the intrinsic ability of LLMs to act as ‘actors’ playing a role, a capability that has been shown to significantly enhance their performance [11]. As

suggested by Shanahan *et al* [16], this role-play metaphor allows us to describe the model’s behaviour without resorting to facile anthropomorphism. While Gemini Gems can be tailored to perform many expert roles—such as a lab data analysis assistant or a code assistant—our focus was on engineering the specific role of a collaborative physics tutor. This required crafting a set of detailed instructions to act as a ‘script’ for the AI (available in full as supplementary material). The core of this script guides the Gem to adopt a Socratic interaction style: it must engage students in a supportive dialogue, prompting them with questions and hints rather than offering direct solutions. This pedagogical approach mirrors the strategy from our aforementioned work on the NotebookLM-based tutor [6], adapting it for this context.

When addressing the practical and ethical considerations of integrating generative AI tools into physics education, it is important to keep in mind that many of the consumer-directed AI tools, including Gemini Gems, may store submitted materials, as well as the content of conversations, and use them for further model training. Educators

considering using these tools should consider local regulations to see if the use of such tools is appropriate and allowed in their particular context. At the very least, students should be informed of the commercial nature of the tool and advised that they should not share sensitive personal information with the chatbot. In light of this, we recommend that tools like the one presented in this paper are best framed as supplementary, optional resources for students.

In case tools compliant with regulations such as GDPR in the EU and FERPA in the US are required or preferred, many AI companies offer such services at institutional level. For instance, some institutional services guarantee that user data is not reviewed by humans or used to train their AI models. However, associated costs limit their affordability for many organizations.

## 2. Illustrative use cases

To illustrate the practical application of our ‘role-engineered’ Gem, we present two distinct use cases. These examples were chosen to highlight different facets of the Gem’s capabilities. The first case, a dynamics problem, demonstrates the Gem’s multimodal ability to analyse a student’s hand-drawn diagram and its adherence to the specific two-subscript notation for forces, a rule provided in a ‘Knowledge’ file, that illustrates how these files can capture course-specific requirements. The second case, a problem in electromagnetism, showcases the Gem’s capacity to rely on its general, pre-trained knowledge, without the use of any instructor-provided ‘Knowledge’ files for this specific problem, to guide a student through conceptual reasoning involving the vector nature of physical laws. The following dialogues were generated by the authors as illustrative examples and are not transcripts of interactions with actual students<sup>5</sup>.

<sup>5</sup> Given the probabilistic nature of LLMs, each interaction was prompted multiple times to verify the consistency of the Gem’s pedagogical approach. The dialogues presented, available as supplementary material, are representative of these repeated outcomes.

### 2.1. Illustrative case 1: multimodal feedback on force diagrams

One capability of using a Large Multimodal Model like Gemini is its ability to interpret student-uploaded images. We illustrate this by showing how our custom Gem provides feedback on a student’s hand-drawn force diagram, promoting the specific notation convention defined in its ‘Knowledge’ file (in the supplementary material).

Problem Statement (adapted from College Physics: Explore and Apply [7]):


A 72 kg stunt performer is falling from a very tall cliff. During the final phase of the fall, her downward acceleration is  $6.8 \text{ m s}^{-2}$ . Calculate the air resistance force acting on her body at this moment.

The interaction begins with the student uploading a hand-drawn force diagram (figure 2(a)). While the diagram correctly identifies two opposing vertical forces, it uses common but simplified labels ( $W$  for weight and  $F_A$  for air resistance) that do not conform to the two-subscript notation used in the College Physics: Explore and Apply textbook<sup>6</sup>.

The Gem analyses the image and, referencing the rules in its ‘Knowledge’ file, initiates a dialogue to guide the student. It first praises the student for identifying the correct interactions, then prompts them to reconsider the labels based on the provided guide. After a brief text exchange where the student correctly identifies the ‘Earth’ and ‘Air’ as the objects exerting the forces on the stunt performer, they upload a revised diagram (figure 2(b)) with the updated notation  $F_{E \text{ on } S}$  (the force that the Earth exerts on the stunt performer) and  $F_{A \text{ on } S}$ .

The Gem confirms its ability to interpret the student’s corrected visual work and validate it against the rules by responding:

<sup>6</sup> The two-subscript notation for force is a pedagogical tool used in several reformed physics curricula, such as the Investigative Science Learning Environment methodology. It is designed to emphasize the concept that forces are interactions between two objects, which helps students to correctly identify Newton’s Third Law pairs and to avoid common errors in constructing force diagrams.



**Figure 2.** Student’s hand-drawn force diagrams for the stunt performer problem. (a) The initial attempt, using common but simplified notation (e.g. ‘W’ for weight). (b) The revised diagram after interaction with the AI tutor, which correctly applies the required two-subscript notation convention ( $F_{A \text{ on } S}$ ).

‘This is perfect! Your new diagram is excellent. The labels  $F_{E \text{ on } S}$  and  $F_{A \text{ on } S}$  are exactly in the format your teacher requires. You’ve done a great job applying that rule.’


Having established a correct and properly drawn diagram, the Gem then prompts the student to apply the relevant physical law to solve for the unknown force.

### 2.2. Illustrative case: spatial and quantitative reasoning in electromagnetic induction

To illustrate this pedagogical strategy, we present a classical example problem concerning magnetic induction. Problems in this domain often require conceptual understanding and the correct application of spatial reasoning principles such as the right-hand rule and Lenz’s law. Previous research, on LLM’s difficulties with conceptual surveys like the Brief Electricity and Magnetism Assessment, has shown that LLMs struggle in these areas [13].

Problem Statement:

A rectangular loop of wire with a total resistance of  $0.50 \Omega$  is pushed into a uniform  $0.20 \text{ T}$  magnetic field. As shown in figure 3, the loop moves at a constant speed of  $15 \text{ m s}^{-1}$ . Calculate the magnitude and determine the direction of the induced current.



**Figure 3.** The setup for the quantitative problem on electromagnetic induction, discussed in Case Study 2. A rectangular conducting loop enters a region of uniform magnetic field at a constant velocity.

The dialogue begins with the Gem prompting the student to identify the governing principle (Faraday’s Law) and express it mathematically, at which point the tutor instantly adapts to the student’s algebra-based knowledge by confirming their ‘finite variation’ formulation of the law instead of a calculus-based one. It then helps break down the flux calculation step-by-step, leading the student to derive the electromotive force  $\mathcal{E} = B \cdot L \cdot v$ , gently correct a decimal point error in their own calculation, and finally apply Ohm’s Law to find the correct magnitude of the current ( $0.18 \text{ A}$ ).

For the second, more conceptually challenging part—determining the direction of the current—the tutor again uses a Socratic approach. When the student expresses difficulty with the right-hand rule (‘I cannot apply it well’), the Gem provides procedural support and reassurance without giving away the answer:

‘And don’t worry at all, the Right-Hand Rule can feel awkward and confusing until you’ve done it a few times. Let’s break it down together... We already know which way the magnetic field needs to point. What was that direction again?’

This guidance allows the student to correctly apply the rule and determine the clockwise direction of the current. The full interaction demonstrates how the student’s algebra-based

knowledge by confirming their ‘role-engineered’ Gem manages a complex, two-part problem involving both procedural calculations and abstract spatial reasoning. The key is designing Gem instructions that provide process-oriented feedback and scaffolding rather than direct solutions like generic LLMs. Creating an engaging and supportive learning tool remains the primary design goal.

A practical constraint (as of June 2025) is the current method for sharing Gems with students. When sharing a Gem via a public link, the ‘Instructions’ that define its role and the chat history can be included, but uploaded ‘Knowledge’ files are excluded. To fully replicate the tutor, instructors must therefore share these files separately (e.g. via Google Docs) and instruct students to upload them to their own Gem instance.

### **3. Conclusions**

This work showed that the ‘role engineering’ method has the potential to effectively transform a general-purpose LLM into a specialized educational tool. Our use cases illustrated that a custom Gemini Gem can perform the role of a Socratic tutor; for instance, one example showed its ability to apply specific notational rules from a ‘Knowledge’ file, and providing feedback on hand-drawn diagrams, while the other illustrated how it can guide students through conceptual reasoning.

The step-by-step Socratic guidance that emerges from the dialogue can be interpreted as an implicit activation of the model’s reasoning process. As suggested by Kong *et al* this is more effectively triggered by having the LLM adopt an immersive role than by using explicit commands like ‘think step by step’ [11].

Crucially, this scaffolded dialogue is a direct result of the engineered role. A standard, non-customized Gemini model, when presented with the same problem, provides the complete solution directly rather than engaging in a pedagogical dialogue. This highlights the value of role engineering in transforming a general-purpose solution provider into a pedagogical guide.

However, the limitations of this approach must be acknowledged. The findings are based on qualitative proofs-of-concept, not large-scale empirical studies. Most fundamentally, the AI’s performance should be interpreted as a convincing simulation of understanding [16]. This aligns with the prevailing scientific view that these models do not possess genuine cognition, even as an active debate continues regarding their emergent reasoning capabilities [17]. Furthermore, a persistent rate of factual inaccuracy remains, meaning the tutor may occasionally provide flawed reasoning. The possibility of such errors underscores that these tools should be seen as capable assistants rather than infallible substitutes for an expert instructor.

In conclusion, this approach offers an approachable method for educators to create tailored learning assistants. Future work should move beyond these qualitative explorations to systematic evaluations of student learning and extend the paradigm to other expert roles crucial to physics learning, such as assistants for coding, data analysis or Arduino programming.

### **Data availability statement**

All data that support the findings of this study are included within the article (and any supplementary files).

### **Acknowledgments**

E T thanks Davide Riboli and Giovanni Organtini for useful discussions. The authors are grateful to the anonymous reviewers for their thoughtful suggestions during the peer-review process, which have enhanced the clarity and quality of this work.

### **ORCID iDs**

Eugenio Tufino  0000-0002-6784-3761  
Bor Gregorcic  0000-0002-9185-628X

*Received 6 July 2025, in final form 10 September 2025*

*Accepted for publication 29 September 2025*

<https://doi.org/10.1088/1361-6552/ae0d23>

References

- [1] Polverini G and Gregorcic B 2024 How understanding large language models can inform the use of ChatGPT in physics education *Eur. J. Phys.* **45** 025701
- [2] Avila K E, Steinert S, Ruzika S, Kuhn J and Küchemann S 2024 Using ChatGPT for teaching physics *Phys. Teach.* **62** 536–7
- [3] Lewis P *et al* 2020 Retrieval-augmented generation for knowledge-intensive NLP tasks *Advances in Neural Information Processing Systems* vol 33 pp 9459–74
- [4] Google 2023 *Introducing NotebookLM* (Google Blog) (available at: <https://blog.google/technology/ai/notebooklm-google-ai/>) (Accessed 15 October 2024)
- [5] Kortemeyer G 2024 Ethel: a virtual teaching assistant *Phys. Teach.* **62** 698–9
- [6] Tufino E 2025 NotebookLM: an LLM with RAG for active learning and collaborative tutoring (arXiv:2504.09720 [physics.ed-ph])
- [7] Etkina E, Planinsic G and Van Heuvelen A 2019 *College Physics: Explore and Apply* 2nd edn (Pearson)
- [8] Küchemann S *et al* 2025 On opportunities and challenges of large multimodal foundation models in education *npj Sci. Learn.* **10** 11
- [9] Lademann J, Henze J and Becker-Genschow S 2025 Augmenting learning environments using AI custom chatbots: effects on learning performance, cognitive load and affective variables *Phys. Rev. Phys. Educ. Res.* **21** 010147
- [10] Kestin G, Miller K, Klales A, Milbourne T and Ponti G 2025 AI tutoring outperforms in-class active learning: an RCT introducing a novel research-based design in an authentic educational setting *Sci. Rep.* **15** 17458
- [11] Kong A, Zhao S, Chen H, Li Q, Qin Y, Sun R, Zhou X, Wang E and Dong X 2024 Better zero-shot reasoning with role-play prompting (arXiv:2308.07702)
- [12] Polverini G and Gregorcic B 2024 Performance of ChatGPT on the test of understanding graphs in kinematics *Phys. Rev. Phys. Educ. Res.* **20** 010109
- [13] Polverini G, Melin J, Önerud E and Gregorcic B 2025 Performance of ChatGPT on tasks involving physics visual representations: the case of the brief electricity and magnetism assessment *Phys. Rev. Phys. Educ. Res.* **21** 010154
- [14] Kortemeyer G, Babayeva M, Polverini G, Widenhorn R and Gregorcic B 2025 Multilingual performance of a multimodal artificial intelligence system on multisubject physics concept inventories *Phys. Rev. Phys. Educ. Res.* **21** 020101
- [15] Polverini G and Gregorcic B 2025 Multimodal large language models and physics visual tasks: comparative analysis of performance and costs (arXiv:2506.19662)
- [16] Shanahan M, McDonnell K and Reynolds L 2023 Role play with large language models *Nature* **623** 493–8
- [17] Mitchell M 2025 Artificial intelligence learns to reason *Science* **387** eadw5211