



# LLM-generated tips rival expert-created tips in helping students answer quantum-computing questions

Lars Krupp<sup>2,1\*</sup>, Jonas Bley<sup>3</sup>, Isacco Gobbi<sup>4</sup>, Alexander Geng<sup>4</sup>, Sabine Müller<sup>4</sup>, Sungho Suh<sup>1</sup>, Ali Moghiseh<sup>4</sup>, Arcesio Castaneda Medina<sup>4</sup>, Valeria Bartsch<sup>5</sup>, Artur Widera<sup>3</sup>, Herwig Ott<sup>3</sup>, Paul Lukowicz<sup>1,2</sup>, Jakob Karolus<sup>1,2</sup> and Maximilian Kiefer-Emmanouilidis<sup>1,2,3</sup>

\*Correspondence: [kruppl@rptu.de](mailto:kruppl@rptu.de)

<sup>2</sup>Department of Computer Science and Research Initiative QC-AI, RPTU Kaiserslautern-Landau, Kaiserslautern, 67663, Germany

<sup>1</sup>Embedded Intelligence, German Research Center for Artificial Intelligence, Kaiserslautern, 67663, Germany

Full list of author information is available at the end of the article

## Abstract

Alleviating high workloads for teachers is crucial for continuous high quality education. To evaluate if Large Language Models (LLMs) can alleviate this problem in the quantum computing domain, we conducted two complementary studies exploring the use of GPT-4 to automatically generate tips for students. (1) A between-subject survey in which students (N = 46) solved four multiple-choice quantum computing questions with either the help of expert-created or LLM-generated tips. To correct for possible biases, we additionally introduced two deception conditions. (2) Experienced educators and students (N = 23) directly compared the LLM-generated and expert-created tips. Our results show that the LLM-generated tips were significantly more helpful and pointed better towards relevant concepts while also giving away more of the answers. Furthermore, we found that participants in the first study performed significantly better in answering the quantum computing questions when given tips labeled as LLM-generated, even if they were expert-created. This points towards a placebo effect induced by the participants' biases for LLM-generated content. Ultimately, we contribute that LLM-generated tips can be used instead of expert tips to support teaching of quantum computing basics.

**Keywords:** Large Language Models; Quantum computing; Education; Survey

## 1 Introduction

For a long time, Quantum Computing (QC) was a niche area within the field of physics that received little to no attention in other Science, Technology, Engineering & Mathematics (STEM) fields. The recent hype in quantum technologies also known as 'the second quantum revolution' [1–4] has led to significant advances in hardware [5–7] and formulations of technological standards [8–10] in the areas of quantum communication, sensing and computation. Today, QC is one of the fastest-growing interdisciplinary research areas at the intersection of computer science, mathematics, and physics [11–14]. Thus, it is no surprise that quantum computing has become an interesting topic for industry and academia [15–18], higher education [19–23], and even high school education [24–27].

© The Author(s) 2025. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

However, the amount of experts and educators is still quite limited and mostly intertwined in academia and large companies who have taken the lead in QC such as Google, IBM and Microsoft [3, 5–7, 28, 29]. The initial barrier to learn quantum algorithms is high, demanding experienced educators and individual teaching. This poses a burden for educators, adding to their workload and making QC mostly inaccessible to learners. Consequently, there is a need to make learning and teaching about QC more accessible. In our work, we want to address these challenges and investigate ways to support educators.

In this work, we explore the use of generative educational content from Large Language Models (LLMs) as a support mechanism for educators, alleviating some of their workload and allowing for more personalized teaching. LLMs are a promising new technology that is already successfully used in education [30]. However, how, if and when LLMs should be used in education and what possible consequences might arise from that use for students and educators is a strongly debated topic that is studied intensively at the moment [31]. While some results demand caution [32–34], others report more neutral [35, 36] or positive results [37]. This shows a need for research on the use of LLMs for quantum computing education.

LLMs first became a topic of public interest with the publication of ChatGPT in November of 2022 [38]. Since then, they have been used successfully in a number of different domains [39–42] and a considerable amount of resources has been invested to improve this technology, supported by large corporations [43–47]. This led to significant scientific progress being made in a short amount of time. Now, with models like GPT4 [48] and Llama3 [44] at the forefront of these developments, many tasks that might have been considered difficult to automate a few years ago, can be done to a satisfying degree by computers. This includes writing code [42] and creative writing [41]. In addition, supporting architectures are being developed, such as Retrieval-Augmented Generation (RAG) [49], further improving the quality of model outputs by providing them with the relevant context to answer specific questions, while reducing the likelihood of bad outcomes, such as hallucinations [50].

This significant recent interest in LLMs has led to the emergence of many new research questions in the field of education. These can be divided into the discovery of possible applications, the evaluation of practicality and effectiveness, and the discussion of challenges and ethicality [51]. As summarized by Yan et al. [51], possible applications of LLMs in education are vast: They include profiling and labeling (classification and structuring of texts and tasks), detection (e.g., of student's and teacher's behaviors and intents), assessment and grading (formative and summative, and also independent of the teacher), teaching support (in- and out-of-classroom), prediction (of student performance, engagement and growth), knowledge representation (graph construction and tracing), feedback (real time or post-hoc and feedback on feedback), content generation (of questions, tasks, code and replies), and recommendation (e.g., of resources and courses).

Scaffolding, the educational practice of providing support and guidance to students when needed [52], is one of these applications. It can help to achieve the desired learning outcome by reducing the effort needed to complete non-essential tasks, allowing students to focus their attention on the important tasks. Providing tips dynamically and individually is a widely used scaffolding method in education [53]. However, this can be time-consuming or might even become unfeasible with growing student-to-educator ratio. LLMs can help alleviate this issue by providing individualized tips on demand [54, 55].

The use of LLMs is seen as natural in the shift to the so-called smart education [54, 55], which is the practice of providing personalized learning opportunities on demand, often associated with digitization [56]. With increasing capability of LLMs, the number of possible educational applications increases and challenge and risk assessments change [57, 58]. Due to this significant change of landscape, there have not yet been many studies assessing the effectiveness of state-of-the-art LLMs for education [51]. In the same sense, a few years ago, using LLMs to aid in the context of quantum education would have been unthinkable.

We contribute to this discourse with two interconnected studies. We evaluated how LLM-generated tips compared to expert-generated tips when answering quantum computing questions in two studies. Our main study simulates a practice scenario for students, tasking them to answer questions about the lecture material providing tips as support. This study was conducted as part of the Quantum-Initiative Rhineland-Palatinate (QUIP) winter school QUIKSTART 2024.<sup>1</sup> The second study complements this initial evaluation by validating the tips in a within-subject design focusing on opinions from both experienced educators and students with regard to the employed tips in the main study. By using the same tips in both studies, we ensure that the results can be aggregated to form a bigger picture.

Our results indicated that LLM-generated tips rivaled those of experts, being perceived as significantly more helpful and focusing more on conveying relevant concepts. Yet, certain design limitations need to be considered, such as a possible placebo effect of artificial intelligence [59] and generally more leading tips when generating them with an LLM.

In this work, we take a step towards answering some of the underlying questions concerning the quality, helpfulness, and correctness of LLM-generated educational content for essential teaching methods [52] such as scaffolding, in the domain of quantum computing. Our findings point towards a paradigm shift of LLM-supported smart education, highlighting its feasibility and providing insights into its successful application.

## 2 Methods and research questions

In this work, we intend to answer two research questions through two interconnected studies. In our main study, we asked how GPT4-generated tips compare to expert-created tips in the domain of quantum computing. We conducted this study with participants (N = 46) from the Quantum-Initiative Rhineland-Palatinate (QUIP) winter school QUIKSTART 2024. Employing a between-subject design with the factors *tip creator* and *tip label*, we tasked the students with four questions about quantum computing. Students were given one tip for each question. Tip creators (LLM, expert) and how the tips were labeled (LLM, expert) were subject to the condition assigned to each participant, including two consistent conditions (creator equaled label) and two contradictory conditions (creator and label differed). The students were not made aware of the real tip creator until after completing the study.

To validate our findings, we conducted a complementary tip evaluation study (N = 23). Here we specifically asked people with the relevant competencies (both experienced educators and students) to rate how LLM-generated tips compare to expert-created tips.

---

<sup>1</sup><https://www.quantum-rlp.de/studierende/sommer-und-winterschulen/>.

*RQ1: Can LLM-generated tips be used instead of expert-created tips to help students answer quantum-physics questions?* This research question is composed of two parts.

*RQ1a: Is there a difference in the score, an objective measure, when using LLM-generated tips compared to expert-created tips?* To answer this question, we evaluated the scores achieved by students in practice scenario of the main study.

*RQ1b: Are there subjective differences in the perception of LLM-generated tips and expert-created tips?* Here we evaluated the subjective measures of our studies. For the main study, we analyzed whether there are differences in perceived correctness, helpfulness, and quality between our conditions. Additionally, we conducted the tip evaluation study, where we asked students and educators to rate the helpfulness and correctness of the tips and if they gave away the answer or pointed to relevant concepts. This was done without giving them any information on the original tip creator.

*RQ2: What are adverse effects of LLM-generated tips?* Recent work in human computer interaction has shown, that the mere label of AI can influence the participants' perception, culminating in a placebo effect if participants' are falsely told that there is an AI system in place [59, 60]. To evaluate if such a bias or placebo effect exists when participants receive LLM-generated tips, we introduced two deception conditions to our main study, where tip creator and tip label are not the same. This method allowed us to analyze if the tip label influenced any of our metrics. Additionally, we used the tip evaluation study to analyze if the LLM-generated tips gave away more of the answers compared to the expert-created questions. This could negatively impact the students learning-gains. Even though we eliminated the possibility to give wrong tips due to hallucination by manually checking each LLM-generated tip, we evaluated the perceived tip correctness and quality in the main study and the tip correctness in the tip evaluation study.

## 2.1 Multiple-choice question creation

We chose topics related to the basics of quantum computing for our questions and made sure that they were covered at QUIKSTART 2024. This way, we established a common ground of knowledge for all study participants taking part in our main study.

We chose two topics and two difficulty levels within each topic. All questions and their answers can be found in the Appendix A. Firstly, the Bloch sphere, a typical graphical visualization of single-qubit systems. However, this representation can still cause some difficulties in learners [61]. The first, simpler, question asks learners to identify the state  $|\psi\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$  on the Bloch sphere. The second, more advanced, question asks the same for the state  $|\psi\rangle = \frac{1}{2}|0\rangle + \frac{\sqrt{3}}{2}e^{i\pi/2}|1\rangle$ .

For the second topic, we decided on some basics of two-qubit quantum computing: a gate identity involving two Hadamard gates and two CNOT gates (the simpler question) and the exact outcome of a circuit meant for the creation of a Bell state (the more advanced question). The gate identity asks participants to know that the property of these quantum gates is their own inverse. Bell states are the most basic two-qubit maximally entangled states that are essential for any teaching course about entanglement. The creation process of these states is therefore an important knowledge domain. In the question we used, participants were asked to predict the outcome of the  $(X \otimes I)CNOT_{12}(H \otimes I)$  operation on the  $|00\rangle$  state.

## 2.2 Tip creation

For each of the four multiple-choice questions of the main study, we acquired one expert and one LLM tip. The expert tips were created by two quantum physics experts with teaching experience in the field who were giving lectures at the winter school. They had access to the questions, answers and the lecture script but not to the LLM-generated tips, to not influence their creative process. A different quantum physics expert confirmed the correctness of the tips and that they do not give away the answer.

The LLM tips were generated by giving the question, answers, and lecture script explaining the quantum computing topics to GPT4 and prompting it to generate a tip for each question in a RAG-like process. Using a zero-shot approach, up to five tips per question were generated. In the next step, we randomly selected a tip for each question. Each of those was later validated and ascertained to be correct by a quantum computing expert. To ensure educators would be able to repeat this process for their lectures, GPT4 was used in its standard configuration through the web interface. The prompt used is available in Sect. C.

## 2.3 Main study

Our main study was a between-subject survey conducted during QUIKSTART 2024. A winter school conducted specifically for students with a background in machine learning or quantum physics. Participation was entirely voluntary and allowed students that participated in the winter school online or in person and students that did not participate in the winter school to test their quantum computing knowledge in the areas of Bloch spheres and Bell states using multiple-choice questions (Sect. 2.1). Participants were tasked to answer the given questions and were additionally provided with a tip for each of the questions. The type of tip was dependent on the randomly assigned conditions covering two factors: tip creator (Sect. 2.2) and tip label. Hence, a tip could either be created by an expert or the LLM. Likewise, it could either be labeled as being created by an expert or generated by an LLM, yielding a total of four different conditions (see Fig. 1). We did not include a condition where students received no tip at all. Since our work is based on scaffolding as a teaching method, we already know that providing tips is an effective improvement during learning [52, 53].

Henceforth, we will use the following structure to differentiate between conditions: CREATOR-LABEL. In the first condition (EXPERT-EXPERT) the tip was created by an expert and labeled as such. LLM-EXPERT refers to LLM-created tips that were displayed to the participant as if they were created by an expert. LLM-LLM refers to LLM-created tips

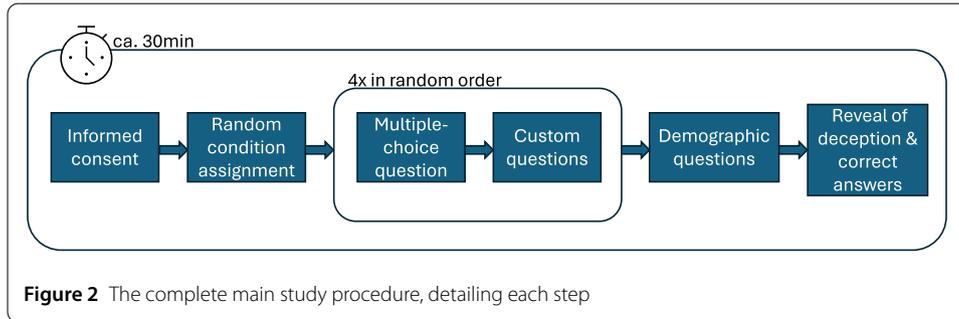
### This tip was created by an Expert:

Realize you are in an equal superposition: Compute the probability of outcome 0 and 1. On the Bloch sphere an outcome is probable if the state vector is closer to the pole corresponding to the outcome. The definition of the relative phase, relative to the x-axis on the Bloch sphere, gives you the representation on the Bloch sphere.

**Figure 1** A screenshot of how tips were displayed in the main study

**Table 1** helpfulness, quality, and correctness asked for each displayed tip; from *strongly disagree* to *strongly agree*; all visual analog scale (0 to 100)

Please rate the following statements regarding the displayed tip.	
CQ1	The tip was helpful.
CQ2	The quality of the tip was high.
CQ3	The tip was correct.



labeled as such. Finally, EXPERT-LLM refers to expert-created tips labeled as if they were created by an LLM.

By including the LLM-EXPERT and EXPERT-LLM conditions, and therefore deceiving participants that were assigned to those conditions, we were able to determine the labels' influence on the participants perception of the tip. This allowed us to evaluate the existence of possible placebo effects that have been shown to occur in some settings where humans interact with Artificial Intelligence (AI) [59]. The deceived participants were informed immediately after answering their demographic questions and shown the correct answers, mitigating possible negative effects of the deception [62].

We measured the students' scores and further evaluated the perceived usefulness of the displayed tips using custom questions (see Table 1), covering correctness, helpfulness, and quality of the tip. Additionally, we asked participants to rate the perceived difficulty of the question. Ethics approval was granted by the Ethics Department of the German Research Center for Artificial Intelligence (DFKI).

### 2.3.1 Main study procedure

The complete study procedure is depicted in Fig. 2. After providing informed consent and being assigned a condition, students were shown the four quantum computing question in randomized order. Each question was displayed individually including its tip. After providing the answer, we inquired about the qualities of the display tips using our custom questions (see Table 1). This process was repeated for each question.

Finally, participants were asked to provide demographics (age, gender), got informed about the deception (if any) and shown the correct answer to each question. The whole study took approximately 30 minutes.

### 2.3.2 Main study participants

In total, 46 participants (Age:  $\bar{x} = 28.9y$ ,  $s = 7.35y$ ; 35 male, 8 female, 2 no comment, 1 other)<sup>2</sup> completed the study. Of these participants, 25 attended the winter school in per-

<sup>2</sup> $\bar{x}$  denotes the mean of a sample;  $s$  denotes the standard deviation of a sample.

**Table 2** The distribution of participants onto the four conditions

Tip creator	Tip label	#Participants
Expert	Expert	9
LLM	Expert	12
LLM	LLM	10
Expert	LLM	15

**Table 3** Additional questions asked for each displayed tip; from *strongly disagree* to *strongly agree*; all visual analog scale (0 to 100)

Please rate the following statements regarding the displayed tip.	
CQ4	The tip was correct.
CQ5	The tip gave away the correct answer
CQ6	The tip was helpful.
CQ7	The tip effectively points towards relevant concepts.

son, 18 online, and 3 did not participate. Due to the random condition selection, the participants are not evenly distributed onto the four conditions. The participant distribution and demographics can be seen in Table 2.

## 2.4 Tip evaluation study

After we completed the main study, we conducted a separate tip evaluation study. The purpose of this study was to form a better understanding of the quality of the tips used in the main study. In this within-subjects study, we had participants rate each tip on a custom scale adjusted to facilitate the evaluation of measures not tested in the main survey (see Table 3) and allowed for direct comparison of LLM-generated and expert-created tips. We acquired participants by asking educators (at least postdoctoral researchers with teaching experience in relevant fields) and students from related fields to participate. This way, we made sure that each participant had the necessary knowledge to rate the tips. The study was not part of the official program at QUIKSTART 2024, unlike our main study. Ethics approval was granted by the Ethics Department of the German Research Center for Artificial Intelligence (DFKI).

### 2.4.1 Tip evaluation study procedure

After providing their informed consent, participants were asked to quantify their perceived knowledge of quantum physics and to answer some demographic questions. After this, the participants were shown each of the four multiple-choice questions used in the main study, their answer options, the associated correct answer, and both their LLM-generated and their expert-created tips. They were not informed about the creator of each tip, only that they should rate and compare them. For each tip, participants answered if it was helpful, correct, gave away the correct answer, and effectively points towards relevant concepts on a scale of 0 to 100, with 0 being the maximum disagreement, and 100 the maximum agreement (see Table 3). Furthermore, for each question, the participants decided which tip, if any, was better. Finally, we asked the participants to explain their reasoning for that decision. This allowed us to qualitatively evaluate their responses. The whole study took approximately 10 minutes.

### 2.4.2 Tip evaluation study participants

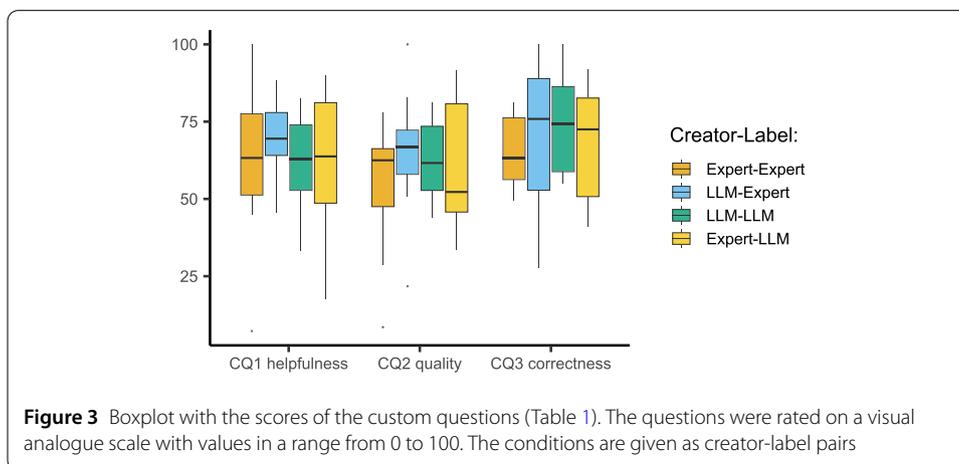
In total 23 participants completed the survey. The study participants were hand-selected to fulfill the required minimum competencies in QC to be able to rate the tips. We allowed participants from the main study to participate in this study (eight of 23 participated in both studies). They were divided into two groups, students and educators. The 13 participants designated as educators (Age:  $\bar{x} = 42.9 y$ ,  $s = 7.7 y$ ; 10 male, 3 female) have obtained a Ph.D. and are working and teaching in academia in a multitude of related fields. While not all of them are experts in quantum computing, all have a basic level of understanding that allowed them to judge the tips. This allowed us to get the opinion of multiple (at least four) professors on the topic. Seven educators are working in the area of physics, with two of them specializing in quantum physics, two are working in computer science, two in artificial intelligence, one in mathematics, and one in applied mathematics and physics. The other ten participants were designated as students (Age:  $\bar{x} = 27.6 y$ ,  $s = 5.16 y$ ; 7 male, 3 female). Of them six studied physics (four physics, two quantum physics), three computer science (one quantum machine learning, one machine learning, one computer science) and one gave no comment.

## 3 Results

### 3.1 Main study results

To evaluate the results of the main study, we conducted two-way analyses of variance (ANOVA) [63] using the factors tip creator and tip label. We constructed models for each of the measures in our main study, including the students' score, their rating of our custom questions (see Table 1) and their rating of the questions difficulty. In cases where normality of the data was violated, we first used aligned rank transform (ART) [64] on the data. We found no interaction effects regarding the two factors for any dependent variable and report statistics for the additive models only. Effect sizes are given using  $\eta^2$  (Partial Eta Squared): small ( $> 0.01$ ), medium ( $> 0.06$ ), large ( $> 0.14$ ). All statistics are reported in Sect. D.

*Custom questions* We found no significant effects of the tip label or tip creator on our custom questions of correctness, helpfulness, or quality. A visualization of the custom questions can be seen in Fig. 3. In general, tips were of high quality ( $\bar{x} = 61$ ,  $s = 19.18$ ),



**Figure 3** Boxplot with the scores of the custom questions (Table 1). The questions were rated on a visual analogue scale with values in a range from 0 to 100. The conditions are given as creator-label pairs

helpfulness ( $\bar{x} = 64.41, s = 18.66$ ) and correctness ( $\bar{x} = 69.91, s = 17.44$ ). Furthermore, there seems to be a tendency to rate questions as less difficult when the tip was labeled as LLM-generated ( $\bar{x} = 40.2, s = 17.9$ ) compared to expert-created ( $\bar{x} = 49.9, s = 16.8$ ).

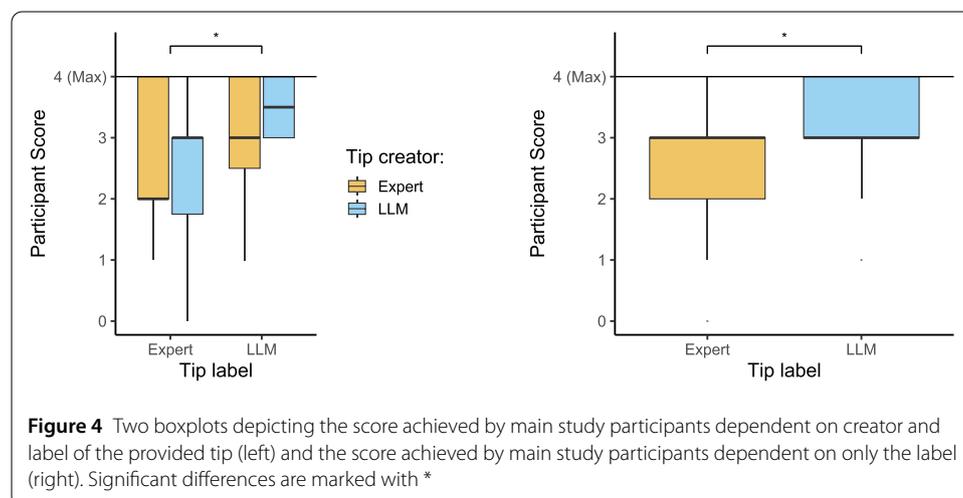
**Question difficulty** The overall mean for the perceived question difficulty was  $\bar{x} = 44.63$  ( $s = 17.91$ ), showing adequate difficulty levels for our question selection. Furthermore, we found that the simpler questions ( $\bar{x} = 39.87, s = 20.12$ ) were perceived as such, compared to the more advanced questions ( $\bar{x} = 49.39, s = 19.02$ ). Please refer to Sect. 2.1 for our selection process.

While we found no significant effect of tip label or creator on the perceived question difficulty, we found a tendency towards perceiving questions as less difficult when participants' were provided with a tip labeled as LLM-generated ( $\bar{x} = 40.2, s = 17.9$ ) compared to expert-labeled tips ( $\bar{x} = 49.9, s = 16.8$ ). There was no such tendency found for the tip creator, with only minimal differences between LLM-generated ( $\bar{x} = 45.2, s = 15.9$ ) and expert-created ( $\bar{x} = 44.1, s = 19.9$ ) tips.

**Participant score** An analysis of our objective measure showed a significant effect (medium effect) for the participants score ( $p < 0.05, \eta^2 = 0.12$ ) for the factor tip label. The factor tip creator did not significantly affect the participant score. We can conclude that there is a positive influence on the score ( $\bar{x}_{Label:Expert} = 2.48, s_{Label:Expert} = 1.21, \bar{x}_{Label:LLM} = 3.28, s_{Label:LLM} = 0.84$ ), when the tip is labeled as LLM-generated (see Fig. 4).

### 3.2 Tip evaluation results

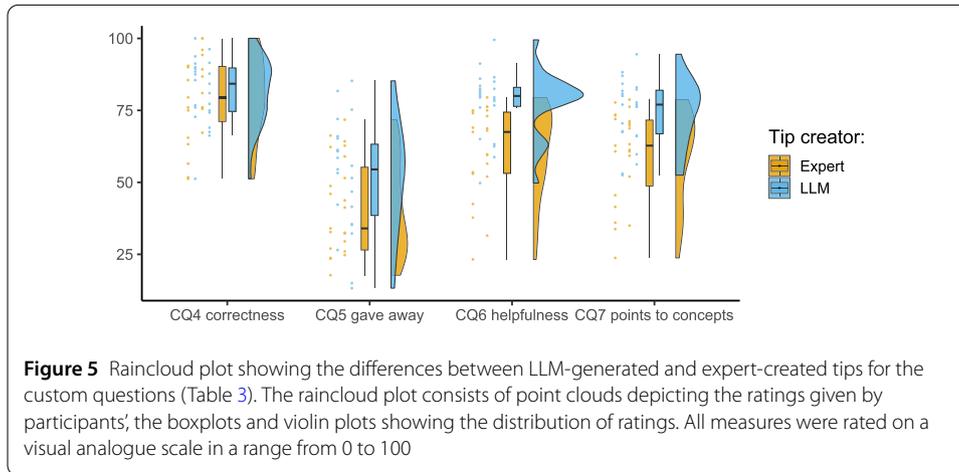
Similarly to our main study, we conducted one-way ANOVAs using the factor tip creator. Note that this study only included one factor. Again, we constructed models for each of our measures (Table 3). In cases where normality of the data was violated, we first align rank transformed the data using ART. Effect sizes are given using  $\eta^2$  (Partial Eta Squared). We first evaluated if the previous study participation influenced how participants' rated the tips, but found no significant effect, omitting it for the analysis. The complete statistics are shown in Table 4.



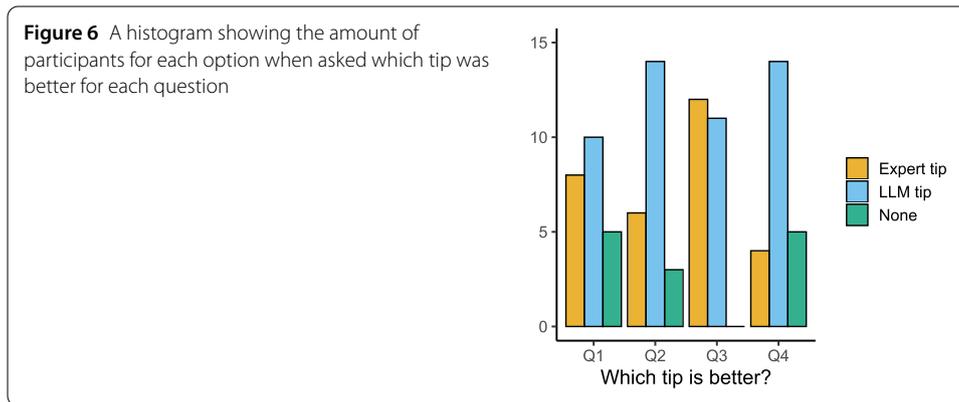
**Figure 4** Two boxplots depicting the score achieved by main study participants dependent on creator and label of the provided tip (left) and the score achieved by main study participants dependent on only the label (right). Significant differences are marked with \*

**Table 4** The ANOVA results of the tip evaluation study for the custom questions. Significance levels marked as ns := not significant;  $p < 0.05$  := \*;  $p < 0.01$  := \*\*;  $p < 0.001$  := \*\*\*

Measure	Df	Df.res	F value	p-value	Signif.	$\eta^2$ (partial)
correctness	1	22	0.15	0.702	ns	6.77e-03
gave away	1	22	9.42	0.006	**	0.30
helpfulness	1	22	22.81	9.075e-05	***	0.51
points to concepts	1	22	12.42	0.002	**	0.36



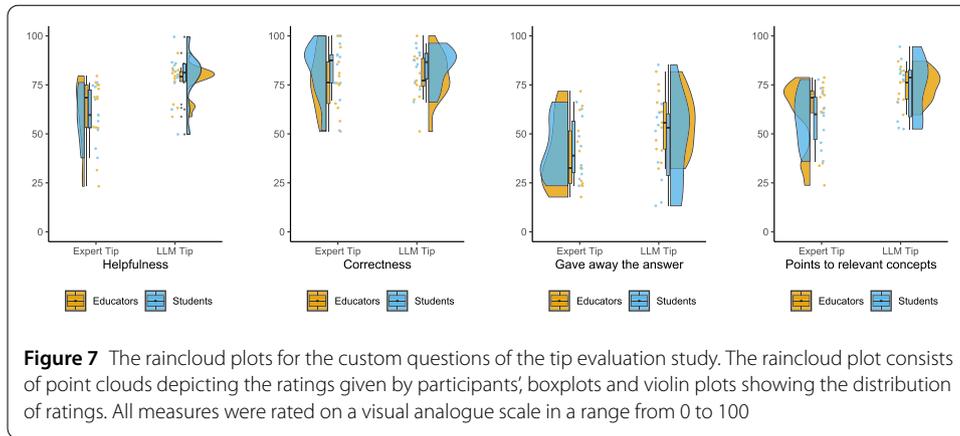
**Figure 5** Raincloud plot showing the differences between LLM-generated and expert-created tips for the custom questions (Table 3). The raincloud plot consists of point clouds depicting the ratings given by participants, the boxplots and violin plots showing the distribution of ratings. All measures were rated on a visual analogue scale in a range from 0 to 100



**Figure 6** A histogram showing the amount of participants for each option when asked which tip was better for each question

*Custom questions* We found no significant difference between LLM-generated and expert-created tips regarding the correctness. However, we found that LLM-generated tips scored significantly higher on giving away the answer ( $p < 0.05$ ,  $\eta^2 = 0.3$ ) compared to the expert-created tip. Furthermore, the LLM-generated tip was rated as significantly more helpful ( $p < 0.0001$ ,  $\eta^2 = 0.51$ ) and as pointing significantly better towards relevant concepts ( $p < 0.002$ ,  $\eta^2 = 0.36$ ) to answer the question. A visualization of these results is shown in Fig. 5.

*Direct tip comparison* We also asked participants to decide for each question, which tip they preferred (see Fig. 6), not showing the labels. In three out of four questions, the participants rated the LLM-generated tip as better when directly compared to the expert-created tip.



**Table 5** The length of the tips (in words) given for each question in comparison

	Q1	Q2	Q3	Q4	Average
Expert	58	48	26	23	38.75
LLM	110	148	162	182	150.5

*Differences between experts and students* When evaluating educators and students separately (see Fig. 7), we observed differences in the evaluation between the groups. As seen in Fig. 7, both educators and students rate the helpfulness of LLM-generated tips more highly compared to expert-generated tips. Furthermore, both groups are in agreement for the rating of correctness and which tip better points to the relevant concepts. However, when looking at which tip gave away the answer more, the ratings of educators and students differ. While educators rate LLM-generated tips as giving more away, for students the ratings between LLM-generated tips and expert tips are relatively similar with only a slight tendency towards the LLM-generated tips giving away more.

*Tip length* There are differences in the length of expert-created and LLM-generated tips. While the expert-created tips were generally short and on point, the LLM-generated tips tended to be lengthy and verbose (see Table 5). As one participant in the tip evaluation study put it, the LLM-generated tip “is more like a lecture note instead of a tip” (P20), highlighting the influence of the presentation on the perception of tips. In our qualitative evaluation (Sect. 3.3), the increased length of LLM-generated tips was also an issue.

### 3.3 Qualitative analysis

In the tip evaluation study, we asked participants for their reasoning on which tip was better for each question (see Table 6). We encoded that reasoning in multiple steps. First, we reduced the arguments written by the participants to their key terms, which were then aggregated into six groups. We then counted the amount of occurrences for the expert tip and the LLM tip.

The codings show that participants said more positive and more negative things about the LLM-generated tips. While participants mentioned a positive helpfulness 35 times for LLM-generated tips, they only mentioned it nine times for expert-created tips. Additionally, they mentioned that the LLM-generated tip contained good amount of detail ( $n_{LLM} = 8, n_{Expert} = 4$ ) and points to relevant concepts ( $n_{LLM} = 8, n_{Expert} = 4$ ) more of-

**Table 6** The codings of the reasons given by participants' on why one of the tips is better than the other. It shows the amount of times the tips were rated that way

Group	#Expert tip	#LLM tip
bad content	4	7
points to concepts	4	8
adequate detail	3	8
adequate helpfulness	9	35
gives away more	9	21
adequate length	3	3

ten. However, they also mentioned that the LLM-generated tip gave away parts of the answer ( $n_{LLM} = 21$ ,  $n_{Expert} = 9$ ) or produced bad content ( $n_{LLM} = 7$ ,  $n_{Expert} = 4$ ) more often.

## 4 Discussion

Our results highlighted potential benefits of using LLMs for tip creation in the domain of quantum computing. In the following, we discuss these benefits and expose pitfalls that should be avoided.

### 4.1 LLM-generated tips rival expert-created tips

We found that participants in the tip evaluation study (Sect. 3.2) rated the helpfulness (CQ6) of LLM-generated tips significantly higher than that of expert-created tips. Yet, we have not found the same effect for helpfulness (CQ1) in our main study (Sect. 3.1), possibly due to the lack of direct comparison. Participants in the main study might have rated it on how helpful the tip was to answer the question (CQ1), while participants of the tip evaluation study rated the helpfulness (CQ6) of the tips in comparison. As such, our results showcase that the helpfulness of LLM-generated tips is on par with expert-created tips and even prevail in a direct ranking (RQ1b). This is in line with other research about LLM-generated content [65].

However, LLM-generated tips were more prone to give away the correct answer (CQ5), possibly correlating with their perceived helpfulness (CQ6) and them pointing to relevant concepts (CQ7). Yet, this caveat was only present for the educator group in our tip evaluation study (Fig. 7). We believe that their different perception stems from the educators' teaching experience which the students are lacking. Consequently, there is a real possibility that LLM-generated tips more easily give away the correct answer (RQ2). This is a pitfall that needs to be avoided through proper prompt engineering [66] or possibly using an LLM as a judge [67] that is specifically trained to evaluate whether the LLM-generated tip gives away the answer. Other options would be few-shot prompting [68] where example tips are given to the LLM to increase output consistency or chain-of-thought prompting to improve the models reasoning capabilities [69] which would allow the model to construct tips more carefully. These methods could further improve tip quality and reduce the likelihood of the LLM-generated tip giving away the answer to a question.

We found no significant differences between the correctness of LLM-generated tips and expert-created tips in the main study (CQ3) or in the tip evaluation study (CQ4). Participants rated all tips — regardless of creator or label — as correct (Sect. 3.1) contributing to answering RQ1b. Finally, the LLM-generated tips pointed to relevant concepts (CQ7) significantly better compared to the expert-created tips, a characteristic that is highly valuable in teaching [70, 71]. In addition to the positive subjective perception of

LLM-generated tips, students did not score worse in our multiple-choice questions (main study, see Sect. 2.3). Here, LLM-generated tips did not negatively impact students' scores (Sect. 3.1), answering *RQ1a*.

Our results show that LLMs can be used successfully to generate tips in the context of basic quantum computing education, rivaling expert-created tips in all metrics, answering our first research question (*RQ1*).

#### 4.2 Novelty of AI possibly leads to increased scrutiny

While the tip creator did not influence the students' scores, a tip labeled as being LLM-generated (regardless of its original creator) elicited significantly higher scores. We believe that this phenomenon occurred due to an increase of attention given to tips labeled as LLM-generated. Due to the exhaustive media attention given to LLMs since the inception of ChatGPT, participants came into contact with these models. While some might have only heard of them, others might have tried them out or may even be actively using LLMs for some tasks. This knowledge could have made them aware of the persisting issues attributed to this technology. Namely hallucinations and confidently worded verbose texts lacking substance [72]. As a consequence, participants who were told that they were given an LLM-generated tip, might have examined the tip more carefully, thinking critically about its contents. This critical thinking might have led to a better understanding and, therefore, an improved score [73]. The measurable impact of introducing LLM labels is an important design consideration to make and needs to be handled carefully (*RQ2*).

We envision this effect due to two interacting factors impacting the students' perceptions of the offered support (the AI-generated tip) and the resulting perceived difficulty of the task.

Firstly, a measurable placebo effect has been shown to exist in previous work on human-computer interaction [59, 60, 74]. While most works were only able to find this placebo effect for subjective metrics, Kloft et al. [60] observed a significant objective performance improvement in a study when comparing participants using an active sham AI with participants using an inactive sham AI. As such, AI-enabled support is seen as something desirable.

Secondly, studies in various contexts have found an impact of perception of difficulty on test performance [75–77]. High difficulty leads to dissatisfaction, while perceived low difficulty can be motivating. Labeling a tip as LLM-generated actually had the impact of lowering the perceived difficulty of the question (see Sect. 3.1) compared to labeling a tip as expert-created, possibly reducing reluctance and increasing motivation, resulting in increased performance.

#### 4.3 Opportunities and challenges of LLMs for teaching

While our work has revealed design pitfalls to avoid when working with LLMs, such as a possible placebo effect as well as lengthy and leading tips (*RQ2*), there is a real opportunity to use LLMs for generating content about the basics of quantum computing. In general, LLM-generated tips have performed equally well, if not better, than their expert-created counterparts. Especially in a direct comparisons, LLM-generated tips were favored and conveyed the concepts for students to understand and learn better than expert-created tips.

Whilst supporting educators for tip creation, we believe that LLMs can shine specifically when used to individualize content for education. Recent works have already successfully

used LLMs to generate math questions and answers based on a student's interests [78] through prompt engineering [66]. It can be a powerful tool strongly dictating the quality of the generated content. If done right, prompt engineering can solve some of the design caveats of LLM-generated tips we found in our work, such as reducing the generated tip length and minimizing the occurrence of leading tips.

Ultimately, there will always be a trade-off with respect to correctness and quality of LLM-generated tips. We made sure that the LLM-generated tips in our study were correct, which was also confirmed by participants. However, there remains a risk that LLM-generated tips might be incorrect or misleading, especially for more complex questions [32]. Although we can reduce this risk, there is no way to ensure that each tip an LLM generates will be correct. While this is also true for expert-created tips, there is an interesting discrepancy regarding acceptable error rates for humans and artificial intelligence [79], indicating a general mistrust of AI. Consequently, there is a need for formal validation of LLM-produced content. While methods to validate AI models exist, validating them is a complex task requiring novel methods [80].

Especially in the domain of education, misinformation and incorrect content are detrimental for the learner's progress [81]. As such, we advise for a step-by-step integration of LLM-generated content for education, starting with human-in-the-loop evaluations [82], applications in low-risk scenarios (such as tips) and potentially applying LLM judges [67] as a secondary safety net.

#### 4.4 Limitations

Our results are limited to our study setup. They only show that LLM-generated tips can be used instead of expert-created tips for basic quantum computing exercises using GPT4 with a RAG-like approach to tip-creation. The results may change for other LLMs, creating tips for different domains, or increasing the question complexity. Before applying this method in practice, experts should ensure that the generated tips are of adequate quality for the questions at hand. Even using a RAG-like approach, the LLM might hallucinate, leading to wrong or misleading tips. This underlines the importance of quality control and formal verification.

Ensuring an equal level of quantum computing knowledge for our participants was challenging. In this work, we used the QUIKSTART 2024 winterschool to ensure that all participants had adequate knowledge of the topic and also used it as distribution platform, ensuring a consistent experience level among participants. Bigger cohorts can potentially be achieved by integrating the study within dedicated study programs.

#### 5 Conclusion

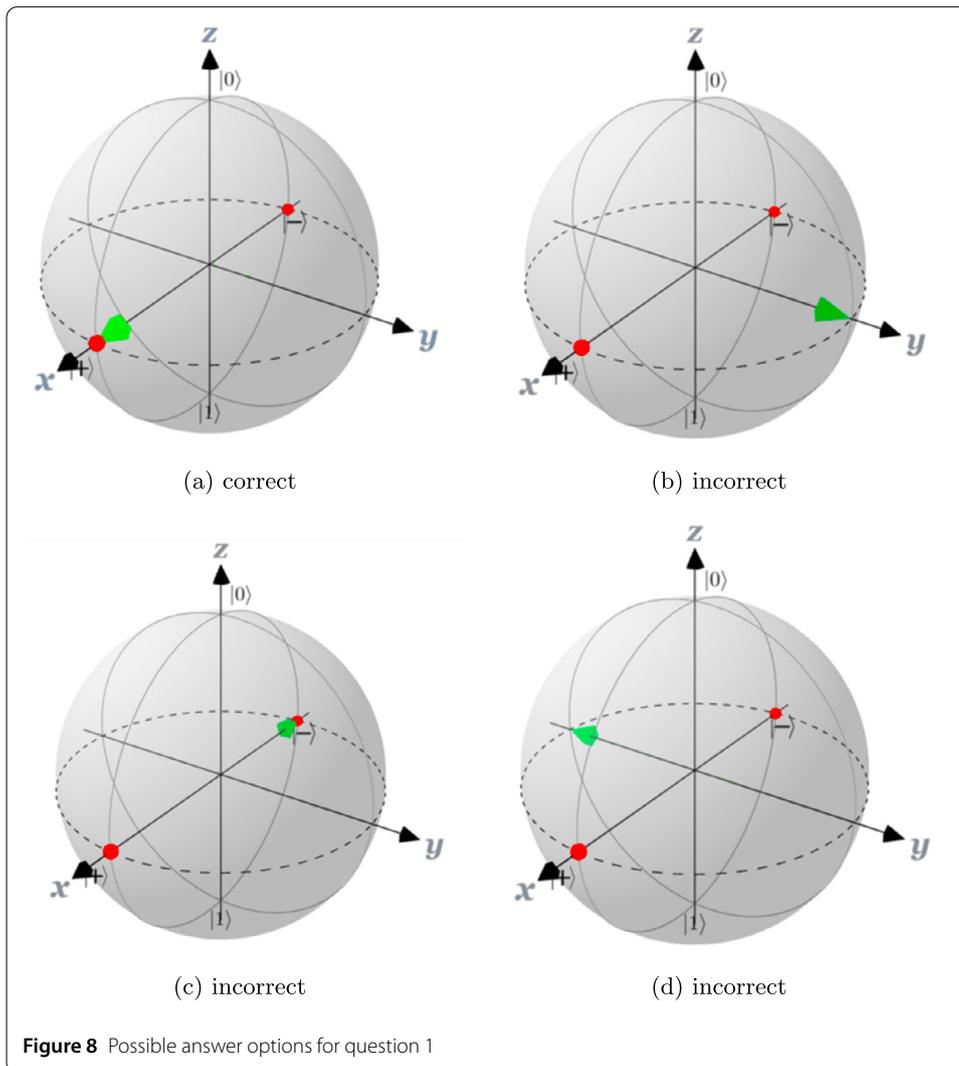
LLM-generated content has permeated almost every area in our everyday lives. In this work, we investigated the potential of LLMs to generate adequate educational content. In particular, we focused on how LLMs can support — not substitute — educators in domains with low educator-to-student ratios. Our results are promising, highlighting that LLMs are capable of producing high quality tips on par with expert-created tips, albeit are subject to design considerations including a possible placebo effect of AI and a tendency to generate leading tips. We discuss how these findings can support teaching methods pointing towards LLM-supported smart education, and expose challenges still to overcome when employing LLM-generated content.

**Appendix A: Multiple choice questions**

**A.1 Question 1**

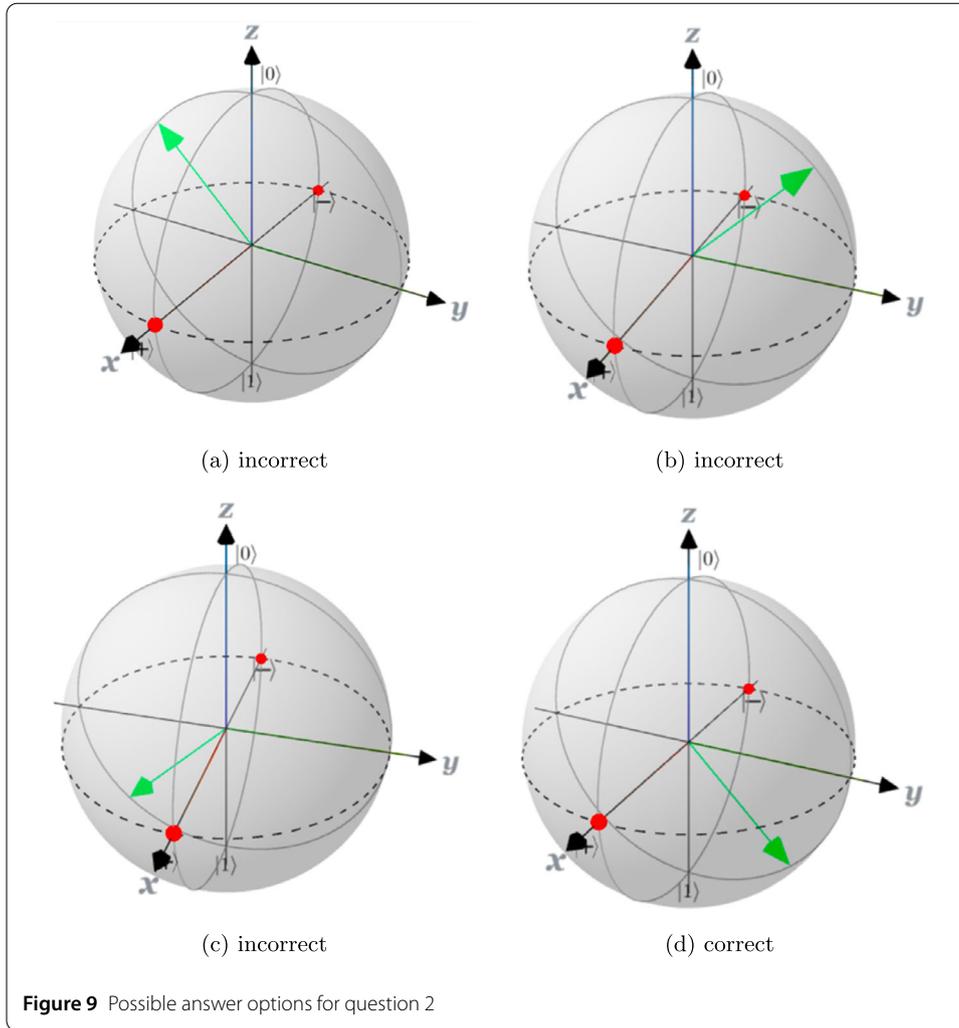
What is the following state on the Bloch sphere?

$$|\psi\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$$



**A.2 Question 2**

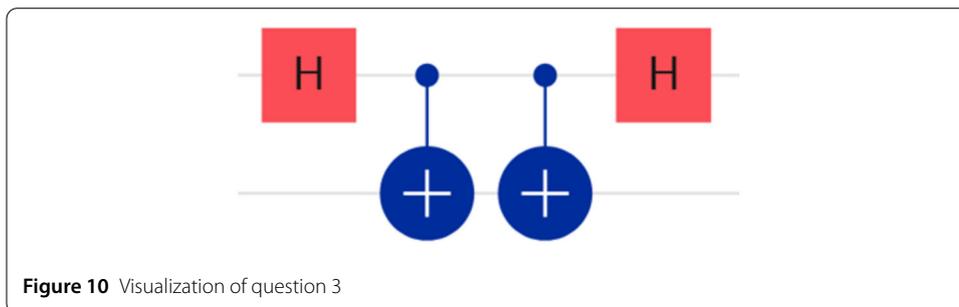
What is the following state on the Bloch sphere?  $|\psi\rangle = \frac{1}{2}|0\rangle + \frac{\sqrt{3}}{2}e^{i\frac{\pi}{2}}|1\rangle$



### A.3 Question 3

What is equal to the following two-qubit operation:

$$(H \otimes I) \text{CNOT} \text{CNOT} (H \otimes I)?$$

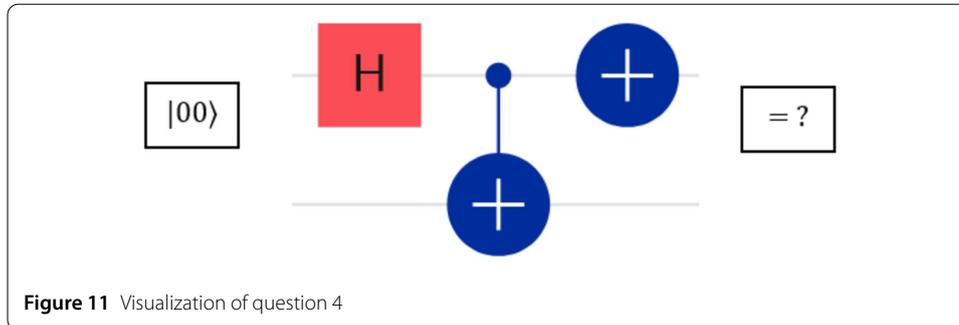


- a)  $I \otimes I$  (correct)
- b)  $H \otimes I$
- c) CNOT
- d)  $(H \otimes I) \text{CNOT}$

#### A.4 Question 4

Which Bell-state is the outcome of the following operation on the  $|00\rangle$  state:

$$(X \otimes I) \text{CNOT}((H \otimes I)|00\rangle)?$$



- a)  $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$
- b)  $\frac{1}{\sqrt{2}}(|00\rangle - |11\rangle)$
- c)  $\frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$  (correct)
- d)  $\frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$

### Appendix B: Tips

#### B.1 Question 1

*Expert tip:* Realize you are in an equal superposition: Compute the probability of outcome 0 and 1. On the Bloch sphere an outcome is probable if the state vector is closer to the pole corresponding to the outcome. The definition of the relative phase, relative to the  $x$ -axis on the Bloch sphere, gives you the representation on the Bloch sphere.

*LLM tip:* To understand where the given state  $|\psi\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$  is on the Bloch sphere, remember that each point on the Bloch sphere represents a possible quantum state of a qubit. The state vector  $|0\rangle$  corresponds to the north pole of the Bloch sphere, and  $|1\rangle$  corresponds to the south pole. A state like  $|\psi\rangle$  which is an equal superposition of  $|0\rangle$  and  $|1\rangle$  with no complex phase (meaning both coefficients are real and positive) will be on the equator of the Bloch sphere. Look for a state vector that is equally between  $|0\rangle$  and  $|1\rangle$  on the equator, without any phase difference indicated by complex numbers.

#### B.2 Question 2

*Expert tip:* On the Bloch sphere, the amplitude of state  $|0\rangle$  is related to the sine of  $\vartheta$ , the polar angle. The amplitude of  $|1\rangle$  is related to the cosine of  $\vartheta$  but one also needs to account for the relative phase, which is represented by the azimuthal angle  $\varphi$ .

*LLM tip:* When trying to identify a state on the Bloch sphere, you'll want to consider both the magnitude and the phase of the qubit state coefficients. For the state  $|\psi\rangle = \frac{1}{2}|0\rangle + \frac{\sqrt{3}}{2}e^{i\frac{\pi}{2}}|1\rangle$ , note the coefficients in front of  $|0\rangle$  and  $|1\rangle$ . The magnitude squared of these coefficients will give you the probabilities of measuring each state. The phase factor  $e^{i\frac{\pi}{2}}$  corresponds to a phase of  $\frac{\pi}{2}$  radians or 90 degrees. On the Bloch sphere, states are represented as points on the surface of the sphere, where the angles theta ( $\theta$ ) and phi

( $\phi$ ) determine the position. Theta is the angle from the positive z-axis and phi is the angle from the positive x-axis in the x-y plane. To find the correct picture, look for the one where the vector representing the qubit state has these properties.

### B.3 Question 3

*Expert tip:* Many quantum gates are their own inverse. Think of the outcome of applying the CX-gate twice. Think of the outcome of applying the Hadamard gate twice.

*LLM tip:* Let's break down the operation given in your question, step by step, using concepts from the document. The Hadamard gate (H) puts a qubit into a superposition state, meaning it creates a combination of all possible states. The Identity gate (I) does nothing to the qubit it acts upon. The CNOT gate is a conditional gate that flips the target qubit if the control qubit is in the state  $|1\rangle$ .

Here's a hint to think about how each component of the operation affects the state of the qubits:

- Start by understanding the effect of the Hadamard gate on a single qubit and how the Identity gate leaves the other qubit unchanged.
- Then, consider what happens when a CNOT gate acts on two qubits, specifically how it changes the state of the qubits based on the condition mentioned above.
- Finally, think about the sequence of operations and their reversibility, especially how applying the same operations in reverse order might affect the qubits' state.

### B.4 Question 4

*Expert tip:* Recall the bell state generator (Hadamard + CX-gate). The additional X-gate in the circuit changes the state of one qubit in the outcome.

*LLM tip:* To understand which Bell state is produced by the given quantum circuit, let's break down the operation step by step:

- The Hadamard gate ( $H \otimes I$ ) acts on the  $|00\rangle$  state. The Hadamard gate puts the first qubit into a superposition state, transforming  $|0\rangle$  into  $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$ .
- The CNOT gate acts on this new state. The CNOT gate is a conditional gate that flips the second qubit if the first qubit is in the  $|1\rangle$  state.
- Finally, the Pauli-X gate ( $X \otimes I$ ) acts only on the first qubit, flipping it from  $|0\rangle$  to  $|1\rangle$  and vice versa.

Your task is to apply these operations successively, keeping in mind the effects of each gate. The Hadamard gate creates a superposition, the CNOT entangles qubits by flipping the second based on the first, and the X gate flips the qubit from  $|0\rangle$  to  $|1\rangle$  or  $|1\rangle$  to  $|0\rangle$ . The final state will be a Bell state, an entangled state of two qubits. Look at the transformation of each qubit at each step, and this will guide you to the correct Bell state.

## Appendix C: Tip generation

We used the following prompt to generate our tips with GPT-4.

You are a passionate and patient quantum-physics teacher. Based on the information given in the pdf, give the struggling student a hint that helps him answer the following

question, without giving away the answer. Be more specific and formulate your tip as if you were talking to a novice. Be as concise as possible.

{Question + answers + correct answer}

Also adding the pdf of the quantum computing basics used in QUICKSTART 2024 and the images of the questions as context.

## Appendix D: Additional results

**Table 7** The ANOVA results of the main study for the participants' score

Measure	Factor	Df	Df.res	F value	p-value	Signif.	$\eta^2$ (partial)
score	tip creator	1	42	0.34	0.56	ns	8.06e-03
score	tip label	1	42	5.89	0.02	*	0.12

**Table 8** The ANOVA results of the main survey for the perceived question difficulty

Measure	Factor	Df	Df.res	F value	p-value	Signif.	$\eta^2$ (partial)
difficulty	tip creator	1	42	0.14	0.71	ns	3.44e-03
difficulty	tip label	1	42	3.08	0.09	ns	0.07

**Table 9** The ANOVA results of the main survey for the custom questions

Measure	Factor	Df	Df.res	F value	p-value	Signif.	$\eta^2$ (partial)
correctness	tip creator	1	42	1.60	0.21	ns	0.04
correctness	tip label	1	42	0.22	0.64	ns	5.14e-03
helpfulness	tip creator	1	42	0.22	0.64	ns	5.30e-03
helpfulness	tip label	1	42	0.43	0.52	ns	0.01
quality	tip creator	1	42	0.58	0.45	ns	0.01
quality	tip label	1	42	0.03	0.86	ns	7.81e-04

**Table 10** The ANOVA results of the main survey for the custom questions

Measure	Factor	Df	Df.res	F value	p-value	Signif.	$\eta^2$ (partial)
difficulty	tip creator	1	42	0.14	0.71	ns	3.44e-03
difficulty	tip label	1	42	3.08	0.09	ns	0.07

### Abbreviations

AI, Artificial Intelligence; ANOVA, Analysis of Variance; ART, Aligned Rank Transform; DFKI, German Research Center for Artificial Intelligence; LLM, Large Language Model; QC, Quantum Computing; QUIP, Quantum-Initiative Rhineland-Palatinate; RAG, Retrieval-Augmented Generation; STEM, Science, Technology, Engineering & Mathematics.

### Acknowledgements

The authors thank all the study participants and all guest speakers of the QUIP quantum machine learning winter school QUIKSTART 2024: Antonio Macaluso, Nico Piatkowski, Carsten Blank, Bob Coecke, Shai Machnes, Iraitz Montalban, Caitlin Jones, and Sebastian Stern. Also, we are grateful to Antonio Macaluso, Christina Jörg, Alexey Bochkarev, Nikolaos Palaiodimopoulos, and Thomas Niederprüm for their additional expertise on the topic. Finally, we would like to thank the helpers beyond the scientific program, especially Verena Paff, Jörn Gukenbiehl and Luis Sanchez. A.W., H.O., P.L., V.B. and M.K-E. acknowledge support from the Quantum Initiative Rhineland-Palatinate (QUIP). A.W., H.O., P.L., and M.K-E. acknowledge support from the Research Initiative Quantum Computing for Artificial Intelligence (QC-AI). J.B., A.W., P.L., and M.K-E. acknowledge support by the QuanTUK project at the RPTU in Kaiserslautern, supported by the Federal Ministry of Education and Research (BMBF) under grant number FKZ13N15995. L.K., J.K., and P.L. acknowledge support by the KI4TUK project at the RPTU in Kaiserslautern, supported by the Federal Ministry of Education and Research (BMBF) under grant number 16DHBKI058.

### Author contributions

The main study was performed during the QUIP winter school QUIKSTART 2024 at RPTU in Kaiserslautern Germany and was mainly organized by V.B. and M.K-E.. QUIKSTART 2024 was mainly possible due to contributed lectures, expertise and supervision of I.G., A.G., S.B., S.S., A.M., A.C.M., A.W., H.O., and P.L.. L.K., and J.K. designed the studies, J.B. and M.K-E. added quantum computing related questions. L.K., J.B., J.K. and M.K-E. finalized the design. I.G. and A.G. provided expert tips and presented the corresponding elementary lectures on quantum computing. S.M., and S.S. provided lectures on advanced quantum and machine learning. A.M., A.C.M, P.L., and V.B. supervised the elementary and advanced quantum computing, machine learning lectures and scripts. L.K., J.B., I.G., A.G., V.B. and M.K-E. collected the data. L.K. and J.K. evaluated and analyzed the data. L.K., J.B., J.K., and M.K-E. discussed the data. L.K., J.B., J.K., and M.K-E. wrote the first draft of the manuscript. P.L., and J.K. supervised the study. All authors have read, reviewed and approved the final manuscript.

### Funding

Open Access funding enabled and organized by Projekt DEAL.

### Data Availability

The datasets generated and analyzed during the current study are available in the following repository: <https://doi.org/10.5281/zenodo.12723463>.

### Declarations

#### Ethics approval

Ethics approval was granted by the Ethics Department of the German Research Center for Artificial Intelligence (DFKI).

#### Competing interests

The authors declare no competing interests.

#### Author details

<sup>1</sup>Embedded Intelligence, German Research Center for Artificial Intelligence, Kaiserslautern, 67663, Germany.

<sup>2</sup>Department of Computer Science and Research Initiative QC-AI, RPTU Kaiserslautern-Landau, Kaiserslautern, 67663,

Germany. <sup>3</sup>Department of Physics and Research Center OPTIMAS, RPTU Kaiserslautern-Landau, Kaiserslautern, 67663,

Germany. <sup>4</sup>Fraunhofer Institute for Industrial Mathematics, Kaiserslautern, 67663, Germany. <sup>5</sup>Fraunhofer Center for

Maritime Logistics and Services, Hamburg, 21073, Germany.

Received: 24 July 2024 Accepted: 24 February 2025 Published online: 04 March 2025

### References

1. MacFarlane AGJ, Dowling JP, Milburn GJ. Quantum technology: the second quantum revolution. *Philos Trans R Soc Lond A, Math Phys Eng Sci.* 2003;361(1809):1655–74. <https://doi.org/10.1098/rsta.2003.1227>.
2. Roberson TM. On the social shaping of quantum technologies: an analysis of emerging expectations through grant proposals from 2002–2020. *Minerva.* 2021;59(3):379–97. <https://doi.org/10.1007/s11024-021-09438-5>.
3. Scheidsteger T, Haunschild R, Ettl C. Historical roots and seminal papers of quantum technology 2.0. *NanoEthics.* 2022;16(3):271–96. <https://doi.org/10.1007/s11569-022-00424-z>.
4. Roberson T, Raman S, Leach J, Vilkins S. Assessing the journey of technology hype in the field of quantum technology. *Z Technikfolgenabschätzung Theor Praxis.* 2023;32(3):17–21. <https://doi.org/10.14512/tatup.32.3.17>.
5. Arute F, Arya K, Babbush R, Bacon D, Bardin JC, Barends R, Biswas R, Boixo S, Brandao FGSL, Buell DA, Burkett B, Chen Y, Chen Z, Chiaro B, Collins R, Courtney W, Dunswoth A, Farhi E, Foxen B, Fowler A, Gidney C, Giustina M, Graff R, Guerin K, Habegger S, Harrigan MP, Hartmann MJ, Ho A, Hoffmann M, Huang T, Humble TS, Isakov SV, Jeffrey E, Jiang Z, Kafri D, Kechedzhi K, Kelly J, Klimov PV, Knysh S, Korotkov A, Kostritsa F, Landhuis D, Lindmark M, Lucero E, Lyakh D, Mandrà S, McClean JR, McEwen M, Megrant A, Mi X, Michielsen K, Mohseni M, Mutus J, Naaman O, Neeley M, Neill C, Niu MY, Ostby E, Petukhov A, Platt JC, Quintana C, Rieffel EG, Roushan P, Rubin NC, Sank D, Satzinger KJ, Smelyanskiy V, Sung KJ, Trevithick MD, Vainsencher A, Villalonga B, White T, Yao ZJ, Yeh P, Zalcman A, Neven H, Martinis JM. Quantum supremacy using a programmable superconducting processor. *Nature.* 2019;574(7779):505–10. <https://doi.org/10.1038/s41586-019-1666-5>.
6. Kim Y, Eddins A, Anand S, Wei KX, Berg E, Rosenblatt S, Nayfeh H, Wu Y, Zaletel M, Temme K, Kandala A. Evidence for the utility of quantum computing before fault tolerance. *Nature.* 2023;618(7965):500–5. <https://doi.org/10.1038/s41586-023-06096-3>.
7. Aboy M, Minssen T, Kop M. Mapping the patent landscape of quantum technologies: patenting trends, innovation and policy implications. *Int Rev Intellect Prop Compet Law.* 2022;53(6):853–82. <https://doi.org/10.1007/s40319-022-01209-3>.

8. van Deventer O, Spethmann N, Loeffler M, Amoretti M, van den Brink R, Bruno N, Comi P. Towards European standards for quantum technologies. *EPJ Quantum Technol.* 2022;9(1):33. <https://doi.org/10.1140/epjqt/s40507-022-00150-1>.
9. Mazur T. Quantum Technologies and Standardization. Strategic Perspectives on Cybersecurity Management and Public Policies. 34.
10. Jenet A. Standards4quantum: making quantum technology ready for industry. Putting science into standards. In: Putting science into standards (May 1, 2020). JRC conference and workshop report; 2020.
11. Sood SK, Pooja: quantum computing review: a decade of research. *IEEE Trans Eng Manag.* 2024;71:6662–76. <https://doi.org/10.1109/TEM.2023.3284689>.
12. Coccia M, Roshani S. New directions in quantum computing technology. *J Knowl Econ SSRN.* 2024. <https://doi.org/10.2139/ssrn.4871417>.
13. Houssein EH, Abohashima Z, Elhoseny M, Mohamed WM. Machine learning in the quantum realm: the state-of-the-art, challenges, and future vision. *Expert Syst Appl.* 2022;194:116512. <https://doi.org/10.1016/j.eswa.2022.116512>.
14. Preskill J. Quantum computing in the NISQ era and beyond. *Quantum.* 2018;2:79. <https://doi.org/10.22331/q-2018-08-06-79>.
15. Bova F, Goldfarb A, Melko RG. Commercial applications of quantum computing. *EPJ Quantum Technol.* 2021;8(1):2.
16. Bayerstadler A, Becquin G, Binder J, Botter T, Ehm H, Ehmer T, Erdmann M, Gaus N, Harbach P, Hess M, et al. Industry quantum computing applications. *EPJ Quantum Technol.* 2021;8(1):25.
17. Greinert F, Müller R. European Competence Framework for Quantum Technologies. Zenodo. 2024. <https://doi.org/10.5281/zenodo.10976836>.
18. Sáez-Ortuño L, Huertas-García R, Forgas-Coll S, Sánchez-García J, Puertas-Prats E. Quantum computing for market research. *J Innov Knowl.* 2024;9(3):100510.
19. Carberry D, Nourbakhsh A, Karon J, Jones MN, Jadidi M, Shahriari K, Beenfeldt C, Andersson MP, Mansouri SS. Building knowledge capacity for quantum computing in engineering education. In: Türkay M, Gani R, editors. 31st European symposium on computer aided process engineering. Computer aided chemical engineering. vol. 50. Amsterdam: Elsevier; 2021. p. 2065–70. <https://doi.org/10.1016/B978-0-323-88506-5.50319-3>.
20. Seegerer S, Michaeli T, Romeike R. Quantum computing as a topic in computer science education. In: Proceedings of the 16th Workshop in Primary and Secondary Computing Education. WiPSCe '21. New York: Association for Computing Machinery; 2021. <https://doi.org/10.1145/3481312.3481348>.
21. Goorney S, Bley J, Heusler S, Sherson J. A Framework for Curriculum Transformation in Quantum Information Science and Technology Education. 2023. [arXiv:2308.10371](https://arxiv.org/abs/2308.10371).
22. Goorney S, Sarantinou M, Sherson J. The quantum technology open master: widening access to the quantum industry. *EPJ Quantum Technol.* 2024;11(1):7. <https://doi.org/10.1140/epjqt/s40507-024-00217-1>.
23. Meyer JC, Passante G, Pollock SJ, Wilcox BR. Introductory quantum information science coursework at us institutions: content coverage. *EPJ Quantum Technol.* 2024;11(1):16. <https://doi.org/10.1140/epjqt/s40507-024-00226-0>.
24. Angara PP, Stege U, MacLean A. Quantum computing for high-school students an experience report. In: 2020 IEEE international conference on Quantum Computing and Engineering (QCE). IEEE; 2020. p. 323–9.
25. Angara PP, Stege U, MacLean A, Müller HA, Markham T. Teaching quantum computing to high-school-aged youth: a hands-on approach. *IEEE Trans Quantum Eng.* 2021;3:1–15.
26. Walsh JA, Fenech M, Tucker DL, Riegle-Crumb C, La Cour BR. Piloting a full-year, optics-based high school course on quantum computing. *Phys Educ.* 2021;57(2):025010.
27. Merzel A, Bitzenbauer P, Krijtenburg-Lewerissa K, Stadermann K, Andreotti E, Anttila D, Bondani M, Chiofalo MLM, Faletić S, Frans R, Goorney S, Greinert F, Jurčić L, Koupilová Z, Malgieri M, Müller R, Onorato P, Pospiech G, Ubben M, Woitzik A, Pol H. The core of secondary level quantum education: a multi-stakeholder perspective. *EPJ Quantum Technol.* 2024;11(1):27. <https://doi.org/10.1140/epjqt/s40507-024-00237-x>.
28. Langston J. In a historic milestone, Azure Quantum demonstrates formerly elusive physics needed to build scalable topological qubits. <https://news.microsoft.com/source/features/innovation/azure-quantum-majorana-topological-qubit/>. Accessed: 22.07.2024.
29. Quantum Supremacy. The Race Among Google, Microsoft, Intel, and IBM. [https://www.defianceetfs.com/quantum-supremacy-the-race-among-google-microsoft-intel-and-ibm/#\\_ftn3](https://www.defianceetfs.com/quantum-supremacy-the-race-among-google-microsoft-intel-and-ibm/#_ftn3). Accessed: 22.07.2024.
30. Kasneci E, Seßler K, Küchemann S, Bannert M, Dementieva D, Fischer F, Gasser U, Groh G, Günemann S, Hüllermeier E, et al. Chatgpt for good? On opportunities and challenges of large language models for education. *Learn Individ Differ.* 2023;103:102274.
31. Ahmad N, Murugesan S, Kshetri N. Generative artificial intelligence and the education sector. *Computer.* 2023;56(6):72–6. <https://doi.org/10.1109/MC.2023.3263576>.
32. Krupp L, Steinert S, Kiefer-Emmanouilidis M, Avila KE, Lukowicz P, Kuhn J, Küchemann S, Karolus J. Unreflected acceptance—investigating the negative consequences of chatgpt-assisted problem solving in physics education. In: HHAI 2024: Hybrid Human AI systems for the social good. Amsterdam: IOS Press; 2024. p. 199–212.
33. Gregorcic B, Pendrill A-M. Chatgpt and the frustrated Socrates. *Phys Educ.* 2023;58(3):035021.
34. Santos RPD. Enhancing physics learning with chatgpt, bing chat, and bard as agents-to-think-with: a comparative case study. 2023. [arXiv preprint. arXiv:2306.00724](https://arxiv.org/abs/2306.00724).
35. Krupp L, Steinert S, Kiefer-Emmanouilidis M, Avila KE, Lukowicz P, Kuhn J, Küchemann S, Karolus J. Challenges and opportunities of moderating usage of large language models in education. 2023. [arXiv preprint. arXiv:2312.14969](https://arxiv.org/abs/2312.14969).
36. Sallam M, Salim NA, Barakat M, Ala'a B. Chatgpt applications in medical, dental, pharmacy, and public health education: a descriptive study highlighting the advantages and limitations. *Narra J.* 2023;3(1).
37. Gregorcic B, Polverini G. Chatgpt-4 and the satisfied Socrates. 2024. [arXiv preprint. arXiv:2401.1987](https://arxiv.org/abs/2401.1987).
38. OpenAI. Introducing ChatGPT. 2022. <https://openai.com/blog/chatgpt>.
39. Latif E, Parasuraman R, Zhai X. PhysicsAssistant: an LLM-Powered Interactive Learning Robot for Physics Lab Investigations. 2024.
40. Wermelinger M. Using github copilot to solve simple programming problems. In: Proceedings of the 54th ACM technical symposium on computer science education. vol. 1. 2023. p. 172–8.
41. Gómez-Rodríguez C, Williams P. A confederacy of models: a comprehensive evaluation of llms on creative writing. 2023. [arXiv preprint. arXiv:2310.08433](https://arxiv.org/abs/2310.08433).

42. Thirunavukarasu AJ, Ting DSJ, Elangovan K, Gutierrez L, Tan TF, Ting DSW. Large language models in medicine. *Nat Med*. 2023;29(8):1930–40.
43. Gemini Team Google, Anil R, Borgeaud S, Alayrac J-B, Yu J, et al. Gemini: a family of highly capable multimodal models. 2024.
44. Meta. Introducing Llama 3. 2024. <https://ai.meta.com/blog/meta-llama-3/>.
45. Microsoft. Microsoft and OpenAI extend partnership. 2023. <https://blogs.microsoft.com/blog/2023/01/23/microsoftandopenaiextendpartnership/>.
46. Parmar J, Prabhumoye S, Jennings J, Patwary M, Subramanian S, Su D, Zhu C, Narayanan D, Jhunjhunwala A, Dattagupta A, et al. Nemotron-4 15b technical report. 2024. arXiv preprint. [arXiv:2402.16819](https://arxiv.org/abs/2402.16819).
47. Templeton A, Conerly T, Marcus J, Lindsey J, Bricken T, Chen B, Pearce A, Citro C, Ameisen E, Jones A, Cunningham H, Turner NL, McDougall C, MacDiarmid M, Freeman CD, Summers TR, Rees E, Batson J, Jermyn A, Carter S, Olah C, Henighan T. Scaling monosemanticity: extracting interpretable features from Claude 3 sonnet. *Transform Circuits Thread*. 2024.
48. OpenAI, Achiam J, Adler S, Agarwal S, Ahmad L, et al. GPT-4 Technical Report. 2024.
49. Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, Küttler H, Lewis M, Yih W-t, Rocktäschel T, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Adv Neural Inf Process Syst*. 2020;33:9459–74.
50. Gao Y, Xiong Y, Gao X, Jia K, Pan J, Bi Y, Dai Y, Sun J, Wang M, Wang H. Retrieval-Augmented Generation for Large Language Models: a Survey. 2024.
51. Yan L, Sha L, Zhao L, Li Y, Martinez-Maldonado R, Chen G, Li X, Jin Y, Gašević D. Practical and ethical challenges of large language models in education: a systematic scoping review. *Br J Educ Technol*. 2024;55(1):90–112. <https://doi.org/10.1111/bjet.13370>.
52. Belland BR. *Instructional scaffolding in STEM education: strategies and efficacy evidence*. London: Springer; 2017.
53. Schartel SA. Giving feedback – an integral part of education. *Best Pract Res, Clin Anaesthesiol*. 2012;26(1):77–87. <https://doi.org/10.1016/j.bpa.2012.02.003>. Challenges in Anaesthesia Education.
54. Bajaj R, Sharma V. Smart education with artificial intelligence based determination of learning styles. *Proc Comput Sci*. 2018;132:834–42. <https://doi.org/10.1016/j.procs.2018.05.095>.
55. Gan W, Qi Z, Wu J, Lin JC-W. Large language models in education: vision and opportunities. In: 2023 IEEE international conference on Big Data (BigData). 2023. p. 4776–85. <https://doi.org/10.1109/BigData59044.2023.10386291>.
56. Zhu Z-T, Yu M-H, Riezebos P. A research framework of smart education. *Smart Learn Environ*. 2016;3(1):4. <https://doi.org/10.1186/s40561-016-0026-2>.
57. Weidinger L, Uesato J, Rauh M, Griffin C, Huang P-S, Mellor J, Glaese A, Cheng M, Balle B, Kasirzadeh A, Biles C, Brown S, Kenton Z, Hawkins W, Stepleton T, Birhane A, Hendricks LA, Rimell L, Isaac W, Haas J, Legassick S, Irving G, Gabriel I. Taxonomy of risks posed by language models. In: Proceedings of the 2022 ACM conference on Fairness, Accountability, and Transparency. FAccT '22. New York: Association for Computing Machinery; 2022. p. 214–29. <https://doi.org/10.1145/3531146.3533088>.
58. Milano S, McGrane JA, Leonelli S. Large language models challenge the future of higher education. *Nat Mach Intell*. 2023;5(4):333–4. <https://doi.org/10.1038/s42256-023-00644-2>.
59. Kosch T, Welsch R, Chuang L, Schmidt A. The placebo effect of artificial intelligence in human–computer interaction. *ACM Trans Comput-Hum Interact*. 2023;29(6):1–32.
60. Kloft AM, Welsch R, Kosch T, Villa S. “AI enhances our performance, I have no doubt this one will do the same”: the placebo effect is robust to negative descriptions of AI. In: Proceedings of the 2024 CHI conference on human factors in computing systems. CHI '24. New York: ACM; 2024. <https://doi.org/10.1145/3613904.3642633>.
61. Hu P, Li Y, Mong RSK, Singh C. Student understanding of the Bloch sphere. *Eur J Phys*. 2024;45(2):025705. <https://doi.org/10.1088/1361-6404/ad2393>.
62. Boynton MH, Portnoy DB, Johnson BT. Exploring the ethics and psychological impact of deception in psychological research. *IRB*. 2013;35(2):7.
63. Girden ER. ANOVA: repeated measures. vol. 84. London: Sage; 1992.
64. Wobbrock JO, Findlater L, Gergle D, Higgins JJ. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In: Proceedings of the SIGCHI conference on human factors in computing systems. 2011. p. 143–6.
65. Jung LB, Gudera JA, Wiegand TL, Allmendinger S, Dimitriadis K, Koerte IK. Chatgpt passes German state examination in medicine with picture questions omitted. *Dtsch Arztebl Int*. 2023;120(21–22):373.
66. Lee U, Jung H, Jeon Y, Sohn Y, Hwang W, Moon J, Kim H. Few-shot is enough: exploring chatgpt prompt engineering method for automatic question generation in English education. *Educ Inf Technol*. 2023:1–33.
67. Zheng L, Chiang W-L, Sheng Y, Zhuang S, Wu Z, Zhuang Y, Lin Z, Li Z, Li D, Xing E, et al. Judging Llm-as-a-judge with mt-bench and chatbot arena. *Adv Neural Inf Process Syst*. 2024;36.
68. Brown TB, 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 DM, Wu J, Winter C, Hesse C, Chen M, Sigler E, Litwin M, Gray S, Chess B, Clark J, Berner C, McCandlish S, Radford A, Sutskever I, Amodei D. Language Models are Few-Shot Learners. 2020. [arXiv:2005.14165](https://arxiv.org/abs/2005.14165).
69. Wei J, Wang X, Schuurmans D, Bosma M, Xia F, Chi E, Le QV, Zhou D, et al. Chain-of-thought prompting elicits reasoning in large language models. *Adv Neural Inf Process Syst*. 2022;35:24824–37.
70. Krathwohl DR. A revision of bloom’s taxonomy: an overview. *Theory Pract*. 2002;41(4):212–8.
71. Erickson HL. *Concept-based curriculum and instruction: teaching beyond the facts*. Thousand Oaks: Corwin Press; 2002.
72. Tonmoy S, Zaman S, Jain V, Rani A, Rawte V, Chadha A, Das A. A comprehensive survey of hallucination mitigation techniques in large language models. 2024. arXiv preprint. [arXiv:2401.01313](https://arxiv.org/abs/2401.01313).
73. Pithers RT, Soden R. Critical thinking in education: a review. *Educ Res*. 2000;42(3):237–49.
74. Villa S, Kosch T, Grelka F, Schmidt A, Welsch R. The placebo effect of human augmentation: anticipating cognitive augmentation increases risk-taking behavior. *Comput Hum Behav*. 2023;146:107787.
75. Hong E. Test anxiety, perceived test difficulty, and test performance: temporal patterns of their effects. *Learn Individ Differ*. 1999;11(4):431–47. [https://doi.org/10.1016/S1041-6080\(99\)80012-0](https://doi.org/10.1016/S1041-6080(99)80012-0).

76. Kukla A. Performance as a function of resultant achievement motivation (perceived ability) and perceived difficulty. *J Res Pers.* 1974;7(4):374–83. [https://doi.org/10.1016/0092-6566\(74\)90059-2](https://doi.org/10.1016/0092-6566(74)90059-2).
77. Sarrazin P, Roberts G, Cury F, Biddle S, Famose J-P. Exerted effort and performance in climbing among boys: the influence of achievement goals, perceived ability, and task difficulty. *Res Q Exerc Sport.* 2002;73(4):425–36.
78. Christ BR, Kropko J, Hartvigsen T. Mathwell: Generating educational math word problems at scale. 2024. arXiv preprint. [arXiv:2402.15861](https://arxiv.org/abs/2402.15861).
79. Lenskjold A, Nybing JU, Trampedach C, Galsgaard A, Brejnebøl MW, Raaschou H, Rose MH, Boesen M. Should artificial intelligence have lower acceptable error rates than humans? *BJR Open.* 2023;5(1):20220053.
80. Huang X, Ruan W, Huang W, Jin G, Dong Y, Wu C, Bensalem S, Mu R, Qi Y, Zhao X, et al. A survey of safety and trustworthiness of large language models through the lens of verification and validation. *Artif Intell Rev.* 2024;57(7):175.
81. Holmes V-L, Miedema C, Nieuwkoop L, Haugen N. Data-driven intervention: correcting mathematics students' misconceptions, not mistakes. *Math Educ.* 2013;23(1).
82. Amirizani M, Yao J, Lavergne A, Okada ES, Chadha A, Roosta T, Shah C. Developing a framework for auditing large language models using human-in-the-loop. 2024. arXiv preprint. [arXiv:2402.09346](https://arxiv.org/abs/2402.09346).

### Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Submit your manuscript to a SpringerOpen<sup>®</sup> journal and benefit from:**

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

---

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)

---