

SciEx: Benchmarking Large Language Models on Scientific Exams with Human Expert Grading and Automatic Grading

Tu Anh Dinh, Carlos Mullov, Leonard Bärmann, Zhaolin Li, Danni Liu, Simon Reiß, Jueun Lee, Nathan Lerzer, Jianfeng Gao, Fabian Ternava, Tobias Röddiger, Alexander Waibel, Tamim Asfour, Michael Beigl, Rainer Stiefelhagen, Carsten Dachsbacher, Klemens Böhm, Jan Niehues
Karlsruhe Institute of Technology, Karlsruhe, Germany
{firstname}.{lastname}@kit.edu

Abstract

With the rapid development of Large Language Models (LLMs), it is crucial to have benchmarks which can evaluate the ability of LLMs on different domains. One common use of LLMs is performing tasks on scientific topics, such as writing algorithms, querying databases or giving mathematical proofs. Inspired by the way university students are evaluated on such tasks, in this paper, we propose SciEx - a benchmark consisting of university computer science exam questions, to evaluate LLMs' ability on solving scientific tasks. SciEx is (1) multilingual, containing both English and German exams, and (2) multi-modal, containing questions that involve images, and (3) contains various types of freeform questions with different difficulty levels, due to the nature of university exams. We evaluate the performance of various state-of-the-art LLMs on our new benchmark. Since SciEx questions are freeform, it is not straightforward to evaluate LLM performance. Therefore, we provide human expert grading of the LLM outputs on SciEx. We show that the free-form exams in SciEx remain challenging for the current LLMs, where the best LLM only achieves 59.4% exam grade on average. We also provide detailed comparisons between LLM performance and student performance on SciEx. To enable future evaluation of new LLMs, we propose using LLM-as-a-judge to grade the LLM answers on SciEx. Our experiments show that, although they do not perform perfectly on solving the exams, LLMs are decent as graders, achieving 0.948 Pearson correlation with expert grading.

1 Introduction

In recent years, Large Language Models (LLMs) have proven their usefulness across a wide range of tasks, from conversational agents to code generation (Rajkumar et al., 2022; Abbasian et al., 2023; Liao et al., 2023). Given the fast pace of development in the field, with an increasing number of

LLMs being trained and released, it is important to have indicators of LLM performance on different domains. This can be achieved by establishing evaluation benchmarks that assess the capabilities of LLMs across diverse use cases.

One use case of LLMs is to handle scientific tasks. Some previous works have introduced benchmarks containing questions on science topics (Welbl et al., 2017; Lu et al., 2022; Gilson et al., 2022; Schubert et al., 2023; Zhang et al., 2024). However, these benchmarks are limited to multiple-choice questions. This restricts the variability of questions, such as instruction-follow ones like "write a mathematical proof for this statement ...". Additionally, it is difficult to ask certain types of questions in a multiple-choice way without including the answer in the question itself. Multiple-choice benchmarks therefore create a gap between testing and actual usage, since they only evaluate whether the LLMs choose the correct answer, whereas in real life, the users are more likely to ask open-ended questions to the LLMs. In contrast, some other works have introduced freeform question benchmarks. These works either convert multiple-choice questions to freeform questions (Bhakhavatsalam et al., 2021), or focus on a specific type of problem such as answering questions related to a paper (Dasigi et al., 2021), thus still limiting the variability of the questions.

In this paper, we introduce a new benchmark, termed SciEx (Scientific Exams), designed to evaluate this capability. Inspired by the way students are evaluated in university, we created the benchmark by evaluating the performance of LLMs on university computer science exams. SciEx's questions are in various formats, from multiple choice to open-ended, thus making it suitable to evaluate LLM's capabilities of generating free-text answers that fit the requirements of the questions. It is multilingual, containing exams in both German and English. It is multimodal, as exam questions can

also contain figures. The set of questions is a good mix of different difficulty levels since they are designed for university exams. This enables us to evaluate LLMs on different levels, and we found that stronger LLMs tend to perform better on more difficult questions.

Unlike the previous multiple-choice benchmarks, the questions in SciEx are freeform, making it non-trivial how to evaluate the LLM output. Therefore, we make use of expert grading, i.e., having the lecturers grade the LLM output the same way they would grade student answers. We also ask the experts to perform qualitative analysis of the LLM output. With expert grading, we provide a highly reliable way of evaluating LLMs, which is more reliable than previous work that uses crowd-sourced evaluation. Expert grading by lecturers also provides an opportunity to compare LLMs' performance to university student performance in a similar setting. We find that the stronger LLMs, i.e., Claude and GPT-4V, are able to outperform the student average. However, they are still far from perfect, achieving only 59% across SciEx exams.

Since new LLMs are constantly being released, we cannot fully rely on expert grading for evaluation. Therefore, we provide an automatic grading scheme by using LLM as a judge so that future LLMs can also be evaluated on SciEx. Interestingly, we find that, although LLMs do not perform too well as examinees, they perform well as graders, achieving over 0.948 Pearson correlation to expert grading in the best setting.

In summary, our contributions are as follows:

- SciEx¹ - a freeform, multimodal, multilingual benchmark consisting of university computer science exams, outputs of various LLMs on the exams, and expert grading of the LLM output.
- Detailed quantitative and qualitative analysis comparing LLM to student performance.
- Automatic grading with 0.948 Pearson correlation to expert grading

2 Related Work

General-Purpose LLM Benchmarks In order to rank different LLMs, there are several commonly used public benchmarks. For example, [Zheng et al.](#)

(2024) introduced MT-bench and Chatbot Arena. MT-bench is a multi-turn question set; and Chatbot Arena is a crowdsourced battle platform for LLMs where the users can ask their questions and vote for the better LLM answer. Another benchmark is MMLU ([Hendrycks et al., 2020](#)), which is a multitask dataset covering multiple domains such as mathematics, US history and law.

Scientific LLM Benchmarks To specifically focus on the scientific domain, previous studies have established benchmarks, such as SciQ ([Welbl et al., 2017](#)) and ScienceQA ([Lu et al., 2022](#)), which feature questions spanning various scientific subjects. More recent works have focused on benchmarking LLMs on solving exam questions on some narrow science domains such as medical ([Gilson et al., 2022](#)) or neurology ([Schubert et al., 2023](#)). M3Exam ([Zhang et al., 2024](#)), in contrast, provides exam questions to benchmark LLMs which span over multiple topics and multiple educational levels (primary, middle, and high school). However, all benchmarks mentioned above are limited to multiple-choice questions. While this simplifies the evaluation process, it does not allow us to assess the LLMs' capability to generate natural text.

Other studies have instead provided scientific benchmarks with open-ended questions. Some examples are Qasper ([Dasigi et al., 2021](#)) and ARC-DA ([Bhakhavatsalam et al., 2021](#)). However, Qasper only focuses on questions about NLP papers rather than on general computer science topics. ARC-DA is closer to our work, since it contains open-ended questions taken from science exams and quiz sources. However, these are created by converting questions that were originally multiple-choice, thus not covering certain types of typical freeform questions (e.g. those that require mathematical proofs, or long explanations).

Different from these works, SciEx is created from university computer science exams, thus naturally providing diversity in the types of questions as well as having freeform format.

Freeform Answer Evaluation Compared to benchmarks with multiple-choice questions, evaluating LLMs' performance on freeform questions is not straightforward. Similar to evaluation conditions in tasks such as machine translation or summarization, there are multiple correct answers, or multiple ways to express a correct answer for a single input. Therefore, it is insufficient to evaluate a model's output by comparing it to a

¹We release SciEx under CC BY-NC-SA 4.0 license. Code: <https://github.com/TuAnh23/SciEx>. Data: <https://huggingface.co/datasets/tuanh23/SciEx>.

gold standard answer. Ideally, in these cases, we can evaluate by human judgment. For example, the ARC-DA benchmark (Bhakhavatsalam et al., 2021) uses a crowdsourcing pipeline for evaluation. Chatbot Arena (Zheng et al., 2024) also uses crowdsourcing, where the users vote between pairs of LLM output. However, human evaluation is inherently non-scalable. Therefore, previous works have used automated metrics. Some traditional metrics such as BLEU (Papineni et al., 2002) or ROUGE (Lin, 2004) compare the model’s output to some gold-standard answer on the surface level, i.e., word matching. More advanced metrics, such as s BERTScore (Zhang* et al., 2020), BLEURT (Sellam et al., 2020), and BARTScore (Yuan et al., 2021), are model-based, thus being able to evaluate answers on the semantic level.

One recent approach is to use LLMs for evaluation, termed “LLM-as-a-judge”. Liu et al. (2023); Chiang and Lee (2023a); Zheng et al. (2024) find that, although still prone to biases, LLM-as-a-judge for textual modality has high agreement with human scoring when a strong judge LLM is used. However, when including images, Chen et al. (2024) find that the performance of LLM-as-a-judge is no longer as well correlated to human judgment. Nevertheless, LLM-as-a-judge is a promising way to perform scalable evaluation.

In our work, we make use of LLM-as-a-judge for automatic grading of LLM answers on SciEx exams, and find that they have good correlation to human expert grading on both text-only and image-related questions.

3 The SciEx Benchmark

The components of SciEx are as follows.

Univeristy Exams SciEx contains university computer science exams in a unified JSON format. The exams are taken from the following computer science courses at the Karlsruhe Institute of Technology from the 2022/2023/2024 semesters:

- Natural Language Processing (NLP)
- Advanced Artificial Intelligence (AI2)
- Deep Learning and Neural Networks (DLNN)
- Deep Learning for Computer Vision (DL4CV2)
- Human–Computer Interaction (HCI)
- Databases (DBS) for the years 2022 and 2023

- Computer Graphics (CG)
- Theoretical Foundations of Computer Science (TGI)
- Algorithms (ALGO)

Descriptions of the exams are in Appendix A. In total, SciEx contains 10 exams, among which 5 exams are in English and 7 exams are in German (some exams are provided in both languages).

There are in total 154 unique questions. Each question is annotated with (1) the maximum points that can be achieved and (2) a difficulty level among *Easy, Medium, Hard*. Most questions are provided with gold reference answers and average student performance. The detailed per-question statistics are shown in Table 1.

	Question count
Total	154
English / German*	95 / 97
Text-only / Image-related	121 / 33
Easy / Medium / Hard	51 / 71 / 32
With / Without reference	120 / 34
With / Without student average	117 / 37

*: Some questions are provided bilingually.

Table 1: Question-level statistics for SciEx.

LLM-Generated Answers SciEx contains answers produced by 7 LLMs on the exam questions. The details of the LLMs are shown in Table 2. In Table 2, only Llama3 was not used to solve the exam, since it was released at a later point of conducting this paper. In total, we obtained 1120 question-answer pairs.

Expert Grading and Automatic Grading Each question-answer pair is assigned a score by an expert. In order to guide future work to evaluate new LLMs on SciEx without relying on human expert grading, we also provide automatic grading generated by Mixtral, Llama3 and GPT4V.

3.1 Data Creation

The data creation process is as follows.

Exam Collection We collect university exams from different courses. We additionally ask the lecturers to provide us with the reference answers, the difficulty level of each question, and the average student grades on each question.

Exam Formatting We convert every exam into a unified JSON format. Each exam includes a list

	Full name	# Params	Quant.	Handle Image
Proprietary				
Claude	Claude-3-opus-20240229	-	-	yes
GPT-4v	gpt-4-vision-preview	-	-	yes
GPT-3.5	gpt-3.5-turbo-0125	-	-	no
Open source				
Llama3	MazyarPanahi/Meta-Llama-3-70B-Instruct-GGUF	70B	4 bit	no
Mixtral	Mistralai/Mixtral-8x7B-Instruct-v0.1	8x7B	5 bit	no
Qwen	Qwen/Qwen-72B	72B	2 bit	no
Mistral	Mistralai/Mistral-7B-Instruct-v0.2	7B	-	no
Llava	Llava-hf/Llava-v1.6-Mistral-7b-hf	7B	-	yes

Table 2: Details of the LLMs in consideration.

of questions, where each question includes an index, its content, and potentially path to any related images. An example is shown in Appendix B.

LLM-Generated Answers We pass the exams to the LLMs listed in Table 2 (except Llama3 due to later release), one question at a time. Questions that contain images are handled differently depending on the LLM. For the text-only LLMs, we exclude the images and only pass the question text to the models. For Llava, since it is trained to handle only 1 image at a time, we concatenate the images into one, with blank padding around the images as separators before feeding it to the model. Claude and GPT-4V can take multiple images, however, there is no pre-defined way of referencing the image within the text. In our work, we reference the image by mentioning the image caption within the question text, and add the text caption to the image.

Since the considered LLMs can only output text, for questions asking to draw on images, we ask the LLMs to describe in text what should be drawn.

The detailed prompts for LLMs to generate the answers are shown in Appendix C.

Expert Grading We then give the LLM answers back to the lecturers, who proceed with grading the LLM output the same way they would grade student answers. We anonymized the LLMs’ names in order to avoid bias during exam grading. We also build a user interface for collecting the grades (see Appendix E for more details).

With expert grading, the evaluation of the LLM output is highly reliable. Most importantly, the expert graders are generally the ones who designed the exam questions. We additionally ask the expert graders to provide their comments on the LLM output to further understand LLMs’ behaviors when solving the exams.

3.2 Automatic Grading

In addition to expert grading, we also provide automatic grading using LLM-as-a-judge, so that we can evaluate future LLMs on SciEx. We use the stronger models, i.e., Mixtral, Llama3 and GPT-4V, to conduct the grading. Given a tuple containing question, answer, and maximum score, we ask the LLMs to output a single score between 0 and the maximum. We include reference answers to the grading prompt. We ask the LLMs to provide chain-of-thought reasoning (Wei et al., 2022; Chiang and Lee, 2023b) before giving the grade. We also include examples for grading in the prompt, so-called few-shot judge (Zheng et al., 2024). Each example is a tuple consisting of a question, an answer, and the expert-provided grade. We try out different settings to select the examples, as described below.

Let’s say we want to grade Question M from Exam A, answered by Examinee X. Then the shot examples can be chosen in one of the three ways:

- **Same question:** Select examples from the same Question M from Exam A, but answered by a different Examinee Y. This mimics the real-life scenario where we use the expert resource to grade some answers of the same exam, then use it to guide the LLM graders.
- **Same exam:** select examples from a different Question N from the same Exam A, Answered by a different Examinee Y. Here the examples are in the same domain as the question-answer pair in consideration. This mimics the real-life scenario where, e.g., we have expert grading on exams of the same course from previous years to guide LLM graders.
- **Different exam:** select examples from a different Question N, from a different Exam B, Answered by a different Examinee Y. This mimics the real-life scenario where, e.g., we

have expert grading for an exam of another course to guide the LLM graders.

Examinee Y and Question N are chosen randomly. For Exam B, we opt to select the exams that do not heavily require images for simplicity in the prompt.

Intuitively, the example-selection settings above have decreasing levels of relevance to the actual grading query, but increasing easiness to collect. The detailed prompts for LLMs to grade answers are shown in Appendix D.

4 Experiments

In this section, we describe our experiments and results. For prompting the proprietary LLMs, we use their APIs, namely OpenAI² and Anthropic³. For the open-source models, we obtain model checkpoints from the Huggingface⁴ model hub. We perform inference with the LLMs using llama.cpp⁵ with the default sampling strategy. The experiments with open-sourced models are conducted on an NVIDIA RTX A6000 GPU with 48GB VRAM.

For our analysis, we consider the exam-level and question-level grades. An exam-level grade is the sum of the grades of all questions in the exam.

4.1 Quantitative Analysis

We analyze the performance of the LLMs on SciEx with expert grading. For both exam level and question level, we normalize the grade to be between 0 and 100%, since they have different scales. The normalization is done by taking the scores obtained by the examinee divided by the maximum score possible per exam/question, where the maximum scores possible are predefined by the lecturers.

We also report on the German grade scale. In the German scale, the grades range from 1.0 to 5.0, where 1.0 is the highest grade and 4.0 is the passing threshold. The detailed mapping from the scores to the German grade scale is defined by the lectures, adjusted based on the overall performance of the students taking the exams.

We compare the performance of the LLMs to students from different aspects: language, difficulty level, and modality, i.e., questions with or without images.

²<https://platform.openai.com/>

³<https://console.anthropic.com/>

⁴<https://huggingface.co/>

⁵<https://github.com/ggerganov/llama.cpp>

4.1.1 General Observations

SciEx is Challenging The performance of the LLMs on SciEx provided by expert grading is shown in Table 3. The bigger-sized LLMs (Claude, GPT-4V, GPT-3.5, Mixtral and Qwen) can achieve exam passing grades (i.e., grades that are better than 4.0 in the German scale). However, the best-performing model (Claude) only achieves 59.4% of the maximum points, which is far from perfect.

Compared to the student average, most LLMs have worse performance. Only the strongest proprietary LLMs, i.e., Claude and GPT-4V, can achieve grades that are better than the students’.

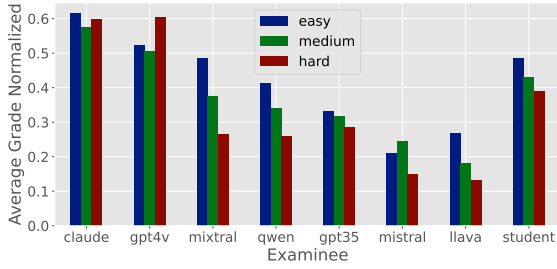
	Grade (%) ↑	German Scale ↓
Proprietary		
Claude	59.4	2.4
GPT-4V	58.2	2.5
GPT-3.5	32.8	3.9
Open source		
Mixtral	41.1	3.5
Qwen	35.4	3.7
Mistral	25.9	4.2
Llava	21.5	4.3
Student avg.	45.3	3.1

Table 3: Average performance of LLMs, exam level.


SciEx Versus Other Benchmarks The ranking of the LLMs on SciEx in Table 3 generally agrees with other public benchmarks. However, SciEx seems to be more challenging. For example, the best LLM accuracy achieved on MMLU’s various tasks is 88.8%. The best accuracy achieved on M3Exam multiple choice questions is 72.92%. Although these scores are not directly comparable, it indicates that SciEx provides a more challenging test set for future LLMs.

4.1.2 Influential Factors

Difficulty Levels Figure 1a shows the influence of the difficulty level on the examinee grades. As can be seen, the student performance aligns with the difficulty level of the questions: they perform better on easier questions. Some weaker LLMs, e.g., Mixtral, Qwen, GPT-3.5, Llava, align with the students. However, the stronger LLMs, i.e., Claude and GPT-4V, perform better on harder questions. This is an indication that difficulty levels from human perspective do not always align with LLMs’ perspective. This is also confirmed by looking at the Pearson correlations between the LLMs’ grades and the student average grades on the question level. These correlations are between 0.4 and



(a) LLMs' and students' scores.



(b) Difference between LLM scores and student scores.

Figure 1: Question-level scores grouped by difficulty.

0.6, indicating that LLM grades and the student grades are not highly correlated.

One possible explanation for the mismatch between LLMs performance and question difficulty level could be that, in some exams, there can be some “template questions”, i.e., questions that are repeated over the years, where students can just learn by heart how to systematically solve them. While this would be marked as “easy” by the lecturer, it might not be as easy for the LLMs, since the LLMs are not previously exposed to these “template questions”. Another potential explanation is that math-type *easy* questions are hard for the LLMs, while long-text *hard* questions are easy for them.

In Figure 1b, we plot the difference between LLM scores and student scores. The stronger LLMs, i.e., Claude and GPT-4V, outperform the students the most on hard questions. Weaker LLMs, on the other hand, generally fall behind students the most on hard questions. Looking at each difficulty level independently, we observe that the ranking of the LLMs changes across different levels. This aligns with the findings made by Li et al. (2024), where they show that the LLM rankings change on a subset of evaluation prompts that are artificially labeled as *hard*.

Text-only versus Image-related Questions Figure 2 shows the influence of images on the differ-




Figure 2: Difference between LLM scores and student scores, question-level, grouped by with/without images. Only Claude, GPT-4V and Llava can handle images.

ence between LLM and student scores. Recall that for the text-only LLMs, we exclude the images and only pass the question text to the models. Trivially, the text-only LLMs perform poorly on the image-related questions. The strong, multi-modal LLMs, i.e., Claude and GPT4, outperform the students on both image-related questions and text-only questions, but the performance gap is still larger for text-only questions. Llava, although can handle images, still falls behind student performance by a large margin on image-related questions. This shows that LLMs’ image-handling capability is still not as advanced as for text.

Language Figure 3 shows the influence of languages on the difference between LLM scores and student scores. When the questions are in English, all LLMs, except for GPT-3.5, outperform the student average. However, for German, either the LLMs outperform students by a smaller gap, or fall behind student performance. It can be concluded that LLMs are still superior in English than other languages like German, although German can be considered a high-resource language.




Figure 3: Difference between LLM scores and student scores, question-level, grouped by languages.

Since some models are not made to deal with images, or with languages other than English, we additionally analyze LLMs’ performance on text-

only and English-only questions. On this subset of questions, the grades obtained by the models are generally better, and more models would outperform the student average. More details can be found in Appendix F.

4.2 Qualitative Analysis

In this section, we summarize the observations made by the graders while grading the LLMs.

General Behaviours The graders observed some common behaviors made by the LLMs. Some solutions of the LLMs were good language-wise but low-quality content-wise. For students, good language usually correlates strongly with good content. The LLMs tend to output lengthy answers, since, unlike the students, they do not have a time constraint. Some LLMs even ignore when the question specifies that they should "answer briefly". There are also some failure cases, although not frequent: (1) Claude refuses to answer the question with "I apologize, but I do not feel comfortable providing answers related to ..." or (2) some LLMs get stuck in decoding loops. Sometimes, instead of answering the question, LLMs give some text that is (or seems) related to the task; rephrase the task; or describe how a task of this nature may be approached in general.

Knowledge-type Questions On some exams such as [AI2](#), [DL4CV2](#), [DLNN](#), [CG](#), questions which students can answer by learning the lecture content by heart are quite easy for the LLM. For the [DL4CV2](#) exam, very specific questions about neural network architectures which are covered in our lecture seem to be quite common knowledge in the LLMs, which might be due to those papers being included in the training data. However, for other exams such as [HCI](#), the models lacked specific course context, which was important for answering many theoretical and open-ended questions.

Math-related Ability The LLMs tend to fail on the math-related questions, even the basic ones. For example, they miscount the number of words in a piece of text, or have trouble comparing numbers. For questions that require writing mathematical proof in the [TGI](#) exam, all LLMs except for GPT-4V and Claude failed. For GPT-4V and Claude, they are able to pass the [TGI](#) exam. Their mistakes are more in line with those that students would make. That is, they are often not successful when making actual proof, and the points where the proof

breaks sometimes are the same as the students. Even the better models handle simple geometry questions poorly and/or struggle to follow the instructions of a simple algorithm.

Reasoning Ability The LLMs do not perform well on questions that require deep thinking and reasoning. For questions of the type "is this statement true or false; reason for your solution", the LLMs often said "true" and then just repeated the statement or reasoned for the opposite of their claim. This is a similar behavior often seen in students. Sometimes they make self-contradicting arguments: making a statement and then providing arguments for the other side.

Image Handling GPT-4V, Claude and Llava can handle images. However, only GPT-4V and Claude have reasonable performance. When the question is about drawing on top of the figures, sometimes the LLMs successfully describe in words what needs to be drawn, but occasionally they just hallucinate a non-existing figure file path.

4.3 Automatic Grading

In this section, we evaluate the performance of LLM-as-a-judge approach to automatic grading. We use the expert grades as the gold standard to evaluate automatic graders. We use Pearson correlation on the normalized scores as our metric. Since the LLMs are asked to provide the scores on the same scale as the expert scores, we also provide the Root Mean Squared Error (RMSE) on the originally-scaled scores as a secondary metric. Note that RMSE would correctly put more weight on the questions that have more points, however, it is not as easily interpretable as the Pearson correlation. Therefore, we only report RMSE in Appendix G.2. The main results are discussed as follows.

4.3.1 General

LLMs Perform Well as Graders On the exam level, LLM-as-a-judge performs well for automatic grading. The best Pearson correlation to expert grading on the exam level, at 0.948, is achieved by GPT-4V. The open-source Llama3 achieves 0.883 Pearson correlation to expert grading.

The LLM ranking based on average exam-level grades provided by the GPT-4V grader in comparison to expert grading is shown in Table 5. As can be seen, the ranking is quite identical, except for Mixtral and Qwen's positions being swapped.

		without ref			with ref		
		same question	same exam	diff exam	same question	same exam	diff exam
Text-only questions							
Mixtral	0 shot	0.232			0.311		
	1 shot	0.352	0.377	0.364	0.395	0.333	0.275
	2 shot	0.395	0.299	0.316	0.398	0.271	0.255
Llama3	0 shot	0.452			0.603		
	1 shot	0.573	0.547	0.500	0.672	0.581	0.645
	2 shot	0.598	0.522	0.546	0.644	0.596	0.575
GPT-4V	0 shot	0.607			0.696		
	1 shot	0.605	0.679	0.616	0.653	0.693	0.701
	2 shot	0.672	0.648	0.674	0.717	0.727	0.678
Image-related questions							
GPT-4V	0 shot	0.677			0.539		
	1 shot	0.640	0.661	0.611	0.642	0.539	0.749
	2 shot	0.613	0.632	0.673	0.712	0.465	0.696

Table 4: LLM grading’s Pearson correlation to expert grading on the question level. Note that there are only single scores for zero-shot, since they do not have different shot settings.

Expert grader		GPT-4V grader	
Examinee	Avg. grade (%, sorted)	Examinee	Avg. grade (%, sorted)
Claude	59.4	Claude	57.7
GPT-4V	58.2	GPT-4V	56.2
Mixtral	41.1	Qwen	42.0
Qwen	35.4	Mixtral	38.2
GPT-3.5	32.8	GPT-3.5	38.0
Mistral	25.9	Mistral	24.6
Llava	21.5	Llava	24.2

Table 5: LLM examinees ranking with expert grader and GPT-4V grader.

The high correlations between expert grading and LLM-as-a-judge grading indicate that, although being far from perfect in solving SciEx exams (discussed in Section 4.1.1), the stronger LLMs are quite reliable for grading the exams. This is useful since we would have to rely less on expert grading to evaluate newly developed LLMs’ performance on SciEx. The details of graders’ performance under different settings on the exam level are in Appendix G.1.

On the question level, the performance of LLM-as-a-judge is shown in Table 4. The highest Pearson correlation to expert grading achieved by the LLMs is now around 0.7, which is lower than on the exam level, but still quite high. Surprisingly, the performance of GPT-4V on grading image-related questions is quite comparable to grading text-only questions. This contradicts the finding made by Chen et al. (2024). This could potentially be due

to the small number of image-related questions in SciEx, thus the results might not be generalizable.

Few-shot and References Help The performance of the graders on the question level is shown in Table 4. We observe that adding examples (shots) and adding reference answers in the prompt generally increases the performance of the LLM graders. GPT-4V is the strongest grader, followed by Llama3 and Mixtral. This shows that proprietary LLMs are still stronger as judges, aligning with previous studies (Zheng et al., 2024).

4.3.2 Grader-specific Behaviours

Mixtral Grader Tends to Give Full Points As can be seen from Table 4, Mixtral has the worst performance on grading the exams. We observe that Mixtral tends to give full points to the answers. Without reference and without examples (0-shot), the portion of answers where Mixtral outputs full points is 67.6%, significantly higher than Llama3 and GPT-4V, at 19.1% and 15.1%, respectively. As a result, Mixtral’s precision on giving full points, at 0.181, is much lower than Llama3 and GPT-4V, at 0.380 and 0.527 respectively. As we add more examples and/or add the reference answer to the prompt, the problem is lessened. More details can be found in Appendix G.3.

Mixtral and GPT-4V Copy Grade of Example

For Mixtral and GPT-4V graders, when having one example (shot) from the same question in the prompt without reference, the performance is worse than having the example from the same exam or

from a different exam. We hypothesize that this is due to these graders tend to copy the grades of the examples when having a chunk of duplicated text (i.e., the question description) in the example. This is verified when looking at the statistics: Mixtral and GPT-4V copy the grade of the example 25% of the time, whereas Llama3 does it 13% of the time. As a result, Llama3 can best make use of examples from the same question. The problem is reduced when having more than 1 shot or when the reference answer is included.

4.3.3 Influential Factors

Different Examinees As can be seen in Table 6, GPT-4V grader has better performance than others, but is more inconsistent: it does worse on grading some LLMs, especially Claude. This is potentially due to Claude being a better examinee than GPT-4V itself, as shown in Section 4.1.1. When using the scores from GPT-4V grader to rank the LLMs, we find that, without reference answer, GPT-4V always ranks itself higher than Claude. This emphasizes the importance of reference answers for grading, especially when the grader is weaker than the examinee.

	Graders		
	Mixtral	Llama3	GPT-4V
Claude	0.304	0.460	0.482
GPT-4V	0.353	0.528	0.612
Mixtral	0.251	0.472	0.564
Qwen	0.351	0.556	0.736
GPT-3.5	0.333	0.522	0.697
Mistral	0.291	0.467	0.601
Llava	0.387	0.716	0.812

Table 6: LLM graders performance (i.e., Pearson correlation to expert grading) on different examinees.

Difficulty Levels Looking at Table 7, the weaker graders, i.e., Mixtral and Llama3, perform better on grading easier questions. In contrast, GPT-4V performs better in grading harder questions.

	Graders		
	Mixtral	Llama3	GPT-4V
Easy	0.374	0.602	0.628
Medium	0.293	0.524	0.690
Hard	0.224	0.496	0.732

Table 7: LLM grader’s performance (i.e., Pearson correlation to expert grading) on different difficulty levels.

5 Conclusion

In this paper, we proposed SciEx - a benchmark consisting of scientific university exams, along with expert grading and automatic grading, to evaluate the abilities of LLMs on science topics. SciEx is multilingual, multi-modal, and contains a variety of free-form questions. Our experiments show that SciEx is still quite challenging for current LLMs, where the best LLM can only achieve 59.4% of the exam score on average. Despite that, the LLMs perform well as graders, achieving 0.948 Pearson correlation to the expert grades. This is a promising observation, since we can use strong LLMs for automatic grading of new LLM examinees on SciEx, rather than relying on expert grading. We encourage the research community as well as LLM developers and users to make use of SciEx for evaluating LLMs’ scientific capabilities.

Limitations

There are certain biases that can occur for SciEx. Firstly, the LLMs do not have time pressure. Therefore, they can output longer answers, which helps them get better grades, as there is a higher likelihood that something will be correct. Secondly, the grading process can not be fully anonymized. It is not easy to mix the LLM answers with student answers for the lecturer to grade, since student answers are usually handwritten. Additionally, the LLMs’ answers content itself might also be easily distinguishable from the students’, since the LLMs tend to, e.g., give longer answers or repeat the questions. Therefore, the lecturers know when they are grading an LLM, thus can bias the score they give. Thirdly, the comparison between the LLMs and the students might be unfair, since the students studied the centralized course material specifically for the exams, while this is not the case for the LLMs. Lastly, due to the reliance on expert resources, the size of SciEx is quite small compared to other scientific benchmarks.

Ethics

Our work makes use of student statistics to compare against LLMs’ performance. However, we only use the average of the student grades, without disclosing any individual student’s information. The student answers are never directly used, as we only ask for the average graders from the lecturers.

Automatic grading, regardless of the high correlation to expert grading, can still be imperfect.

We are not suggesting to use LLMs to evaluate students, but to evaluate new models coming out when it is not possible to do human evaluation.

Regarding data consent, we had group meetings and email exchanges to come to an agreement from all lecturers that the data would be made public under the CC BY-NC-SA 4.0 license.

Acknowledgments

This work was supported by the Helmholtz Programme-oriented Funding, with project number 46.24.01, project name AI for Language Technologies. It was also supported by funding from the pilot program Core-Informatics of the Helmholtz Association (HGF).

We thank the lecturers for their contribution during the creation of the dataset: Kunyu Peng, Alexander Jaus, David Schneider, Ruiping Liu, Zdravko Marinov, Yufan Chen, Miklós Borsi, Florian Kalinke, Federico Matteucci, Fabian Richter, Bela Böhnke, Jose Cribeiro-Ramallo, Daniel Ebi, Florian Kalinke, Adrian Feilhauer, Wendy Yi, Laura Merker, Miriam Goetze, Jean-Pierre von der Heydt, Max Göttlicher, Thomas Bläsius, Marcus Wilhelm, Michael Zündorf

References

- Mahyar Abbasian, Iman Azimi, Amir M Rahmani, and Ramesh Jain. 2023. Conversational health agents: A personalized llm-powered agent framework. *arXiv preprint arXiv:2310.02374*.
- Sumithra Bhakthavatsalam, Daniel Khashabi, Tushar Khot, Bhavana Dalvi Mishra, Kyle Richardson, Ashish Sabharwal, Carissa Schoenick, Oyvind Tafjord, and Peter Clark. 2021. Think you have solved direct-answer question answering? try arca, the direct-answer ai2 reasoning challenge. *arXiv preprint arXiv:2102.03315*.
- Dongping Chen, Ruoxi Chen, Shilin Zhang, Yinuo Liu, Yaochen Wang, Huichi Zhou, Qihui Zhang, Pan Zhou, Yao Wan, and Lichao Sun. 2024. Mllm-as-a-judge: Assessing multimodal llm-as-a-judge with vision-language benchmark. *arXiv preprint arXiv:2402.04788*.
- Cheng-Han Chiang and Hung-yi Lee. 2023a. [Can large language models be an alternative to human evaluations?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15607–15631, Toronto, Canada. Association for Computational Linguistics.
- Cheng-Han Chiang and Hung-yi Lee. 2023b. [A closer look into using large language models for automatic evaluation.](#) In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8928–8942, Singapore. Association for Computational Linguistics.
- Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith, and Matt Gardner. 2021. [A dataset of information-seeking questions and answers anchored in research papers.](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4599–4610, Online. Association for Computational Linguistics.
- Aidan Gilson, Conrad Safranek, Thomas Huang, Vimig Socrates, Ling Chi, R Andrew Taylor, and David Chartash. 2022. How does chatgpt perform on the medical licensing exams? the implications of large language models for medical education and knowledge assessment. *MedRxiv*, pages 2022–12.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Tianle Li, Wei-Lin Chiang, and Lisa Dunlap. 2024. [Introducing hard prompts category in chatbot arena.](#) Published on LMSYS.
- Lizi Liao, Grace Hui Yang, and Chirag Shah. 2023. [Proactive conversational agents in the post-chatgpt world.](#) In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23*, page 3452–3455, New York, NY, USA. Association for Computing Machinery.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries.](#) In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. [G-eval: NLG evaluation using gpt-4 with better human alignment.](#) In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *The 36th Conference on Neural Information Processing Systems (NeurIPS)*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation.](#) In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

- Nitarshan Rajkumar, Raymond Li, and Dzmitry Bahdanau. 2022. Evaluating the text-to-sql capabilities of large language models. *arXiv preprint arXiv:2204.00498*.
- Marc Cicero Schubert, Wolfgang Wick, and Varun Venkataramani. 2023. Performance of large language models on a neurology board-style examination. *JAMA network open*, 6(12):e2346721–e2346721.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Johannes Welbl, Nelson F. Liu, and Matt Gardner. 2017. [Crowdsourcing multiple choice science questions](#). In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pages 94–106, Copenhagen, Denmark. Association for Computational Linguistics.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. [Bartscore: Evaluating generated text as text generation](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 27263–27277. Curran Associates, Inc.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Wenzuan Zhang, Mahani Aljunied, Chang Gao, Yew Ken Chia, and Lidong Bing. 2024. M3exam: A multilingual, multimodal, multilevel benchmark for examining large language models. *Advances in Neural Information Processing Systems*, 36.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36.

A Exam Description

The overall description of each exam in SciEx is as follows:

1. Natural Language Processing (NLP): exam contains questions about word and sequence representation, language modeling, and pre-trained models.
2. Advanced Artificial Intelligence (AI2): exam contains questions about natural language processing, signal processing, automatic speech recognition and cognitive robotics.
3. Deep Learning and Neural Networks (DLNN): exam contains questions about neural network fundamentals, in-depth questions about multi-head self attention and calculation questions on backpropagation.
4. Deep Learning for Computer Vision (DL4CV2): exam contains questions about semi-supervised learning, weakly supervised learning, multi-modal text-image models, continual learning, representation learning, interactive segmentation, transfer learning and generative models from recent literature.
5. Human-Computer Interaction (HCI): exam encompasses fundamental HCI subjects like observational studies, human perception and information processing, user studies, and system design and design analysis. It requires students to utilize theoretical knowledge and perform brief analyses based on the given context.
- 6/7. Databases (DBS): 2 exams from 2022 and 2023, containing questions about ER (Entity Relationship) modeling, SQL writing and comprehension, relational algebra, and transaction management.
8. Computer Graphics (CG): exam contains questions about color and perception, raytracing, shading, data structures, transformations, textures, OpenGL, blending, shaders, procedural modeling, and bezier curves.
9. Theoretical Foundations of Computer Science (TGI): exam contains questions about finite automata, regular languages, push down automata, grammars (Chomsky hierarchy), Turing machine, formal languages,

NP-completeness, approximation algorithms, decidability. Most questions require writing mathematical proofs.

10. Algorithms (ALGO): exam contains questions on writing proofs (correctness of algorithms, asymptotic run time analysis), knowing basic algorithms and data structures, designing simple new algorithms, selecting the right data structure or algorithm for a given task at hand.

B Exam Formatting

Originally, exams were in different formats, depending on their creator. We convert the exams into JSON format, with file paths to images if any. An example is shown in Figure 4.

C LLM Answer Generation Prompts

We provide the prompt in the same language as the exam question to the LLMs to generate answers. The English prompt is shown in Figure 5, and the German prompt is shown in Figure 6.

D LLM Grader Prompts

We provide the prompt in the same language as the exam question to the LLMs to perform automatic grading. The English prompt is shown in Figure 7, and the German prompt is shown in Figure 8.

E User Interface for Expert Grading

We instructed the expert grader to use our user interface (UI) for grading. Figure 9 shows the open page of the UI, where the grader can choose their exam and enter their password. Figure 10 shows the page for the grading, where the expert is shown with the question, the LLM answer to the question, and a text box to enter the grade. The expert can choose the examinee to grade from the dropdown on the left-hand side. Figure 11 shows the page to enter additional information about the exam questions, including the maximal achievable score, average student performance, gold answer, and difficulty level.

Once the data is collected, we also ask the experts and have their consent to make the data public.


F Performance on Text-only, English-only Questions

The performance of the LLMs on text-only, English-only questions is shown in Table 8. On this

Question 2 (5 points): Dense Passage Retrieval

As introduced in the lecture, the Dense Passage Retriever (DPR) is a model for open-domain question-answering. The model consists of two encoders, a passage encoder and a question encoder. During training, the passage encoder maps text passages to fixed-size vectors optimized for retrieval. At inference time, the question is first encoded to a fixed-size vector by the question encoder. Then the nearest encoded passages are retrieved as potential answers.

- a. (3 points) When creating fixed-size vectors for passages and questions, DPR uses sentence representations from the BERT CLS token. What are the differences between this approach and pooling (e.g. meanpool or maxpool) the word vectors in the passages or questions as the sentence representation? What are the advantages? As supplementary information, a figure about the BERT CLS token is shown below.



- b. (2 points) In DPR, the training objective pushes relevant question-passage pairs closer to each other, while irrelevant pairs away from each other. Why do we need to include the irrelevant/negative pairs? What would happen if we leave out the irrelevant/negative pairs in the training objective?

(a) Exam in original PDF format.

```
{
  "Index": "4-2",
  "Description": "Dense Passage Retrieval. As introduced in the lecture, the Dense Passage Retriever (DPR) is a model for open-domain question-answering. The model consists of two encoders, a passage encoder and a question encoder. During training, the passage encoder maps text passages to fixed-size vectors optimized for retrieval. At inference time, the question is first encoded to a fixed-size vector by the question encoder. Then the nearest encoded passages are retrieved as potential answers.",
  "Subquestions": [
    {
      "Index": "a",
      "Content": "When creating fixed-size vectors for passages and questions, DPR uses sentence representations from the BERT CLS token. What are the differences between this approach and pooling (e.g. meanpool or maxpool) the word vectors in the passages or questions as the sentence representation? What are the advantages? As supplementary information, a figure about the BERT CLS token is shown below.",
      "Figures": ["figures/bert_class_bw.png"]
    },
    {
      "Index": "b",
      "Content": "In DPR, the training objective pushes relevant question-passage pairs closer to each other, while irrelevant pairs away from each other. Why do we need to include the irrelevant/negative pairs? What would happen if we leave out the irrelevant/negative pairs in the training objective?"
    }
  ]
}
```

(b) Exam converted to JSON format.

Figure 4: Exam question before and after being converted to JSON format.

subset of questions, besides GPT-4V and Claude, we can see that Mixtral and Qwen also have better performance than the student average.

	Grade (%)	German Scale
Proprietary		
GPT-4V	70.8	1.4
Claude	69.2	1.6
GPT-3.5	47.8	2.9
Open source		
Mixtral	61.2	2.0
Qwen	56.8	2.4
Mistral	48.0	3.2
Llava	42.4	3.5
Student avg.	56.5	2.4

Table 8: Average performance of the LLMs on the exam level, provided by expert grading, text-only and English-only questions.

G Grader Performance

G.1 Pearson Correlation on Exam Level

The performance of LLM-as-a-judge for automatic grading on the exam level is shown in Table 9. Note that Mixtral and Llava3 graders have disadvantage since they cannot take image input for image-related questions.

G.2 RMSE on Question-level

Since the LLM graders are asked to output the scores in the original scale, RMSE would be the most informative metric, since it also reflects the

importance of the questions that have higher maximum scores. The LLM graders' performance in RMSE is shown in Table 10.

G.3 Performance on Giving Full Points

The performance of the LLM graders on assigning full points to the answers is shown in Table 11.

You are a university student. Please answer the following JSON-formatted exam question. The subquestions (if any) are indexed. The provided figures (if any) each contains its path at the bottom, which matches the path provided in the JSON.

Please give the answers to the question and subquestions that were asked, and index them accordingly in your output.

You do not have to provide your output in the JSON format. If you are asked to draw on the figure, then describe with words how you would draw it.

Please provide all answers in English.

Here is the question:

<input text>

Figure 5: Answer generation prompt in English.

Sie sind Student. Bitte beantworten Sie die folgende JSON-formatierte Prüfungsfrage. Die Unterfragen (falls vorhanden) sind indiziert. Die bereitgestellten Abbildungen (falls vorhanden) enthalten jeweils unten ihren Pfad, der mit dem im JSON bereitgestellten Pfad übereinstimmt.

Bitte geben Sie die Antworten auf die gestellten Fragen und Unterfragen an und indizieren Sie diese in Ihrer Ausgabe entsprechend.

Sie müssen Ihre Ausgabe nicht im JSON-Format bereitstellen. Wenn Sie aufgefordert werden, auf der Figur zu zeichnen, beschreiben Sie mit Worten, wie Sie sie zeichnen würden.

Bitte geben Sie alle Antworten auf Deutsch an.

Hier ist die Frage:

<input text>

Figure 6: Answer generation prompt in German.

You are a university professor. Please grade the following exam question.

The exam question, examinee's answer, correct answer, and the maximum possible score are provided in the format:

- [question] <exam_question> [/question]
- [answer] <answer> [/answer]
- [correct_answer] <correct_answer> [/correct_answer]
- [max_score] <max_score> [/max_score]

The question is provided in JSON format, but the answer can be freeform text. The provided figures in the question (if any) each contain its path at the bottom, which matches the path provided in the JSON. The answer is text-only. If the question asks to draw on the figure, then the answer should contain a text description of how the drawing should be.

Please provide the grade between [0, <max_score>]. Please provide the reasoning for your grade. Please provide your output in the format:

- [reason] <reasoning> [/reason]
- [grade] <grade> [/grade]

Below you are provided with examples on how to perform the grading:

<example text>

Here is your input:

<input text>

Figure 7: Grading prompt in English.

Sie sind Universitätsprofessor. Bitte bewerten Sie die folgende Prüfungsfrage.
Die Prüfungsfrage, die Antwort des Prüflings, die richtige Antwort und die maximal mögliche Punktzahl werden im Format bereitgestellt::

- [question] <Prüfungsfrage> [/question]
- [answer] <Antwort> [/answer]
- [correct_answer] <korrekteAntwort> [/correct_answer]
- [max_score] <maxPunkt> [/max_score]

Die Frage wird im JSON-Format bereitgestellt, die Antwort kann jedoch Freiformtext sein. Die bereitgestellten Abbildungen in der Frage (falls vorhanden) enthalten jeweils unten ihren Pfad, der mit dem im JSON bereitgestellten Pfad übereinstimmt. Die Antwort ist nur Text. Wenn es sich bei der Frage darum handelt, auf der Abbildung zu zeichnen, sollte die Antwort eine Textbeschreibung darüber enthalten, wie die Zeichnung aussehen soll.

Bitte geben Sie die Note zwischen [0, <maxPunkt>] an. Bitte begründen Sie Ihre Note. Bitte geben Sie Ihre Ausgabe im Format an:

- [reason] <Grundsatz> [/reason]
- [grade] <Note> [/grade]

Nachfolgend finden Sie ein Beispiel für die Durchführung der Benotung:

<example text>

Hier ist Ihre Eingabe:

<input text>

Figure 8: Grading prompt in German.

LLM exam grading

You can fill in the additional information by choosing the `additional_information` option on the left side bar. You can grade each LLM outputs by choosing `llm_<index>` on the left side bar. Your input is automatically saved after you press `Enter` on each text box.

Enter your exam name

Enter your password

Figure 9: Open page of the grading UI.

		without ref			with ref		
		same question	same exam	diff exam	same question	same exam	diff exam
Mixtral	0 shot	0.404			0.445		
	1 shot	0.542	0.549	0.619	0.565	0.390	0.344
	2 shot	0.564	0.505	0.620	0.500	0.466	0.463
Llama3	0 shot	0.649			0.677		
	1 shot	0.812	0.731	0.706	0.883	0.770	0.772
	2 shot	0.771	0.729	0.768	0.788	0.738	0.785
GPT-4V	0 shot	0.911			0.938		
	1 shot	0.886	0.906	0.893	0.902	0.904	0.948
	2 shot	0.921	0.910	0.896	0.917	0.888	0.934

Table 9: LLM graders’ Pearson correlation to expert graders on exam level, scores normalized. Note that there are only a single scores for zero-shot, since they do not have different shot settings.

		without ref			with ref		
		same question	same exam	diff exam	same question	same exam	diff exam
Text-only questions							
Mixtral	0 shot	3.25			2.96		
	1 shot	2.68	2.90	2.83	2.54	2.65	2.79
	2 shot	2.69	2.83	2.86	2.51	2.46	2.70
Llama3	0 shot	2.66			2.09		
	1 shot	1.89	2.30	2.50	1.45	1.92	1.85
	2 shot	1.88	2.36	2.31	1.92	1.77	1.92
GPT-4V	0 shot	1.56			1.20		
	1 shot	1.34	1.32	1.53	1.49	1.34	1.28
	2 shot	1.25	1.29	1.54	1.36	1.31	1.31
Image-related questions							
GPT-4V	0 shot	0.30			0.28		
	1 shot	0.22	0.24	0.29	0.19	0.26	0.17
	2 shot	0.23	0.28	0.30	0.18	0.24	0.19

Table 10: LLM grading’s RMSE compared to expert grading on the question level. Note that there are only single scores for zero-shot, since they do not have different shot settings.

This screenshot shows a grading interface for a question. On the left, a sidebar contains a dropdown menu with the text "Select to fill in additional information, or select an LLM to grade:" and the selected option "llm0". The main content area displays the question "Question Q1a: Warm Up. What is the definition of human-computer interaction according to SIG CHI?". Below the question is the answer: "The Association for Computing Machinery's Special Interest Group on Computer-Human Interaction (SIG CHI) defines human-computer interaction as the study of the design, evaluation, and implementation of interactive computing systems for human use and the study of the impact of these systems on human behavior." Underneath the answer, it says "Points out of 1" and a grade input field containing the number "1". Below this, the next question "Question Q1b: Warm Up. What is the difference between 'usability' and 'user experience'?" is partially visible.

Figure 10: Grading UI showing the question, the answer and textbox to input the grade.

This screenshot shows a grading interface for the same question as Figure 10. The sidebar dropdown menu now shows "additional_information". The main content area displays the question "Question Q1a: Warm Up. What is the definition of human-computer interaction according to SIG CHI?". Below the question, several metadata fields are shown in light blue boxes: "Maximum Points" with the value "1", "Average Student Points" with the value "0.0", "Gold Answer English" with the text "Human-computer interaction is a discipline concerned with the design, evaluation and implementation", "Gold Answer German" with the text "Correct German answer", and "Difficulty Label" with the value "medium".

Figure 11: Grading UI where grader can fill in additional information about the question.

		Max score precision		Max score predicted (%)	
		No ref	Ref	No ref	Ref
Mixtral	0	0.181	0.196	67.0	37.2
	1	0.251	0.325	46.3	26.6
	2	0.258	0.256	45.3	37.1
Llama3	0	0.380	0.636	19.1	5.9
	1	0.438	0.579	19.9	8.5
	2	0.447	0.483	17.8	10.5
GPT-4V	0	0.527	0.405	15.0	11.3
	1	0.422	0.526	16.8	10.2
	2	0.458	0.560	15.9	9.8

Table 11: Performance on giving full points