

# SciEval: A Multi-Level Large Language Model Evaluation Benchmark for Scientific Research

Liangtai Sun, Yang Han, Zihan Zhao, Da Ma, Zhennan Shen, Baocai Chen, Lu Chen\*, Kai Yu\*

X-LANCE Lab, Department of Computer Science and Engineering  
MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University  
Shanghai Jiao Tong University, Shanghai, China  
{slt19990817, csyanghan, zhao\_mengxin, mada123}@sjtu.edu.cn  
{ieec-szn, 15368493547, chenlusz, kai.yu}@sjtu.edu.cn

## Abstract

Recently, there has been growing interest in using Large Language Models (LLMs) for scientific research. Numerous benchmarks have been proposed to evaluate the ability of LLMs for scientific research. However, current benchmarks are mostly based on pre-collected objective questions. This design suffers from data leakage problem and lacks the evaluation of subjective Q/A ability. In this paper, we propose *SciEval*, a comprehensive and multi-disciplinary evaluation benchmark to address these issues. Based on Bloom’s taxonomy, *SciEval* covers four dimensions to systematically evaluate scientific research ability. In particular, we design a “dynamic” subset based on scientific principles to prevent evaluation from potential data leakage. Both objective and subjective questions are included in *SciEval*. These characteristics make *SciEval* a more effective benchmark for scientific research ability evaluation of LLMs. Comprehensive experiments on most advanced LLMs show that, although GPT-4 achieves SOTA performance compared to other LLMs, there is still substantial room for improvement, especially for dynamic questions. The codes and data are publicly available on <https://github.com/OpenDFM/SciEval>.

## 1 Introduction

Large Language Models (LLMs), such as ChatGPT (Schulman et al. 2022), have attracted widespread attention in general scenarios, including information search, code generation, and more. In the field of science, LLMs have also shown preliminary potential in improving scientific research efficiency and transforming scientific research paradigms (Blanco-Gonzalez et al. 2023; WANG and MIAO 2023). In the meanwhile, several scientific LLMs have been proposed by researchers (Taylor et al. 2022; Luo et al. 2022; Frey et al. 2022). In the general field, there are already numerous evaluation benchmarks to evaluate the language understanding, language generation and reasoning capabilities of LLMs, such as MMLU (Hendrycks et al. 2020), AGIEval (Zhong et al. 2023), and C-EVAL (Huang et al. 2023), shown in Table 1. Although these benchmarks cover data of science domain, the data sources are usually confined to educational materials, which can not adequately as-

sess the research ability of LLMs and not align with real-life scientific research scenarios. In addition, some benchmarks have been proposed to evaluate the scientific capability of LLMs, such as MultiMedQA (Singhal et al. 2023), Chem-LLMBench (Guo et al. 2023), and MATH (Hendrycks et al. 2021), while these benchmarks are restricted to a specific scientific discipline, leaving a lack of a more general scientific evaluation benchmark.<sup>1</sup> In addition, these benchmarks (1) lack evaluation systems for scientific capabilities, (2) are all based on objective questions, which are insufficient to assess scientific abilities, and (3) face the risk of data leakage.

In response to this gap, we present SciEval, an English benchmark designed to evaluate advanced abilities of LLMs in the scientific domain. SciEval consists of a total of about 18000 challenging scientific questions, spanning three important basic science fields: chemistry, physics and biology, each of which is further divided into multiple sub-topics. SciEval mainly has the following three characteristics:

- **Multi-level and comprehensive evaluation of the ability of LLMs in the scientific field.** Scientific ability of LLMs needs to be evaluated from multiple aspects. Leveraging cognitive domains of Bloom’s taxonomy (Krathwohl 2002; Forehand 2010), which covers six levels, SciEval evaluates the scientific capabilities of large language models across four dimensions: basic knowledge, knowledge application, scientific calculation, and research ability, where each capability aligns with one or more cognitive levels.
- **Combination of objective and subjective questions.** SciEval is mainly based on objective questions, which allow for quick and standard model evaluations, involving multiple-choice, fill-in-the-blank, and judgment questions. These questions can help us understand whether the model can correctly understand and memorize scientific knowledge. However, objective questions are insufficient to assess scientific capability holistically. To better assess scientific reasoning and application ability, SciEval introduces a small number of subjective questions, involving a total of twelve basic science experiments, which is named Experimental Data.

\*The corresponding authors are Lu Chen and Kai Yu.  
Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup>Due to the page limitation, we only compare some widely used benchmarks. For more information, we refer to (Chang et al. 2023).

Name	Category	Ability	Source	Data Type	Dynamic	#Data
MMLU	humanities, social science, STEM, other	BK, KA, SC	exam, book, course	objective	✗	14079
AGIEval	social science, STEM	BK, KA, SC	exam	objective	✗	8062
C-EVAL	humanities, social science, STEM, other	BK, KA, SC	exam	objective	✗	12342
MultiMedQA	medical	BK, KA, RA	exam, research	objective	✗	13115
ChemLLMBench	chemistry	BK,KA	knowledge base	objective	✗	800
MATH	mathematics	SC	exam	objective	✗	5000
SciEval	science	BK, KA,SC, RA	community QA, knowledge base	objective + subjective	✓	15901

Table 1: Dataset comparison of SciEval and some other datasets covering science domain. “BK” stands for Basic Knowledge, “KA” stands for Knowledge Application, “SC” stands for Scientific Calculation, and “RA” stands for Research Ability.

- **Dynamic data generation based on basic scientific principles.** The huge amount of training data used for pre-training LLMs may cause the risk of data leakage for evaluation. In order to solve this problem, one of the main features of SciEval is the use of Dynamic Data, which can prevent potential data leakage and ensure the fairness and credibility of the evaluation results. The Dynamic Data will be updated regularly, and we will maintain a stable version to make a fair comparison of model performance. And the objective questions other than Dynamic Data are referred to as Static Data.

We conduct experiments to evaluate LLMs on SciEval in answer-only, chain-of-thought and few-shot settings. Results indicate that GPT-4 is the strongest model, with only GPT-4, GPT-3.5-turbo and Claude-v1.3 surpassing 60% average accuracy on Static Data, signifying considerable opportunities for improvement. With the results of Dynamic Data, we find that these LLMs have little knowledge about molecules, and most models could only retain near-random accuracy in the physics subset. As for Experimental Data, some top-tier models could perform satisfactorily in experimental principle and design, while almost all models struggle to analyze the experimental results. With the analysis of experiment results, we claim that training on large-scale scientific corpus is helpful for the scientific ability of LLMs, and most LLMs perform bad on calculation problems, especially in physics domain. We hope SciEval can provide an excellent benchmark for the assessment of scientific capability of LLMs, and promote wide application in science.

## 2 Related Work

### 2.1 General Benchmarks for LLMs

To evaluate the performance of LLMs across different tasks, several benchmarks have been proposed. MMLU (Hendrycks et al. 2020) aims to develop a comprehensive test for evaluating text models in multi-task contexts. HELM (Liang et al. 2022) offers a comprehensive assessment, evaluating LLMs across various aspects, such

as language understanding and common-sense reasoning. Big-Bench (Srivastava et al. 2022) introduces 204 challenging tasks covering various domains, aiming to evaluate tasks beyond the capabilities of existing language models. AGIEval (Zhong et al. 2023) serves as an evaluation framework for assessing the performance of foundation models in human-centric standardized exams. C-Eval (Huang et al. 2023) assesses the advanced knowledge and reasoning capabilities of foundation models in Chinese.

### 2.2 Specific Benchmarks for LLMs

Apart from general tasks, specific benchmarks are designed for certain downstream tasks. MultiMedQA (Singhal et al. 2023) focuses on medical question-answering, evaluating LLMs in terms of clinical knowledge and QA abilities. MATH (Hendrycks et al. 2021) assesses reasoning and problem-solving proficiencies of LLMs in mathematics. ScienceQA (Lu et al. 2022) proposes a multi-modal benchmark with a diverse set of science topics and annotations of their answers with corresponding lectures and explanations, collected from elementary and high school science curricula. SCIBENCH (Wang et al. 2023) examines the reasoning capabilities required for complex scientific problem-solving and proposes two datasets of college-level scientific problems. Compared to these benchmarks, SciEval (1) evaluates scientific capabilities from multiple aspects, having a broader coverage, (2) uses data of community Q&A, which is more flexible and diverse, (3) designs a subset of dynamic data, making an effort to mitigate data leakage.

## 3 The SciEval Dataset

### 3.1 Scientific Research Evaluation System

Scientific research requires different dimensions of knowledge, such as understanding and calculation, thence evaluation of scientific ability should be conducted at multiple levels. Bloom’s taxonomy is a set of three hierarchical methods used for classification of educational learning objectives covering cognitive, affective and psychomotor domains. The

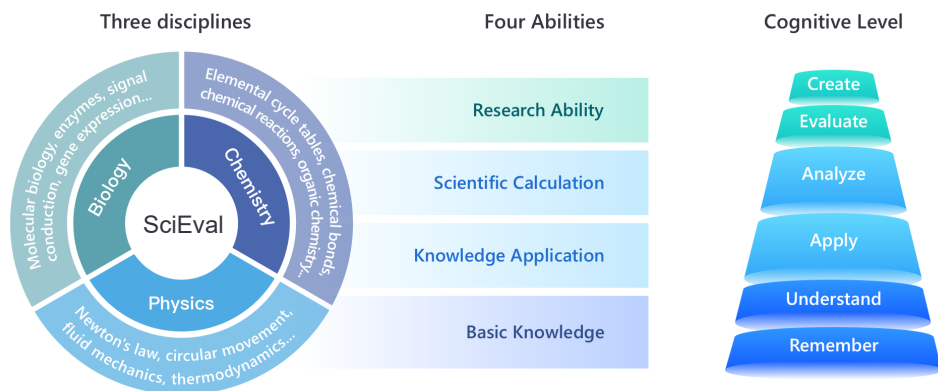


Figure 1: The illustration of the evaluation system. SciEval covers three disciplines with amounts of sub-topics, and investigates four abilities, corresponding to six cognitive levels.

cognitive domain is frequently used to structure curriculum learning objectives, assessments and activities, and is broken into six levels: Remember, Understand, Apply, Analyze, Evaluate and Create, as is shown in Figure 1, which are suitable for the evaluation of scientific capability.

Based on the cognitive domain of Bloom’s taxonomy, the evaluation system of SciEval consists of four knowledge dimensions: *Basic Knowledge (BK)*, *Knowledge Application (KA)*, *Scientific Calculation (SC)*, and *Research Ability (RA)*. As is shown in Figure 1, BK primarily assesses the fundamental scientific knowledge of LLMs. KA focuses on how to apply basic knowledge to solve scientific problems, requiring models to have comprehension, application, and analysis abilities. SC is a specialized application of knowledge that further examines complex reasoning capabilities of LLMs based on their general knowledge application abilities. RA assesses evaluation capabilities at a higher cognitive level, requiring models to participate in various aspects of scientific research, including problem formulation, experimental design, data analysis, and summarization.

Based on the evaluation system, we design three different types of data: *Static Data*, *Dynamic Data*, and *Experimental Data*. The Static Data covers all these four knowledge dimensions and will remain constant throughout, while the Dynamic Data examines from the aspects of Knowledge Application and Scientific Calculation and will be regularly updated to prevent any data leakage. The Experimental Data comprises a set of questions for twelve scientific experiments and can be used to evaluate the Research Ability.

### 3.2 Data Collection

**Static Data** The collection steps of Static Data are shown in Figure 2. The primary source of Static Data is Socratic Q&A<sup>2</sup>, a community-driven website that covers a wide range of subjects such as science and literature. Specifically, we collect data from the fields of biology, chemistry, and physics. To ensure quality, we employ rule-based methods

<sup>2</sup><https://socratic.org>

to preprocess the crawled data. While gathering the questions, we found that not all of them are suitable as titles. To address this, we utilize GPT-4 with the “Task 1” prompt, as depicted in Figure 2, to process these questions. Since most of the collected questions are open-ended and challenging to evaluate, we employ GPT-4 to simplify ground-truth answers and generate three wrong answers to formulate them as multiple-choice questions. Additionally, we classify the questions into their respective knowledge domains. And during this process, we manually check the generated content of GPT-4 to ensure data quality.

To make the dataset more diverse and comprehensive, we further integrate data from some publicly available datasets:

- **MedQA** (Jin et al. 2021) is a free-form multiple-choice OpenQA dataset for solving medical problems, collected from professional medical board exams. We use the test set of USMLE, which is the English subset of MedQA.
- **PubMedQA** (Jin et al. 2019) is a biomedical question-answering dataset collected from PubMed abstracts. The task of PubMedQA is to answer research questions with yes/no/maybe using the corresponding abstracts, which is fit for evaluating the literature comprehension ability. We incorporate 1000 expert-annotated data from it and frame them as judgment questions.
- **Reagent Selection** (Guo et al. 2023) involves the identification and proposal of the most fitting reagents for a specific chemical reaction or process, which is a subset of ChemLLMBench. We randomly select 40% data and formulate them as multiple-choice questions.

**Dynamic Data** The current training of LLMs often uses a large amount of data, resulting in a risk of data leakage for evaluation. In order to solve this problem, we design a “dynamic” subset, which can generate data dynamically according to scientific principles. The dynamic subset covers two disciplines, chemistry and physics. For chemistry data, we use the basic information and properties of molecules

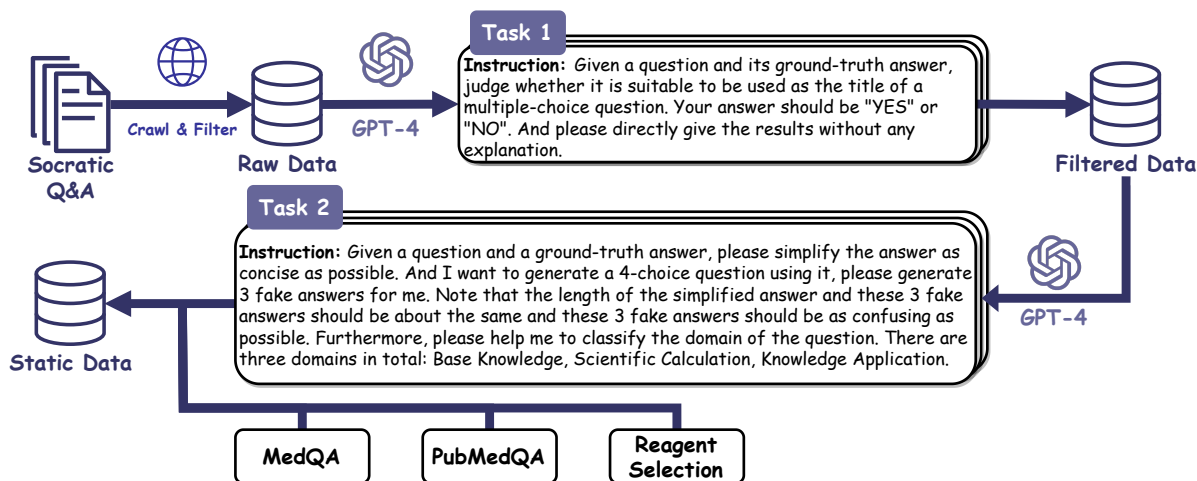


Figure 2: Data Collection steps of Static Data

crawled from PubChem<sup>3</sup> to create data. For physics data, we manually write some Python scripts according to the physics formulas. When obtaining the evaluation dataset, we will provide a regenerated version to users and we will update it regularly, while at the same time, we will maintain a stable version of the dynamic data to make a fair comparison.

**Experimental Data** To better evaluate the scientific thoughts and abilities of LLMs, SciEval introduces a subset of experimental data, involving 12 different basic scientific experiments. These experiments are collected from basic science experiment courses at university, and each experiment conducts a comprehensive investigation of the ability of LLMs in scientific research and experimentation from the perspectives of experimental principle, process, and analysis and summarization of experimental results.

### 3.3 Data Statistics

Summarized statistics are shown in Table 2, where we only count Static Data. For Dynamic Data, the chemistry part examines the KA ability and contains 2000 data, while the physics part evaluates the SC ability and involves 890 data. All these questions are in English and we show some data examples in Appendix D.

Ability	Bio	Chem	Phy
Basic Knowledge	2147	2914	456
Knowledge Application	1379	3720	36
Scientific Calculation	301	3401	1165
Research Ability	1000	0	0
Total	4830	10035	1657

Table 2: Statistics of Static Data

For Static Data, we further split the data into dev, valid, and test set. For each data source, each knowledge domain, and each discipline, we randomly select 5 data to form the

<sup>3</sup><https://pubchem.ncbi.nlm.nih.gov/>

dev set, which can be used for few-shot learning, and we split the remaining data with a ratio of 1:9 to construct the valid set and test set respectively.

## 4 Experiment

Given a question and four options, please select the right answer. Your answer should be "A", "B", "C" or "D". Please directly give the answer without any explanation.

How many atoms are in 3.5 moles of arsenic atoms?

- A.  $1.5 \times 10^{24}$  atoms
- B.  $3.0 \times 10^{24}$  atoms
- C.  $2.7 \times 10^{24}$  atoms
- D.  $2.1 \times 10^{24}$  atoms

Answer: D

Figure 3: An example of the prompt we used for AO setting. The red text is the response from the model, while the black text is the inputted prompt.

Given a question and four options, please select the right answer. Your answer should be "A", "B", "C" or "D".

How many atoms are in 3.5 moles of arsenic atoms?

- A.  $1.5 \times 10^{24}$  atoms
- B.  $3.0 \times 10^{24}$  atoms
- C.  $2.7 \times 10^{24}$  atoms
- D.  $2.1 \times 10^{24}$  atoms

Answer: Let's think step by step: To find the number of atoms ..... Therefore, the answer is D

Figure 4: An example of the prompt we used for CoT setting. The red text is the response from the model, while the blue text and black text are the inputted prompt.

### 4.1 Experiment Setup

**Prompts** We evaluate LLMs in both Answer-Only (AO) and Chain-Of-Thought (CoT) (Kojima et al. 2022) settings. The prompts we used are shown in Figures 3 and 4 respectively. Furthermore, we also evaluate using 3-shot setting, where the three exemplars are selected from the dev set.

Model	Creator	#Parameters	Access	SD	DD	ED
GPT-4	OpenAI	<i>undisclosed</i>	API	✓	✓	✓
GPT-3.5-turbo	OpenAI	<i>undisclosed</i>	API	✓	✓	✓
Claude-v1.3	Anthropic	<i>undisclosed</i>	API	✓	✓	✓
Claude-instant-v1.1	Anthropic	<i>undisclosed</i>	API	✓	✓	✓
ERNIE Bot	Baidu	<i>undisclosed</i>	Web			✓
SparkDesk	iFLYTEK	<i>undisclosed</i>	Web			✓
Vicuna	LMSYS	13B	Weights	✓	✓	
Galactica	Meta	30B, 6.7B	Weights	✓	✓	
ChatGLM2	Tsinghua	6B	Weights	✓	✓	
ChatGLM	Tsinghua	6B	Weights	✓	✓	
Alpaca	Stanford	7B	Weights	✓	✓	
MOSS	Fudan	16B	Weights	✓	✓	
LLaMa	Meta	7B, 13B	Weights	✓	✓	

Table 3: Models evaluated in this paper. The “access” columns show whether we have full access to the model weights or we can only access through API or web. SD stands for Static Data, DD stands for Dynamic Data, and ED stands for Experimental Data. Marking “✓” means we evaluate the corresponding model on this subset.

Model	Static Data				Chemistry(DD)			Physics(DD)	Exp Score
	Biology	Chemistry	Physics	Avg.	Acc.	BLEU	MSE	Acc.	
GPT-4	<b>84.49</b>	<b>69.38</b>	<b>65.22</b>	<b>73.93</b>	<b>11.05</b>	<b>23.78</b>	891.09	25.84	<b>93.31</b>
GPT-3.5-turbo	76.42	64.30	52.30	66.97	7.65	18.86	2008.72	21.80	88.27
Claude-v1.3	72.58	59.72	54.94	63.45	5.75	21.98	1489.87	26.14	85.73
Claude-instant-v1.1	70.43	53.36	52.30	58.92	0.45	16.07	8258.46	21.46	87.50
Galactica-30B	66.48	50.16	44.65	54.96	0.9	4.14	<b>485.99</b>	22.47	-
Vicuna-13B	58.39	53.06	45.13	53.93	0.95	6.50	766.64	21.24	-
Galactica-6.7B	57.84	50.77	30.99	50.87	1.55	6.47	5519.82	20.79	-
ChatGLM2-6B	58.62	44.00	40.26	48.44	0.2	1.86	3449.44	24.83	-
ChatGLM-6B	52.54	45.36	40.80	47.23	0.75	2.44	10303.90	21.01	-
Alpaca-7B	56.66	42.43	37.01	46.54	0.2	2.92	428419.27	26.74	-
MOSS-16B	47.71	33.87	31.73	38.23	0.1	7.37	30505.17	24.27	-
LLaMa-13B	48.59	33.56	19.48	36.96	0.3	5.21	3707.01	7.08	-
LLaMa-7B	36.24	26.38	15.02	28.37	0.5	1.26	11305.65	14.38	-
ERNIE Bot	-	-	-	-	-	-	-	-	61.12
SparkDesk	-	-	-	-	-	-	-	-	33.69

Table 4: Model performances of Answer-Only setting. The leaderboard is sorted by the average accuracy of Static Data.

**Models** In order to comprehensively assess the scientific capabilities of Large Language Models (LLMs), we evaluate 15 high-performing LLMs that are widely accessible. These models are selected to represent a diverse range of organizations and vary in size. The details of these models are summarized in Table 3.

- GPT-3.5-turbo and GPT-4 (Schulman et al. 2022; OpenAI 2023) are the strongest GPT model variants from OpenAI that have undergone pretraining, instruction tuning, and reinforcement learning from human feedback (RLHF, (Ouyang et al. 2022)).
- Claude<sup>4</sup>, developed by Anthropic, is often considered

<sup>4</sup><https://www.anthropic.com/index/introducing-claude>.

comparable to GPT-3.5-turbo. We evaluate both the Claude-v1.3 and Claude-instant-v1.1, a lighter version.

- ERNIE Bot<sup>5</sup> is developed by Baidu, possessing deep semantic understanding and generation capabilities across modalities and languages. SparkDesk<sup>6</sup> is proposed by iFLYTEK. It has cross-domain knowledge and language understanding capabilities and can understand and execute tasks based on natural dialogue.
- LLaMa (Touvron et al. 2023), developed by Meta, is probably the best open-weight foundation model so far. Vicuna (Zheng et al. 2023) and Alpaca (Taori et al. 2023)

<sup>5</sup><https://yiyao.baidu.com/>

<sup>6</sup><https://xinghuo.xfyun.cn/>

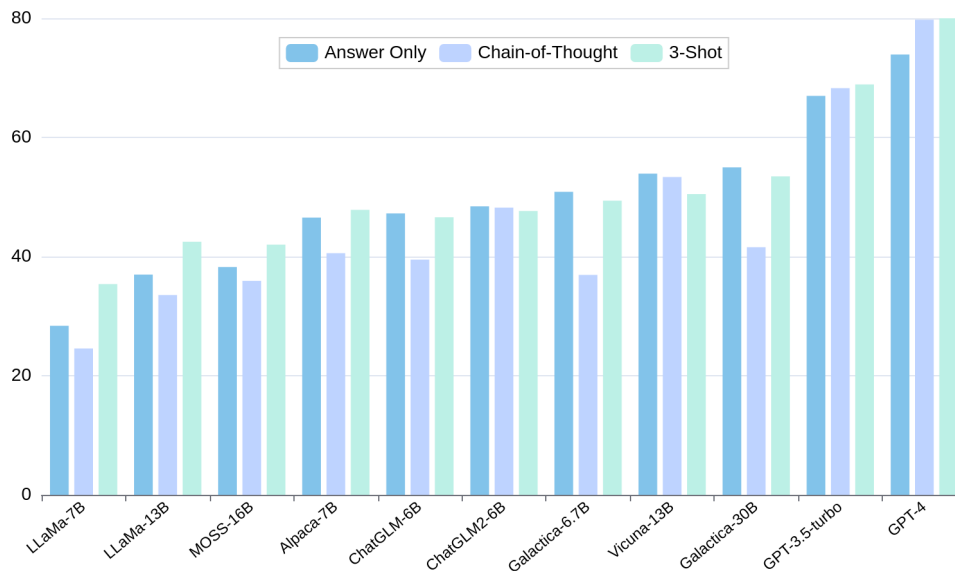


Figure 5: Accuracy on Answer Only, Chain-of-Thought and 3-Shot settings of each LLMs for Static Data.

are both fine-tuned from LLaMa with supervised instruction fine-tuning. It is believed that the performance of Vicuna is better than that of Alpaca.

- Galactica (Taylor et al. 2022) is also developed by Meta, which is trained on a large-scale scientific corpus. It is developed to study the use of language models for the automatic organization of science and can perform numerous scientific tasks, such as citation prediction, scientific QA, and molecular property prediction.
- ChatGLM and ChatGLM2, created by Tsinghua University, are based on GLM architecture (Du et al. 2022), and further adapted on conversational data. MOSS (Sun et al. 2023), developed by Fudan University, is the first publicly available Chinese LLM, and it follows a training procedure similar to ChatGPT.

We evaluate GPT-3.5-turbo, GPT4 and Claude on all three subsets, including Static Data, Dynamic Data, and Experimental Data. Since we can only assess ERNIE Bot and SparkDesk through web interface, we evaluate these two models only on the Experimental Data. And for the rest LLMs with billions or tens of billions of parameters, since the length of the Experimental Data exceeds the length limit of these models<sup>7</sup>, we evaluate them on Static Data and Dynamic Data, as is shown in Table 3.

**Evaluation Metrics** In the case of Static Data, all questions are objective, making accuracy the appropriate evaluation metric. For Dynamic Data, the physics questions are presented as multiple-choice questions, which can also be evaluated using accuracy. Conversely, the chemistry questions involve complex components, such as “What is the

<sup>7</sup>The maximum context length of ChatGLM2 is extended to 32k, while it has limited ability to understand long texts.

molecular weight of A?” and “What is the SMILES expression of B?”. Hence, for questions with numerical answers, we employ MSE<sup>8</sup> as the evaluation metric, while for questions with string answers, we utilize the BELU score (Papineni et al. 2002). Additionally, we also calculate the extract match scores. As for Experimental Data, each experiment consists of multiple open-ended questions. As a result, we assess the model-generated responses manually.

## 4.2 Experiment Results

**Answer-Only Setting** Answer-only results of all the models on the test set are shown in Table 4 and detailed results of Static Data across different knowledge domains are provided in Appendix B. Analyzing the results of Static Data, GPT-4 demonstrates significantly superior performance compared to other models. And only GPT-4, GPT-3.5-turbo, and Claude-v1.3 achieve an average accuracy exceeding 60%, which highlights the challenge posed by SciEval.

For the results of Dynamic Data, GPT-4 performs the best in terms of average accuracy and BLEU score. However, for counting and calculation questions, Galactica-30B yields the best results, indicating its strong aptitude in the field of science. Conversely, models with billions or tens of billions of parameters perform poorly on the chemistry subset, suggesting their limited knowledge about molecules. Regarding the performance of models on the physics subset, since all questions are 4-choices questions, the accuracy should be at least 25%. However, none of these models achieve satisfactory results in this subset.

As for Experimental Data, GPT-series models and

<sup>8</sup>If the predictions do not contain any number, we will regard the MSE as  $1 \times 10^{10}$

Model	Chemistry			Physics		
	AO	CoT	3-Shot	AO	CoT	3-Shot
GPT-4	<b>11.05</b>	<b>11.65</b> ↑	<b>12.42</b> ↑	25.84	17.98 ↓	<b>51.01</b> ↑
GPT-3.5-turbo	7.65	10.20 ↑	8.85 ↑	21.80	<b>47.19</b> ↑	25.39 ~
Galactica-6.7B	1.55	1.75 ↑	3.05 ↑	20.79	23.37 ~	21.12 ~
Vicuna-13B	0.95	1.95 ↑	1.80 ↑	21.24	18.65 ~	23.37 ~
Galactica-30B	0.90	2.60 ↑	3.30 ↑	22.47	14.72 ↓	22.58 ~
ChatGLM-6B	0.75	0.80 ↑	1.15 ↑	21.01	25.39 ~	23.37 ~
LLaMa-7B	0.50	0.10 ↓	1.55 ↑	18.65	9.66 ↓	27.53 ↑
LLaMa-13B	0.30	0.25 ~	2.11 ↑	7.08	5.84 ~	22.70 ↑
ChatGLM2-6B	0.20	2.65 ↑	1.60 ↑	24.83	25.39 ~	26.74 ~
Alpaca-7B	0.20	0.65 ↑	2.10 ↑	<b>26.71</b>	28.43 ~	25.62 ~
MOSS-16B	0.10	0.85 ↑	0.65 ↑	24.27	25.06 ~	26.40 ~

Table 5: Results on Answer-Only, Chain-of-Thought and 3-Shot settings of each LLM for Dynamic Data. ↑ means the performance is slightly better than that under Answer-Only setting, ↓ means worse, and ~ means the performance is nearly the same.

Claude-series models achieve good results, while the other two models are not. The detailed scores models reached in each experiment are shown in Appendix C. However, although some models could get a great performance, during experiments, we find that these models are good at experimental principles and designing, while when it comes to analyzing the experiment results, the performances are not satisfying.

**CoT Setting and 3-Shot setting** Comparison of experiment results among Answer-Only, Chain-of-Thought and 3-Shot settings are shown in Figure 5 and Table 5.<sup>9</sup> And we refer detailed results to Appendix A and B.

The experimental results from Static Data reveal that solely the GPT-series LLMs get performance enhancement within the CoT setting due to the limited CoT capabilities of other LLMs. As for the 3-Shot setting, roughly half of the LLMs analyzed demonstrate superior performances relative to the Answer-Only setting. The performances of the remaining LLMs are closely similar to those observed within the Answer-Only setting.

From the experimental results of Dynamic Data, it is observed that both CoT and 3-Shot significantly enhance the performance of most Language Learning Models (LLMs) in the chemistry subset. However, the performances achieved are still not up to the mark. In the physics subset, the impact of CoT and 3-Shot on most LLMs is less pronounced, resulting in nearly random performances. Under the CoT setting, GPT-3.5-turbo achieves an accuracy of 47.19, suggesting a robust understanding of physical principles. Conversely, the performance of GPT-4 is markedly poor, from which we find that despite its extensive knowledge of physical principles, it frequently employs incorrect formulas to solve problems. Nevertheless, GPT-4 attains an accuracy of 51.01 under 3-Shot setting, the highest among all models, demonstrating its ability to learn from a mere three examples.

<sup>9</sup>When evaluating on CoT and 3-Shot settings, Claude-Instant and Claude are not available for us, due to the limitation of API.

### 4.3 Discussion

#### Training on large-scale scientific corpus is helpful.

Based on experimental results (Table 4), Galactica (Taylor et al. 2022), which has been trained on an extensive scientific corpus, significantly outperforms other LLMs with a comparable number of parameters, although Galactica is trained with a much smaller amount of data. Remarkably, when tested on Dynamic Data, Galactica surpasses the GPT-series and Claude-series LLMs in computational problems.

#### Most LLMs perform bad on calculation problems, especially in physics domain.

Detailed results across various knowledge domains on Static Data (refer to Appendix B) reveal that most LLMs underperform in the Scientific Calculation domain, while demonstrate relatively superior performance in other domains, which is particularly acute in the field of physics. Similar issues are also observed in Dynamic Data and Experimental Data. In the context of Dynamic Data, the mean square error, employed to evaluate calculation abilities within the chemistry subset, is exceedingly high for most LLMs, and almost all LLMs can only achieve nearly random performance within the physics subset. Regarding Experimental Data, our findings indicate that these LLMs struggle with the analysis of experimental results.

## 5 Conclusion

In this paper, we introduce SciEval, a benchmark designed to evaluate scientific capabilities of LLMs. SciEval comprises about 18,000 challenging scientific questions, covering three fundamental fields of science. SciEval assesses the scientific ability of LLMs across four dimensions. It incorporates both objective and subjective questions, and employs dynamic data generation to mitigate potential data leakage. We conduct comprehensive experiments on various advanced LLMs using SciEval and perform thorough analyses. Our experimental results reveal that most LLMs do not perform well on our benchmark, with the exception of the GPT-series and Claude-series LLMs. We hope that SciEval can serve as a robust benchmark for assessing scientific capabilities of LLMs.

## Acknowledgements

This work is funded by the China NSFC Projects (92370206, U23B2057, 62106142 and 62120106006) and Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102).

## References

- Blanco-Gonzalez, A.; Cabezon, A.; Seco-Gonzalez, A.; Conde-Torres, D.; Antelo-Riveiro, P.; Pineiro, A.; and Garcia-Fandino, R. 2023. The role of ai in drug discovery: challenges, opportunities, and strategies. *Pharmaceuticals*, 16(6): 891.
- Chang, Y.; Wang, X.; Wang, J.; Wu, Y.; Zhu, K.; Chen, H.; Yang, L.; Yi, X.; Wang, C.; Wang, Y.; et al. 2023. A survey on evaluation of large language models. *arXiv preprint arXiv:2307.03109*.
- Du, Z.; Qian, Y.; Liu, X.; Ding, M.; Qiu, J.; Yang, Z.; and Tang, J. 2022. GLM: General Language Model Pretraining with Autoregressive Blank Infilling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 320–335.
- Forehand, M. 2010. Blooms taxonomy. *Emerging perspectives on learning, teaching, and technology*, 41(4): 47–56.
- Frey, N.; Soklaski, R.; Axelrod, S.; Samsi, S.; Gomez-Bombarelli, R.; Coley, C.; and Gadepally, V. 2022. Neural scaling of deep chemical models.
- Guo, T.; Guo, K.; Liang, Z.; Guo, Z.; Chawla, N. V.; Wiest, O.; Zhang, X.; et al. 2023. What indeed can GPT models do in chemistry? A comprehensive benchmark on eight tasks. *arXiv preprint arXiv:2305.18365*.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Hendrycks, D.; Burns, C.; Kadavath, S.; Arora, A.; Basart, S.; Tang, E.; Song, D.; and Steinhardt, J. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Huang, Y.; Bai, Y.; Zhu, Z.; Zhang, J.; Zhang, J.; Su, T.; Liu, J.; Lv, C.; Zhang, Y.; Lei, J.; et al. 2023. C-eval: A multi-level multi-discipline chinese evaluation suite for foundation models. *arXiv preprint arXiv:2305.08322*.
- Jin, D.; Pan, E.; Oufattole, N.; Weng, W.-H.; Fang, H.; and Szolovits, P. 2021. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Applied Sciences*, 11(14): 6421.
- Jin, Q.; Dhingra, B.; Liu, Z.; Cohen, W. W.; and Lu, X. 2019. Pubmedqa: A dataset for biomedical research question answering. *arXiv preprint arXiv:1909.06146*.
- Kojima, T.; Gu, S. S.; Reid, M.; Matsuo, Y.; and Iwasawa, Y. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35: 22199–22213.
- Krathwohl, D. R. 2002. A revision of Bloom’s taxonomy: An overview. *Theory into practice*, 41(4): 212–218.
- Liang, P.; Bommasani, R.; Lee, T.; Tsipras, D.; Soylu, D.; Yasunaga, M.; Zhang, Y.; Narayanan, D.; Wu, Y.; Kumar, A.; et al. 2022. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*.
- Lu, P.; Mishra, S.; Xia, T.; Qiu, L.; Chang, K.-W.; Zhu, S.-C.; Tafjord, O.; Clark, P.; and Kalyan, A. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35: 2507–2521.
- Luo, R.; Sun, L.; Xia, Y.; Qin, T.; Zhang, S.; Poon, H.; and Liu, T.-Y. 2022. BioGPT: generative pre-trained transformer for biomedical text generation and mining. *Briefings in Bioinformatics*, 23(6): bbac409.
- OpenAI. 2023. GPT-4 Technical Report. *arXiv:2303.08774*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35: 27730–27744.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 311–318.
- Schulman, J.; Zoph, B.; Kim, C.; Hilton, J.; Menick, J.; Weng, J.; Uribe, J. F. C.; Fedus, L.; Metz, L.; Pokorny, M.; et al. 2022. ChatGPT: Optimizing language models for dialogue. *OpenAI blog*.
- Singhal, K.; Azizi, S.; Tu, T.; Mahdavi, S. S.; Wei, J.; Chung, H. W.; Scales, N.; Tanwani, A.; Cole-Lewis, H.; Pfohl, S.; et al. 2023. Large language models encode clinical knowledge. *Nature*, 1–9.
- Srivastava, A.; Rastogi, A.; Rao, A.; Shoeb, A. A. M.; Abid, A.; Fisch, A.; Brown, A. R.; Santoro, A.; Gupta, A.; Garriga-Alonso, A.; et al. 2022. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*.
- Sun, T.; Zhang, X.; He, Z.; Li, P.; Cheng, Q.; Yan, H.; Liu, X.; Shao, Y.; Tang, Q.; Zhao, X.; Chen, K.; Zheng, Y.; Zhou, Z.; Li, R.; Zhan, J.; Zhou, Y.; Li, L.; Yang, X.; Wu, L.; Yin, Z.; Huang, X.; and Qiu, X. 2023. MOSS: Training Conversational Language Models from Synthetic Data.
- Taori, R.; Gulrajani, I.; Zhang, T.; Dubois, Y.; Li, X.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023. Stanford alpaca: An instruction-following llama model.
- Taylor, R.; Kardas, M.; Cucurull, G.; Scialom, T.; Hartshorn, A.; Saravia, E.; Poulton, A.; Kerkez, V.; and Stojnic, R. 2022. GALACTICA: A Large Language Model for Science.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- WANG, F.; and MIAO, Q. 2023. Novel Paradigm for AI-driven Scientific Research: From AI4S to Intelligent Science. *Bulletin of Chinese Academy of Sciences (Chinese Version)*, 38(4): 536–540.

Wang, X.; Hu, Z.; Lu, P.; Zhu, Y.; Zhang, J.; Subramaniam, S.; Loomba, A. R.; Zhang, S.; Sun, Y.; and Wang, W. 2023. SciBench: Evaluating College-Level Scientific Problem-Solving Abilities of Large Language Models. *arXiv preprint arXiv:2307.10635*.

Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E.; et al. 2023. Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. *arXiv preprint arXiv:2306.05685*.

Zhong, W.; Cui, R.; Guo, Y.; Liang, Y.; Lu, S.; Wang, Y.; Saied, A.; Chen, W.; and Duan, N. 2023. Agieval: A human-centric benchmark for evaluating foundation models. *arXiv preprint arXiv:2304.06364*.