

SciCustom: A Framework for Custom Evaluation of Scientific Capabilities in Large Language Models

Yiyang Gu^{1*}, Junwei Yang^{1*}, Junyu Luo¹, Ye Yuan¹, Bin Feng³, Yingce Xia², Shufang Xie², Kaili Liu⁴, Bohan Wu¹, Qi Shi⁵, Haoran Li⁵, Beier Xiao⁵, Zhiping Xiao^{6†}, Xiao Luo⁷, Weizhi Zhang⁸, Philip S. Yu⁸, Zequn Liu^{2†}, Ming Zhang^{1†}

¹ State Key Laboratory for Multimedia Information Processing, School of Computer Science, PKU-Anker LLM Lab, Peking University ² Zhongguancun Academy ³ IDEA

⁴ Xidian University ⁵ Peking University ⁶ University of Washington

⁷ University of Wisconsin–Madison ⁸ University of Illinois Chicago

{yiyangu, yjwtheonly, mzhang_cs}@pku.edu.cn, patxiao@uw.edu, liuzequn@bza.edu.cn

Abstract

Large language models (LLMs) are increasingly applied to scientific research, yet existing evaluations often fail to reflect the fine-grained capabilities required in practice. Most benchmarks are manually curated or domain-generic, limiting scalability and alignment with real scientific use cases. In this paper, we propose a new framework named SCICUSTOM to address the problem. It enables the custom construction of benchmarks from large-scale scientific data to evaluate application-specific scientific capabilities in LLMs. SCICUSTOM first organizes scientific knowledge into ontology-grounded knowledge units with controlled granularity and trains a tagger to map large-scale data instances into this knowledge space. Given a custom requirement, relevant knowledge units are identified via voting-based multi-model consensus. These units enable relevance-aware benchmark retrieval via binary search, followed by proxy subset selection and data-grounded benchmark generation for efficient evaluation. Experiments in chemistry and healthcare demonstrate that SCICUSTOM reveals fine-grained differences in LLM scientific capabilities that standard benchmarks overlook, while requiring neither expert annotation nor synthetic question generation. This work provides a scalable and application-aware foundation for benchmarking scientific capabilities in LLMs. The source code is available at <https://github.com/yjwtheonly/SciCustom>.

1 Introduction

With the rapid advancement of Large Language Models (LLMs), their applications have significantly broadened across the scientific landscape (Chang et al., 2024; Bommasani et al., 2021; Birhane et al., 2023; Luo et al., 2025; Yuan et al.,

*Equal contribution with order determined by flipping a coin.

†Corresponding authors.

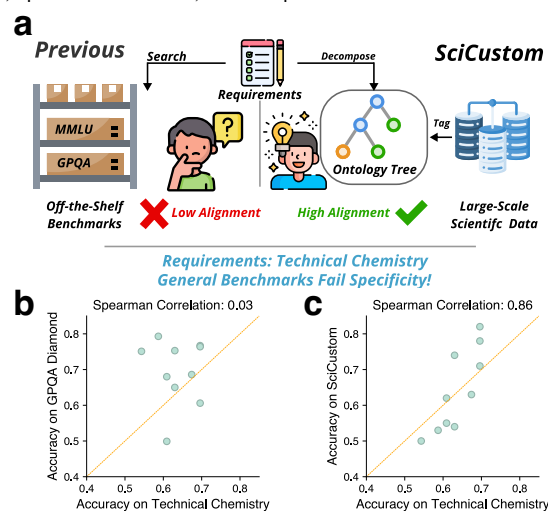


Figure 1: Illustrations of SCICUSTOM. (a) Comparison between traditional off-the-shelf benchmarking and our ontology-driven framework. (b, c) Evaluation of 10 LLMs on different benchmarks, where each dot represents a model. Targeting specific capabilities in Technical Chemistry, (b) the general scientific benchmark (GPQA Diamond) aligns poorly with expert ground truth, whereas (c) the benchmark constructed by SCICUSTOM demonstrates strong alignment.

2024; Liu et al., 2023; Xia et al., 2025). For example, an LLM-based system successfully proposed a novel mechanism of gene transfer in bacteria, mirroring conclusions that took years of experimental validation (Gottweis et al., 2025; Penadés et al., 2025). However, as scientists increasingly seek to leverage LLMs for specific scientific applications, a critical challenge emerges: the performance of a given LLM within a particular scientific context remains largely uncertain and difficult to evaluate (Cai et al., 2025; Singhal et al., 2025; Miret and Krishnan, 2025; Bedi et al., 2025).

While numerous benchmarks exist to evaluate various aspects of model capability, they frequently fail to reflect the requirements of specialized usage (Anjum et al., 2025; Bommasani et al., 2023; Bandel et al., 2024). Empirical observations in-

dicating that performance on widely-used benchmarks often diverges from that in specific scientific tasks (**Figure 1b**), necessitating the customization of benchmarks tailored to particular scientific applications (Singhal et al., 2025; Miret and Krishnan, 2025). Given the vast and ever-expanding range of scientific applications, manual curation of benchmarks for each emerging use case is impractical (Huang et al., 2021; Wu et al., 2018).

A plausible solution is to automate the customized benchmark construction process (Farchi et al., 2024; Cohrs et al., 2025; Boiko et al., 2023; M. Bran et al., 2024). However, developing such automated framework presents two challenges. First, scientific applications are inherently complex and highly interdisciplinary. A single application often requires knowledge from multiple subfields (e.g., drug discovery intertwines organic chemistry, molecular biology, and pharmacology (Lu et al., 2025)), while distinct applications frequently share underlying knowledge (e.g., both drug discovery and clinical decision require the knowledge of pharmacology (Ong et al., 2025; Pathak et al., 2008)). Consequently, constructing benchmarks from scratch for each specific scenario results in massive redundant labor and limits the framework’s scalability (Huang et al., 2021; M. Bran et al., 2024). Second, scientific evaluation requires grounded validity. High-quality evaluation data are frequently derived from costly wet-lab experiments or computational simulations (Bioscience, 2025), rendering simple general-purpose data synthesis methods infeasible (Parker, 2020; Wu et al., 2018; Huang et al., 2021; Miret and Krishnan, 2025).

To this end, we introduce SCICUSTOM, a framework for custom evaluation of scientific capabilities in LLMs (**Figure 1a,c**). The key idea is that a complex scientific application can be approximated by a composition of fine-grained knowledge units. To construct these units, we select 642 representative concepts from a comprehensive scientific ontology to serve as unified knowledge units. Leveraging the rich hierarchical information within the ontology, which links specific scientific concepts to knowledge units, we train a tagger to automatically align real-world scientific corpora with the corresponding knowledge units. This offline process performs a once-for-all population of each knowledge unit with grounded scientific data. When a user presents an evaluation requirement, SCICUSTOM identifies the relevant knowledge units and

orchestrates a hierarchical retrieval process to construct a tailored benchmark. The combinability of these pre-constructed knowledge units enables dynamic adaptation to new requirements via efficient reuse, ensuring scalability for diverse downstream applications.

Experimental results in chemistry and healthcare domains demonstrate that our automatically constructed benchmarks closely align with expert-curated benchmarks. Furthermore, for pericyclic reaction, a novel scientific application lacking prior benchmarks, our framework successfully generates a high-quality benchmark, providing a scalable and reliable solution for assessing LLMs in the ever-expanding landscape of scientific research.

Our contributions can be summarized as follows:

- We propose SCICUSTOM, a framework that models scientific applications as compositions of reusable knowledge units to automate benchmark construction, enabling efficient adaptation to diverse downstream tasks.
- We develop a tagger that leverages ontology structures to map large-scale, real-world scientific corpora into knowledge units, ensuring that evaluations are grounded in verifiable data.
- Experiments in chemistry and healthcare confirm SCICUSTOM’s strong alignment with expert-curated benchmarks and its capability to evaluate novel scientific contexts where no prior benchmarks exist.

2 Methodology

2.1 Overview

We study the problem of customized scientific capability evaluation for LLMs. Given an evaluation requirement r and a large corpus of scientific data \mathcal{D} , the goal is to construct a benchmark \mathcal{B}_r that reflects the capabilities required by r without expert annotation.

SCICUSTOM is an automated framework for this problem, operating through offline data indexing and online data-grounded generation phases. In the offline phase, SCICUSTOM represents large-scale scientific data within shared knowledge units using a tagger. This phase organizes scattered scientific content into reusable units. In the online phase, SCICUSTOM first employs a voting-based multi-model consensus to identify knowledge units

relevant to the user requirement r . Guided by these units, the framework performs a hierarchical data selection process: retrieving candidate data, filtering via binary search, and extracting a representative proxy subset to generate final questions for efficient evaluation. An overview of SCICUSTOM is depicted in **Figure 2**.

2.2 Ontology-Grounded Knowledge Units

To enable fine-grained yet reusable modeling of scientific knowledge, SCICUSTOM is built upon a structured scientific ontology that serves as a unified semantic backbone for both knowledge definition and data organization.

Following Liu et al. (2021), we integrate multiple authoritative repositories to construct a large-scale scientific ontology, organized into a collection of directed acyclic graphs (DAGs) spanning 227 scientific subdisciplines. Within each DAG, concepts are interconnected via *is-a* relations, where descendant nodes represent finer-grained specializations of their parent concepts (e.g., *Organic Chemistry* is a descendant node of *Chemistry*.)

We select concepts from this ontology as knowledge units. The knowledge units should be neither overly coarse nor excessively specific. Overly coarse concepts lack the resolution to distinguish fine-grained model capabilities, while excessively specific ones hinder the effective reuse of knowledge units for broader applications. We empirically select concepts at a granularity comparable to “textbook chapter titles”. We therefore capture these knowledge units via a depth-first traversal over each ontology DAG \mathcal{G}_i . At each node v , an LLM evaluates whether the ontology term aligns with the granularity criteria based on the term name, illustrative examples, and its prior knowledge of scientific taxonomy.

The overall procedure is summarized in Algorithm 1, and the detailed prompting strategy is provided in Appendix G. Through this procedure, we obtain 641 scientific knowledge units. We additionally introduce a dedicated NON-SCIENTIFIC unit, resulting in a total of 642 knowledge units. These units define the ontology-grounded scientific knowledge space that underpins all subsequent benchmark construction in SCICUSTOM.

2.3 Mapping Data to Knowledge Units

To populate the identified knowledge units with grounded scientific data, we collect a large-scale corpus containing high-quality question-answering

Algorithm 1 Ontology-Guided Knowledge Unit Selection

Input: Ontology $\{\mathcal{G}_i\}$, **Output:** Units \mathcal{T}

- 1: $\mathcal{T} \leftarrow \emptyset$
- 2: **for each** node v visited by DFS in $\{\mathcal{G}_i\}$ **do**
- 3: **if** $|\text{Desc}(v)| < 10$ **then**
- 4: Backtrack
- 5: **end if**
- 6: $\text{label} \leftarrow$ LLM classifies v as *coarse*, *moderate*, or *fine*
- 7: **if** label is *coarse* **then**
- 8: Continue traversal (recurse into children)
- 9: **else if** label is *moderate* **then**
- 10: $\mathcal{T} \leftarrow \mathcal{T} \cup \{v\}$
- 11: **else**
- 12: Backtrack {Prune branch if *fine*}
- 13: **end if**
- 14: **end for**
- 15: **return** \mathcal{T}

(QA) pairs across diverse scientific domains. The corpus comprises a total of N data instances, each defined as $d = \langle q, a \rangle$, where q denotes a natural language question and a the corresponding reference answer. These instances are originally scattered, necessitating the development of a specialized scientific tagger to organize them into knowledge units.

To enable efficient large-scale annotation, we train a small language model as the tagger. The model maps each query q to a subset of knowledge units, $\mathcal{T}_q \subseteq \mathcal{T}$. This design implements a once-for-all mapping between data and knowledge units during the construction phase, allowing the same annotated corpus to be reused across diverse evaluation requirements.

Supervised data for training the tagger is constructed using a two-stage strategy. In the first stage, we generate synthetic scientific queries via controlled compositional sampling over the knowledge units \mathcal{T} . Specifically, we randomly sample between 1 and 5 knowledge units to form a target set $\mathcal{T}_q \subset \mathcal{T}$. We fully leverage the rich hierarchical information of the ontology: for each selected unit $v \in \mathcal{T}_q$, we aggregate all its descendant nodes to form a representative keyword set \mathcal{K}_v (e.g., *Ventricular Septum* is a keyword for the unit *Anatomical Entities*). We further sample keywords from each \mathcal{K}_v and combine them to form a composite prompt for an LLM to generate a synthetic query q (**Figure 3**). Each generated query is then paired with its

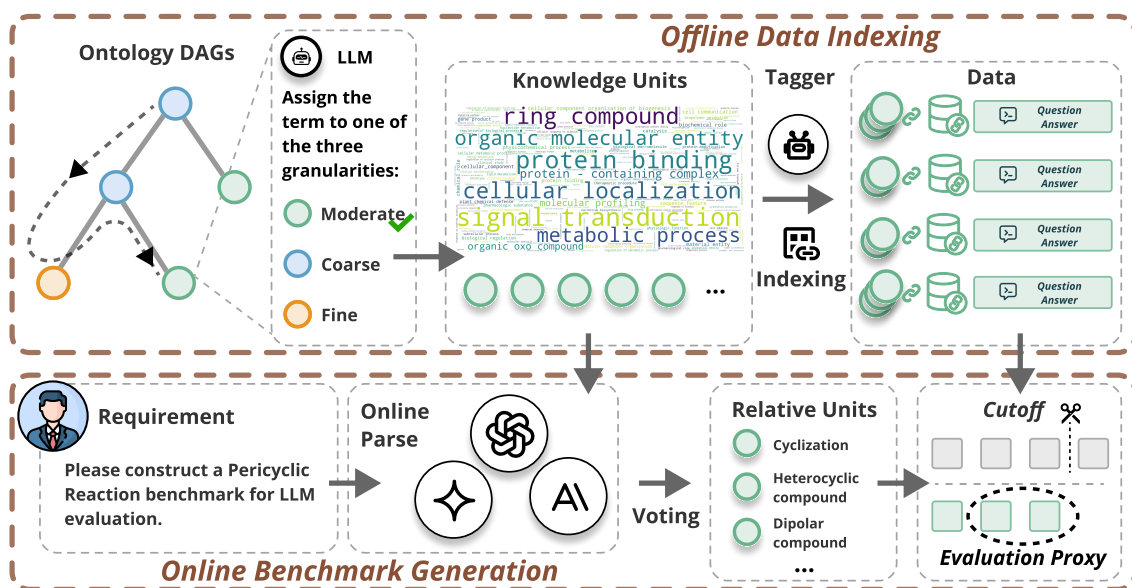


Figure 2: Framework of SCICUSTOM. It consists of an offline phase where scientific data is indexed into ontology-grounded knowledge units via a trained tagger, and an online phase where user requirements are parsed by multi-model voting to identify relevant tags. These tags guide the binary search-based selection and proxy selection of data. The problem set of the benchmark is generated based on these data.

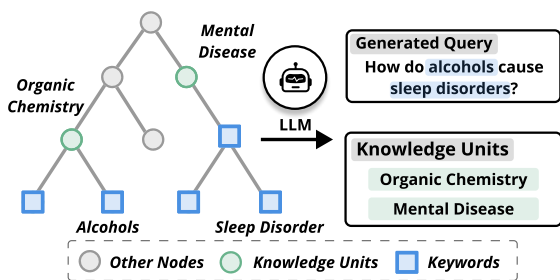


Figure 3: Illustration of the synthetic data construction pipeline for tagger training. We sample knowledge units (blue squares) and extract descendant keywords (green circles) from the ontology. An LLM then generates a natural language query based on these keywords, creating a labeled training instance.

corresponding set of source knowledge units, yielding a labeled training instance of the form $\langle q, \mathcal{T}_q \rangle$. Using this procedure, we construct synthetic scientific training instances covering complex compositional knowledge patterns. In the second stage, we complement the synthetic data with real-world scientific queries sampled from existing instruction-tuning datasets. These queries are annotated with their corresponding knowledge units with an LLM. In addition, we include non-scientific queries to model boundary cases between scientific and non-scientific content.

By introducing a small and efficient tagger for annotating scientific queries, SCICUSTOM achieves scalable and interpretable alignment between natural language queries and scientific knowledge, forming a key foundation for cus-

Algorithm 2 Binary Search for Relevant Benchmark Construction

Input: Sorted list $L_{\mathcal{D}'_r}$, Requirement r , Models \mathcal{M}

Output: Data \mathcal{D}_r

- 1: $low \leftarrow 0, high \leftarrow |L_{\mathcal{D}'_r}| - 1, cutoff \leftarrow 0$
- 2: **while** $low \leq high$ **do**
- 3: $mid \leftarrow low + \lfloor (high - low) / 2 \rfloor; d_{mid} \leftarrow L_{\mathcal{D}'_r}[mid]$
- 4: $vote \leftarrow \sum_{M_i \in \mathcal{M}} \mathbb{I}(M_i \text{ judges } d_{mid} \text{ relevant to } r)$
- 5: **if** $vote > |\mathcal{M}| / 2$ **then**
- 6: $cutoff \leftarrow mid; low \leftarrow mid + 1$
- 7: **else**
- 8: $high \leftarrow mid - 1$
- 9: **end if**
- 10: **end while**
- 11: **return** $\mathcal{D}_r = L_{\mathcal{D}'_r}[0 : cutoff]$

tomized scientific evaluation.

2.4 Voting-Based Knowledge Unit Selection

With the large-scale scientific corpus successfully grounded to ontology-defined knowledge units via the tagger, we possess a structured knowledge base ready for dynamic retrieval. To construct a benchmark tailored to a specific user evaluation requirement r , the critical first step is to accurately identify the subset of knowledge units $\mathcal{T}_r \subset \mathcal{T}$ that underpins the requirement.

To identify the target knowledge units \mathcal{T}_r , we

employ a multi-model consensus mechanism where a set of heterogeneous LLMs independently rank candidate units based on their relevance to r . We aggregate these rankings by calculating the average rank position for each unit across all models, and select the top- K_1 units with the lowest consensus ranks (indicating highest relevance) to form the target set \mathcal{T}_r .

2.5 Hierarchical Benchmark Generation

Binary Search-Based Data Selection. With the requirement-relevant knowledge unit set \mathcal{T}_r established, we aim to retrieve a dataset \mathcal{D}_r from the massive corpus \mathcal{D} that best matches r . For any data sample $\langle q, a \rangle \in \mathcal{D}$, let $\mathcal{T}_{r,q} = \mathcal{T}_r \cap \mathcal{T}_q$ represent the intersection between the data’s inherent knowledge units and the query knowledge units.

Intuitively, we construct a candidate set $\mathcal{D}'_r = \{\langle q, a \rangle \in \mathcal{D} \mid |\mathcal{T}_{r,q}| \geq 1\}$. However, a simple intersection filter is insufficient to guarantee high semantic alignment. While employing LLMs to verify the relevance of each candidate against r would ensure quality, applying this to the entire set is computationally prohibitive. To optimize API usage without compromising retrieval quality, we adopt a binary search-based strategy that relies on a pre-sorted candidate list. Specifically, we define an ordering where data sample $d_i = \langle q_i, a_i \rangle$ is ranked higher than $d_j = \langle q_j, a_j \rangle$ if:

1. $|\mathcal{T}_{r,q_i}| > |\mathcal{T}_{r,q_j}|$; or
2. $|\mathcal{T}_{r,q_i}| = |\mathcal{T}_{r,q_j}|$ and $\bar{\mathcal{R}}_r(d_i) < \bar{\mathcal{R}}_r(d_j)$,

where $\mathcal{T}_{r,q} = \mathcal{T}_r \cap \mathcal{T}_q$ denotes the matching knowledge units, and $\bar{\mathcal{R}}_r(d)$ represents average rank of the knowledge units in $\mathcal{T}_{r,q}$ acquired in Sec. 2.4.

We assume that within the prioritized candidate list, the relevance to the requirement r follows a generally monotonic decreasing trend. Under this assumption, we apply the binary search procedure to efficiently identify the cutoff point—defined as the last position where a majority of the LLM ensemble still judges the sample as relevant to r . This approach enables us to determine the maximal relevant data subset without linearly verifying every instance, significantly reducing the number of required oracle judgments to $O(\log(|\mathcal{D}'_r|))$. The detailed procedure is outlined in Algorithm 2.

Efficient Proxy Subset Selection. The constructed dataset \mathcal{D}_r may still be excessively large (e.g., $|\mathcal{D}_r| > 200,000$ in our experiments). To enable efficient evaluation, we propose a selection

strategy to extract a proxy subset $\mathcal{P}_r \subseteq \mathcal{D}_r$ such that $|\mathcal{P}_r| \leq K_2$, while maintaining the evaluative power of the full set.

For each data sample $\langle q, a \rangle \in \mathcal{D}_r$, we compute two commonly used intrinsic metrics (Saranathan et al., 2025): hardness score characterizing the difficulty and quality score indicating the linguistic quality. Our objective is to select \mathcal{P}_r such that the distributions of these scores closely approximate those of \mathcal{D}_r . We formulate this as minimizing the cumulative Wasserstein distance between the distributions of the subset and the full set. To solve the sampling of \mathcal{P}_r efficiently, we adopt a cluster-based strategy following SubLIME (Saranathan et al., 2025). Intuitively, we encode all pairs $d = \langle q, a \rangle$ into an embedding space using RoBERTa (Liu et al., 2019). We then perform clustering in embedding space and sample representatives from each cluster. Detailed information is in Appendix A.

Data-Grounded Benchmark Generation. To facilitate standardized and deterministic evaluation, we further transform the raw query-answer pairs in the selected proxy subset \mathcal{P}_r into multiple-choice questions (MCQs). We utilize LLM to generate plausible distractors for each instance. The final output is a structured, efficient benchmark \mathcal{B}_r composed of high-quality MCQs, ready for immediate automated model assessment.

3 Experiments

We established a comprehensive experimental framework to evaluate the effectiveness of SCICUSTOM. First, we benchmarked our approach against groundtruth expert-curated datasets across 11 specific evaluation requirements. We focused on measuring the alignment of model rankings produced by SCICUSTOM versus those from human experts. Furthermore, to verify the framework’s utility in real-world exploration, we applied SCICUSTOM to Pericyclic Reaction, a novel evaluation requirement without existing benchmarks, and qualitatively analyzed the validity of the constructed benchmark.

3.1 Experimental Setup

Ground-Truth Benchmarks. We evaluated SCICUSTOM on 6 tasks in chemistry, including analytical chemistry, inorganic chemistry, material science, organic chemistry, physical chemistry and technical chemistry, and 5 tasks in healthcare, including virology, human aging, medical genetics, anatomy and nutrition. The groundtruth bench-

Table 1: Ranking consistency analysis of 10 LLMs on 6 chemistry-specific benchmarks. The table compares SCICUSTOM against baseline benchmarks in preserving model rankings, quantified by Spearman and Kendall correlation coefficients. Higher scores indicate better consistency. Best results are in **bold**.

Metrics	Benchmarks	Analytical chemistry	Inorganic chemistry	Material science	Organic chemistry	Physical chemistry	Technical chemistry
Spearman correlation	IfBench	-0.32	-0.34	-0.43	0.12	-0.61	-0.54
	SimpleQA	-0.67	-0.69	-0.78	-0.48	-0.42	-0.31
	GPQA	0.61	0.52	0.21	0.72	0.21	0.03
	MMLU	0.21	0.27	-0.61	0.21	0.52	0.31
	MedQA	0.31	0.42	0.01	-0.21	0.63	0.72
	GPT-5	-0.11	0.05	-0.04	0.38	0.24	-0.07
	Embedding	-0.34	-0.59	-0.39	0.11	-0.73	-0.41
	SCICUSTOM	0.86	0.67	0.42	0.89	0.74	0.86
Kendall correlation	IfBench	-0.14	-0.19	-0.29	0.10	-0.43	-0.43
	SimpleQA	-0.52	-0.49	-0.56	-0.39	-0.24	-0.23
	GPQA	0.39	0.39	0.10	0.42	0.14	-0.05
	MMLU	0.14	0.18	-0.49	0.10	0.43	0.24
	MedQA	0.24	0.40	-0.09	-0.14	0.43	0.48
	GPT-5	-0.15	0.00	-0.06	0.35	0.16	0.00
	Embedding	-0.31	-0.52	-0.31	0.09	-0.62	-0.32
	SCICUSTOM	0.62	0.48	0.35	0.68	0.57	0.65

Table 2: Ranking consistency analysis of 10 LLMs on 6 healthcare-specific benchmarks. Higher scores indicate better consistency. Best results are in **bold**.

Metrics	Benchmarks	Virology	Human aging	Medical genetics	Anatomy	Nutrition
Spearman correlation	IfBench	0.04	-0.32	-0.64	-0.26	-0.11
	SimpleQA	0.35	-0.50	0.00	-0.37	0.11
	GPQA	-0.11	-0.10	-0.09	0.48	0.18
	MMLU	0.35	0.21	0.04	-0.15	0.56
	MedQA	0.44	0.62	0.35	-0.19	0.45
	GPT-5	0.25	0.20	0.09	0.11	0.52
	Embedding	0.18	0.21	-0.21	-0.32	0.27
	SCICUSTOM	0.55	0.49	0.42	0.62	0.78
Kendall correlation	IfBench	-0.05	-0.14	-0.43	-0.21	-0.00
	SimpleQA	0.31	-0.33	0.00	-0.21	0.10
	GPQA	-0.05	-0.10	-0.11	0.32	0.10
	MMLU	0.31	0.14	0.00	0.00	0.51
	MedQA	0.35	0.42	0.23	-0.10	0.31
	GPT-5	0.20	0.18	0.06	0.06	0.43
	Embedding	0.11	0.10	-0.14	-0.16	0.18
	SCICUSTOM	0.45	0.37	0.33	0.51	0.64

marks for chemistry tasks are from *ChemBench* (Mirza et al., 2024), while the healthcare tasks are sourced from the health subset of *MMLU-Pro* (Wang et al., 2024b).

Baselines. We compared SCICUSTOM-curated benchmarks against a comprehensive set of baselines spanning general-purpose benchmarks, scientific benchmarks and alternative benchmark construction methods. To assess the correlation between general-purpose benchmarks and target tasks, we utilized *IfBench* (Pyatkin et al., 2025) and *SimpleQA* (Wei et al., 2024). *GPQA-Diamond* (Rein et al., 2024) was included to compare general scientific benchmarks. For domain-specific comparisons, we selected the chemistry-related subset of *MMLU-Pro* (Wang et al., 2024b) for chemistry tasks and *MedQA* (Jin et al., 2021) for healthcare tasks. Furthermore, to isolate the contributions of our grounded data and ontology-driven design, we introduced two alternative construction baselines: *GPT-5*, a fully synthetic benchmark where

the model generates multiple-choice questions directly without access to grounded data, and *Embedding*, an embedding-based baseline that selects data instances from the corpus \mathcal{D} solely via k-nearest neighbor (k-NN) retrieval to the requirement, bypassing the ontology-grounded knowledge units.

Evaluation Protocol. We first evaluate 10 representative LLMs on the ground-truth benchmark to obtain the reference ranking. We then evaluate the same 10 LLMs using benchmarks constructed by each baseline method (including ours) and compute the ranking induced by each alternative benchmark. Finally, we measure the consistency between the baseline-induced ranking and the ground-truth ranking using Spearman and Kendall correlation coefficients. Further implementation details are provided in Appendix C.

3.2 Alignment Between SCICUSTOM and Expert-Curated Benchmarks

We report the results of SCICUSTOM against various baselines in **Table 1** and **Table 2**. SCICUSTOM exhibits a robust correlation with expert assessments, achieving the highest Spearman correlation in 10 out of 11 tasks. This strong alignment indicates that the benchmarks generated by SCICUSTOM effectively capture the capabilities required for specific applications, serving as a reliable proxy for expensive expert evaluation. Notably, SCICUSTOM selects the same top-1 model as the ground-truth benchmark on 8 out of 11 tasks, highlighting its high practical utility in assisting users to identify the optimal model for their specific needs.

In contrast, widely adopted benchmarks fails to show consistent alignment with expert rankings. This underscores that neither general instruction

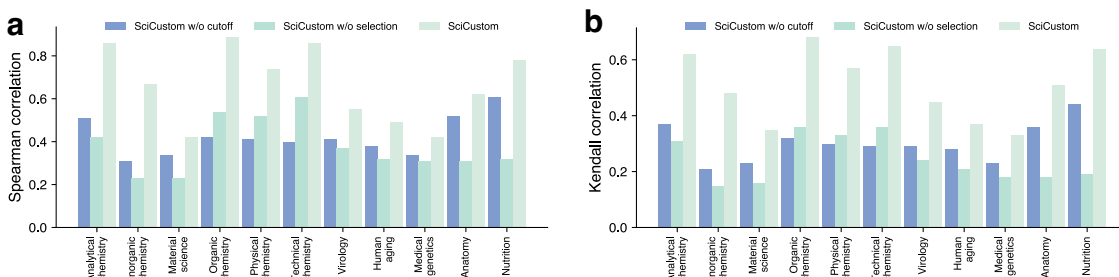


Figure 4: Bar plot showing the effectiveness of relevance cutoff and subset selection strategies.

following skills nor broad scientific reasoning capabilities are reliable predictors of proficiency for specialized scientific requirements. Furthermore, established domain-specific benchmarks also struggle to predict sub-domain-level performance due to insufficient granularity. These findings further necessitate the SCICUSTOM framework, which dynamically constructs benchmarks at the precise granularity of the target requirement.

Comparing SCICUSTOM with alternative construction baselines reveals the critical role of our methodology design. First, our framework significantly outperforms the fully synthetic *GPT-5* baseline, underscoring the necessity of grounded scientific data. Second, SCICUSTOM also surpasses the *Embedding* baseline, which retrieves grounded data via simple semantic similarity (k -NN) rather than knowledge-unit-based mapping. This demonstrates that: by organizing data into knowledge units, we achieve better alignment with user intent than unstructured semantic retrieval. Additional comparison results for *GPT-5-RAG* and *Embedding-filtering* are provided in Appendix D.

Special-Case Analysis. We observed a relatively lower Spearman correlation (0.42) in Material Science compared to other tasks. To investigate this discrepancy, we further evaluated the LLMs using MatSci-NLP (Song et al., 2023), another large-scale expert-curated benchmark in this field (comprising 9,466 classification questions grounded in experimental material science facts). Interestingly, we found that the alignment between the two expert-curated benchmarks themselves was also limited ($\rho = 0.31$, $\tau_b = 0.22$). This suggests that Material Science is an exceptionally broad and multifaceted discipline (Shoghi et al., 2023; Wang et al., 2025), where different evaluation protocols may focus on orthogonal knowledge areas or distinct reasoning types. We identify the fine-grained unification of such complex, heterogeneous domains as a direction for future research.

Human Evaluation. We evaluate the quality of the benchmarks generated by SCICUSTOM using a random sample of 50 chemistry-related questions that are originally not in multiple-choice format. Three Master students in AI4Chemistry annotate the samples with binary scores for *Correctness* (the validity of the MCQ after transformation) and *Relevance* (the alignment with the benchmark query). The resulting average scores are 0.92 for Correctness and 0.7 for Relevance, confirming the reliability of our benchmark generation pipeline.

3.3 Component Analysis

Effectiveness of Tagger. We evaluate the tagger on 1,000 unseen queries generated by compositional sampling over the knowledge units \mathcal{T} , where each query is generated based on the keywords corresponding to the knowledge units. We use fuzzy label matching to assess performance, where predicted tags are considered if their normalized Indel similarity with any gold tag exceeds 85%, using the rapidfuzz library¹. The Macro F1 and Micro F1 scores achieved are 75.2% and 78.6%, respectively, demonstrating the effectiveness of the tagger in mapping queries to relevant scientific knowledge.

Effectiveness of Relevance Filtering and Subset Selection Strategies. We evaluate the efficacy of our retrieval and selection components by comparing the full SCICUSTOM framework against two ablation variants: *w/o cutoff* (removing the binary search relevance filter) and *w/o selection* (replacing the intrinsic score-based subset selection with random sampling). The full SCICUSTOM consistently outperforms both variants across all 11 sub-disciplines (Figure 4). The drop in performance for *w/o selection* (random sampling) is substantial, necessitating the data selection component. Furthermore, removing the binary search cutoff (*w/o cutoff*) also leads to a notable degradation in ranking consistency, suggesting that our cutoff strategy ensures the benchmark remains tightly focused on

¹<https://rapidfuzz.github.io/RapidFuzz/>

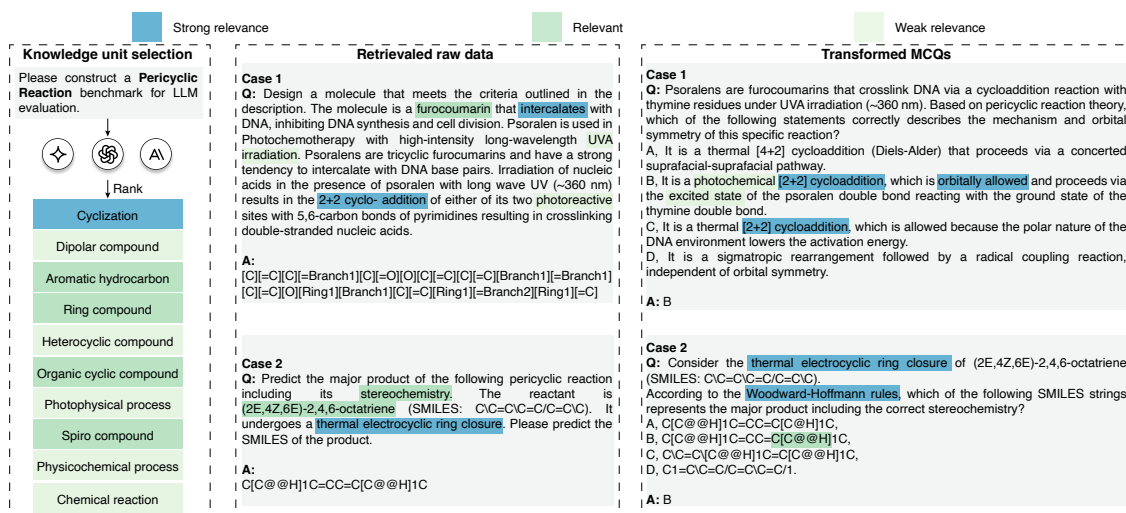


Figure 5: Case study on constructing a benchmark for "Pericyclic Reaction". The pipeline progresses from identifying relevant knowledge units via voting (Left), retrieving grounded raw data with relevance-highlighted terms (Middle), to generating standardized multiple-choice questions for evaluation (Right)

the target capability. These results demonstrate that both precise relevance filtering and representative subset selection are indispensable for achieving efficient and accurate evaluation.

We further report an automatic evaluation of the quality of LLM-generated distractors in MCQs in Appendix F.

3.4 Benchmark for Pericyclic Reaction

To evaluate the capability of SCICUSTOM in handling highly specific and novel scientific requirements, we conducted a case study on *Pericyclic Reactions*. Unlike general organic chemistry queries, pericyclic reactions require a nuanced understanding of orbital symmetry, stereochemistry, and specific reaction conditions (e.g., thermal vs. photochemical) (Woodward and Hoffmann, 1970). No existing LLM benchmark specifically isolates this distinct class of reactions.

As illustrated in Figure 5, the construction process progresses through three key stages, demonstrating high alignment with expert judgment at each step: Upon receiving the abstract requirement "Pericyclic Reaction benchmark," SCICUSTOM first decomposes it into fine-grained knowledge anchors. As shown in the left panel, SCICUSTOM successfully prioritizes highly relevant concepts. The ranking produced by SCICUSTOM highly correlates with expert assessment, where the top-ranked units, such as *Cyclization*, *Aromatic hydrocarbon*, and *Ring compound*, are identified by experts as the most scientifically relevant (indicated by darker shades). Guided by these concepts, SCI-

CUSTOM proceeds to data retrieval and question transformation (Figure 5, middle and right panels). Both the retrieved raw data and the transformed multiple-choice questions exhibit high alignment with expert requirements, accurately capturing critical domain-specific terminology (highlighted in the figure) and complex reaction mechanisms.

4 Related Works

Existing scientific evaluation benchmarks primarily fall into two categories. General scientific benchmarks (Hendrycks et al., 2020; Wang et al., 2024b; Rein et al., 2024; Sun et al., 2024; Wang et al., 2024a) focus predominantly on evaluating broad reasoning capabilities or generalized scientific common sense, lacking the depth required for specific scientific sub-domains. Conversely, domain-specific benchmarks (Mirza et al., 2024; Jin et al., 2021; Singhal et al., 2023; Song et al., 2023) offer more specialized evaluations but are inherently static and labor-intensive to maintain due to their reliance on manual expert curation.

To address the limitations of manual curation, there are several automated benchmarking frameworks in CV (Zhang et al., 2024b) and general NLP (Pombal et al., 2025; Li et al., 2025; Guo et al., 2025). They are ill-suited for complex scientific evaluation, as they typically suffer from restricted requirement scopes (Zhang et al., 2024b), a lack of grounded data (Pombal et al., 2025), or the inability to reuse compositional knowledge units (Guo et al., 2025; Li et al., 2025), thereby failing to meet the critical demands of scientific applications for diversity, intersectionality, and verifiability.

5 Conclusion

In this paper, we introduced SCICUSTOM, an ontology-driven framework that automates the construction of customized benchmarks for scientific LLM evaluation. By reorganizing scattered scientific corpora into fine-grained, reusable knowledge units, our approach overcomes the scalability bottlenecks of manual curation while ensuring evaluations remain grounded in verifiable facts. Extensive experiments demonstrate that SCICUSTOM not only achieves high alignment with expert-curated benchmarks in chemistry and healthcare but also reliably assesses model capabilities in novel, unbenchmarked scientific frontiers. This work establishes a scalable, application-aware evaluation framework for navigating the complex and evolving landscape of scientific research.

Limitations

SCICUSTOM has two primary limitations. First, the scope of our current evaluation is constrained by the underlying ontologies, which predominantly cover biomedical and chemical domains (derived from OBO, BioPortal, and OLS). Consequently, some scientific disciplines such as mathematics and theoretical physics, are not yet included. Future work can address this by integrating broader scientific taxonomies to expand the knowledge space. Second, SCICUSTOM-curated benchmarks depend on the coverage of the source scientific corpus \mathcal{D} . Some knowledge units may currently suffer from data sparsity. With the tagging model, we can identify these low-resource knowledge units by monitoring frequencies of knowledge units. In future work, this framework will continuously evolve and scale as new scientific datasets become available.

Ethics Statement

Our experimental evaluation includes a benchmark subset focused on healthcare. We emphasize that our framework and the constructed benchmarks are strictly designed for educational and evaluative purposes within a textual Question-Answering (QA) context. The benchmarks assess an LLM’s ability to recall and reason about established scientific facts found in public academic literature. They do not involve, facilitate, or encourage any actionable bio-security threats. Furthermore, all scientific data used in SCICUSTOM are aggregated from open-source, publicly available repositories. No private, classified, or patient-identifiable data

were used in the construction of our benchmarks. We believe that rigorous evaluation of LLMs in these domains is a prerequisite for their safe and responsible deployment in scientific research.

Acknowledgments

The authors Ming Zhang, Yiyang Gu, and Junwei Yang are supported by grants from the National Natural Science Foundation of China (NSFC Grant Number 62276002). The authors Zequn Liu, Yingce Xia, and Shufang Xie are supported by Zhongguancun Academy (Grant No. C20250513 and Grant No. P190260302).

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Khizar Anjum, Muhammad Arbab Arshad, Kadhim Hayawi, Efstathios Polyzos, Asadullah Tariq, Mohamed Adel Serhani, Laiba Batool, Brady Lund, Nishith Reddy Mannuru, Ravi Varma Kumar Bevara, and 1 others. 2025. Domain specific benchmarks for evaluating multimodal large language models. *arXiv preprint arXiv:2506.12958*.
- Elron Bandel, Yotam Perlit, Elad Venezian, Roni Friedman, Ofir Arviv, Matan Orbach, Shachar Don-Yehiya, Dafna Sheinwald, Ariel Gera, Leshem Choshen, and 1 others. 2024. Unitxt: Flexible, shareable and reusable data preparation and evaluation for generative ai. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 3: System Demonstrations)*, pages 207–215.
- Suhana Bedi, Yutong Liu, Lucy Orr-Ewing, Dev Dash, Sanmi Koyejo, Alison Callahan, Jason A Fries, Michael Wornow, Akshay Swaminathan, Lisa Soleymani Lehmann, and 1 others. 2025. Testing and evaluation of health care applications of large language models: a systematic review. *Jama*, 333(4):319–328.
- Twist Bioscience. 2025. Ai’s future hinges on the wet lab. *Twist Bioscience Blog*.
- Abeba Birhane, Atoosa Kasirzadeh, David Leslie, and Sandra Wachter. 2023. Science in the age of large language models. *Nature Reviews Physics*, 5(5):277–280.
- Daniil A Boiko, Robert MacKnight, Ben Kline, and Gabe Gomes. 2023. Autonomous chemical research with large language models. *Nature*, 624(7992):570–578.

- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, and 1 others. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Rishi Bommasani, Percy Liang, and Tony Lee. 2023. [Holistic evaluation of language models](#). *Annals of the New York Academy of Sciences*, 1525(1):140–146.
- Hengxing Cai, Xiaochen Cai, Junhan Chang, Sihang Li, Lin Yao, and 1 others. 2025. Sciassess: Benchmarking llm proficiency in scientific literature analysis. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 2335–2357.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, and 1 others. 2024. A survey on evaluation of large language models. *ACM transactions on intelligent systems and technology*, 15(3):1–45.
- Kai-Hendrik Cohrs, Emiliano Diaz, Vasileios Sitokontantinou, Gherardo Varando, and Gustau Camps-Valls. 2025. Large language models for causal hypothesis generation in science. *Machine Learning: Science and Technology*, 6(1):013001.
- Yin Fang, Xiaozhuan Liang, Ningyu Zhang, Kangwei Liu, Rui Huang, Zhuo Chen, Xiaohui Fan, and Hua-jun Chen. 2024. [Mol-instructions: A large-scale biomolecular instruction dataset for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Eitan Farchi, Shmulik Froimovich, Rami Katan, and Orna Raz. 2024. Automatic generation of benchmarks and reliable llm judgment for code tasks. *arXiv preprint arXiv:2410.21071*.
- Juraj Gottweis, Wei-Hung Weng, Alexander Daryin, Tao Tu, Anil Palepu, Petar Sirkovic, Artiom Myaskovsky, Felix Weissenberger, Keran Rong, Ryutaro Tanno, and 1 others. 2025. Towards an ai co-scientist. *arXiv preprint arXiv:2502.18864*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Cheng Guo, Hu Kai, Shuxian Liang, Yiyang Jiang, Yi Gao, Xian-Sheng Hua, and Wei Dong. 2025. Sd-bench: A survey-based domain-specific llm benchmarking and optimization framework. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13492–13506.
- 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*.
- Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf H Roohani, Jure Leskovec, Connor W. Coley, Cao Xiao, Jimeng Sun, and Marinka Zitnik. 2021. [Therapeutics data commons: Machine learning datasets and tasks for drug discovery and development](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Applied Sciences*, 11(14):6421.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*, pages 611–626.
- Xiang Lisa Li, Farzaan Kaiyom, Evan Zheran Liu, Yifan Mai, Percy Liang, and Tatsunori Hashimoto. 2025. [Autobench: Towards declarative benchmark construction](#). In *The Thirteenth International Conference on Learning Representations*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Zequan Liu, Shukai Wang, Yiyang Gu, Ruiyi Zhang, Ming Zhang, and Sheng Wang. 2021. Graphine: A dataset for graph-aware terminology definition generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3453–3463.
- Zequan Liu, Wei Zhang, Yingce Xia, Lijun Wu, Shufang Xie, Tao Qin, Ming Zhang, and Tie-Yan Liu. 2023. [MolXPT: Wrapping molecules with text for generative pre-training](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1606–1616, Toronto, Canada. Association for Computational Linguistics.
- James Lu, Keunwoo Choi, Maksim Ereemeev, Jogarao Gobburu, Srijib Goswami, Qi Liu, Gary Mo, Cynthia J Musante, and Mohamed H Shahin. 2025. Large language models and their applications in drug discovery and development: A primer. *Clinical and Translational Science*, 18(4):e70205.
- Junyu Luo, Weizhi Zhang, Ye Yuan, Yusheng Zhao, Junwei Yang, Yiyang Gu, Bohan Wu, Binqi Chen, Ziyue Qiao, Qingqing Long, and 1 others. 2025. Large language model agent: A survey on methodology, applications and challenges. *arXiv preprint arXiv:2503.21460*.

- Andres M. Bran, Sam Cox, Oliver Schilter, Carlo Baldassari, Andrew D White, and Philippe Schwaller. 2024. Augmenting large language models with chemistry tools. *Nature machine intelligence*, 6(5):525–535.
- Santiago Miret and NM Anoop Krishnan. 2025. Enabling large language models for real-world materials discovery. *Nature Machine Intelligence*, 7(7):991–998.
- Adrian Mirza, Nawaf Alampara, Sreekanth Kunchapu, Martiño Ríos-García, Benedict Emoekabu, Aswanth Krishnan, Tanya Gupta, Mara Schilling-Wilhelmi, Macjonathan Okereke, Anagha Aneesh, and 1 others. 2024. Are large language models superhuman chemists? *arXiv preprint arXiv:2404.01475*.
- Jasmine Chiat Ling Ong, Liyuan Jin, Kabilan Elangovan, Gilbert Yong San Lim, Daniel Yan Zheng Lim, Gerald Gui Ren Sng, Yu He Ke, Joshua Yi Min Tung, Ryan Jian Zhong, Christopher Ming Yao Koh, and 1 others. 2025. Large language model as clinical decision support system augments medication safety in 16 clinical specialties. *Cell Reports Medicine*, 6(10).
- Wendy S Parker. 2020. Evaluating data journeys: Climategate, synthetic data and the benchmarking of methods for climate data processing. In *Data journeys in the sciences*, pages 191–206. Springer.
- Jyotishman Pathak, Thomas M Johnson, and Christopher G Chute. 2008. Modular ontology techniques and their applications in the biomedical domain. In *2008 IEEE International Conference on Information Reuse and Integration*, pages 351–356. IEEE.
- José R Penadés, Juraj Gottweis, Lingchen He, Jonasz B Patkowski, Alexander Shurick, Wei-Hung Weng, Tao Tu, Anil Palepu, Artiom Myaskovsky, Annalisa Pawlosky, and 1 others. 2025. Ai mirrors experimental science to uncover a novel mechanism of gene transfer crucial to bacterial evolution. *bioRxiv*, pages 2025–02.
- José Pombal, Nuno M Guerreiro, Ricardo Rei, and André FT Martins. 2025. Zero-shot benchmarking: A framework for flexible and scalable automatic evaluation of language models. *arXiv preprint arXiv:2504.01001*.
- Valentina Pyatkin, Saumya Malik, Victoria Graf, Hamish Ivison, Shengyi Huang, Pradeep Dasigi, Nathan Lambert, and Hannaneh Hajishirzi. 2025. [Generalizing verifiable instruction following](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2024. Gpqa: A graduate-level google-proof q&a benchmark. In *First conference on language modeling*.
- Gayathri Saranathan, Cong Xu, Mahammad Parwez Alam, Tarun Kumar, Martin Foltin, Soon Yee Wong, and Suparna Bhattacharya. 2025. Sublime: Subset selection via rank correlation prediction for data-efficient llm evaluation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 30572–30593.
- Nima Shoghi, Adeesh Kolluru, John R Kitchin, Zachary W Ulissi, C Lawrence Zitnick, and Brandon M Wood. 2023. From molecules to materials: Pre-training large generalizable models for atomic property prediction. *arXiv preprint arXiv:2310.16802*.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, and 1 others. 2023. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180.
- Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Mohamed Amin, Le Hou, Kevin Clark, Stephen R Pfohl, Heather Cole-Lewis, and 1 others. 2025. Toward expert-level medical question answering with large language models. *Nature medicine*, 31(3):943–950.
- Yu Song, Santiago Miret, and Bang Liu. 2023. Matscinnlp: Evaluating scientific language models on materials science language tasks using text-to-schema modeling. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3621–3639.
- Liangtai Sun, Yang Han, Zihan Zhao, Da Ma, Zhennan Shen, Baocai Chen, Lu Chen, and Kai Yu. 2024. Scieval: A multi-level large language model evaluation benchmark for scientific research. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19053–19061.
- David Wadden, Kejian Shi, Jacob Morrison, Alan Li, Aakanksha Naik, Shruti Singh, Nitzan Barzilay, Kyle Lo, Tom Hope, Luca Soldaini, and 1 others. 2025. Sciriff: A resource to enhance language model instruction-following over scientific literature. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 6072–6109.
- Botian Wang, Yawen Ouyang, Yaohui Li, Mianzhi Pan, Yuanhang Tang, Yiqun Wang, Haorui Cui, Jianbing Zhang, Xiaonan Wang, Wei-Ying Ma, and Hao Zhou. 2025. Moma: A modular deep learning framework for material property prediction. *arXiv preprint arXiv:2502.15483*.
- Xiaoxuan Wang, Ziniu Hu, Pan Lu, Yanqiao Zhu, Jieyu Zhang, Satyen Subramaniam, Arjun R Loomba, Shichang Zhang, Yizhou Sun, and Wei Wang. 2024a. Scibench: Evaluating college-level scientific problem-solving abilities of large language models. In *International Conference on Machine Learning*, pages 50622–50649. PMLR.

- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, and 1 others. 2024b. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37:95266–95290.
- Jason Wei, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese, John Schulman, and William Fedus. 2024. Measuring short-form factuality in large language models. *arXiv preprint arXiv:2411.04368*.
- R. B. Woodward and Roald Hoffmann. 1970. *The Conservation of Orbital Symmetry*. Verlag Chemie, Weinheim. First Edition, First Printing.
- Zhenqin Wu, Bharath Ramsundar, Evan N. Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S. Pappu, Karl Leswing, and Vijay Pande. 2018. [Moleculenet: A benchmark for molecular machine learning](#). *Chemical Science*, 9(2):513–530.
- Yingce Xia, Peiran Jin, Shufang Xie, Liang He, Chuan Cao, Renqian Luo, Guoqing Liu, Yue Wang, Zekun Liu, Yuan-Jyue Chen, Zekun Guo, Yeqi Bai, Pan Deng, Yaosen Min, Ziheng Lu, Hongxia Hao, Han Yang, Jielan Li, Chang Liu, and 27 others. 2025. [Nature language model: Deciphering the language of nature for scientific discovery](#). *Preprint*, arXiv:2502.07527.
- Ye Yuan, Kexin Tang, Jianhao Shen, Ming Zhang, and Chenguang Wang. 2024. Measuring social norms of large language models. *arXiv preprint arXiv:2404.02491*.
- Dan Zhang, Ziniu Hu, Sining Zhoubian, Zhengxiao Du, Kaiyu Yang, Zihan Wang, Yisong Yue, Yuxiao Dong, and Jie Tang. 2024a. [Sciinstruct: a self-reflective instruction annotated dataset for training scientific language models](#). *Advances in Neural Information Processing Systems*, 37:1443–1473.
- Jieyu Zhang, Weikai Huang, Zixian Ma, Oscar Michel, Dong He, Tanmay Gupta, Wei-Chiu Ma, Ali Farhadi, Aniruddha Kembhavi, and Ranjay Krishna. 2024b. [Task me anything](#). *arXiv preprint arXiv:2406.11775*.

A Details of Proxy Subset Selection

For each data sample $\langle q, a \rangle \in \mathcal{D}_r$, we compute two commonly used intrinsic metrics to characterize its difficulty and quality: (i) **Hardness Score (H)**: We utilize open-weights LLMs to calculate the perplexity (PPL) of the answer a_d conditioned on the question q_d :

$$H(d) = \exp \left(-\frac{1}{|a_d|} \sum_{i=1}^{|a_d|} \log P(t_i | q_d, t_{<i}) \right) \quad (1)$$

where t_i represents the i -th token of the answer a_d . (ii) **Quality Score (Q)**: We assess the linguistic quality using the Flesch Readability Score.

Our objective is to select \mathcal{P}_q such that the distributions of these scores closely approximate those of \mathcal{D}_q . We formulate this as minimizing the cumulative Wasserstein distance (\mathcal{W}) between the distributions of the subset and the full set:

$$\min_{\mathcal{P}_q} \left(\mathcal{W}(P_H^{\mathcal{D}_q}, P_H^{\mathcal{P}_q}) + \mathcal{W}(P_Q^{\mathcal{D}_q}, P_Q^{\mathcal{P}_q}) \right) \quad (2)$$

To solve the sampling of \mathcal{P}_q efficiently, we adopt a cluster-based strategy following SubLIME (Saranathan et al., 2025). Intuitively, we encode all pairs $d = \langle q, a \rangle$ into an embedding space using BERT (Liu et al., 2019). We then perform clustering in embedding space and sample representatives from each cluster.

B Details of Ground-Truth Benchmarks

Table 3 presents the detailed statistics of the expert-curated ground-truth benchmarks used in our evaluation.

Task	Size (# Samples)
Chemistry Domain	
Analytical Chemistry	152
Inorganic Chemistry	92
Material Science	84
Organic Chemistry	429
Physical Chemistry	165
Technical Chemistry	40
Healthcare Domain	
Virology	46
Human Aging	86
Medical Genetics	54
Anatomy	79
Nutrition	179

Table 3: Statistics of the ground-truth benchmarks used in our experiments.

C Implementation Details

We finetuned LLaMa-3-8B (Grattafiori et al., 2024) to train a tagging model. We sampled 50,000 synthetic scientific queries and 30,000 real-world scientific queries as the training set. The LLM used for construct the dataset is GPT-4o (Achiam et al., 2023). The model was finetuned for 2 epochs with a learning rate of $2e - 5$. We conducted experiments on 8 NVIDIA A100 GPUs in a standard Linux environment, using the vLLM (Kwon et al., 2023) framework for efficient large-scale model inference.

We constructed a comprehensive scientific data corpus by aggregating diverse high-quality instruction-tuning datasets and benchmarks, including *SciRIF* (Wadden et al., 2025), *SciInstruct* (Zhang et al., 2024a), *Mol-Instruct* (Fang et al., 2024), *MultiMedQA* (Singhal et al., 2023), *SciEval* (Sun et al., 2024), *MMLU-Pro* (Wang et al., 2024b), *GPQA* (Rein et al., 2024), *IfBench* (Pyatkin et al., 2025), and *SimpleQA* (Wei et al., 2024). This collection yielded a total of 2,000,367 data instances. To avoid data leakage, we filtered out any instances that overlap with the expert-curated ground-truth benchmarks used in our experiments.

To facilitate multi-model voting as discussed in both **Section 2.4** and **Section 2.5**, we utilized three advanced models: gpt-5, claude-opus-4.5, and gemini-3-pro-preview. We set the voting parameter $K_1 = 10$. For the proxy subset selection, we set $K_2 = 100$, ensuring the scale remains consistent with the expert-curated ground-truth benchmarks (**Section B**). Furthermore, following the SubLIME (Saranathan et al., 2025) framework, we employed a cluster-based strategy to sample 100 candidate sets and selected the optimal subset based on the criteria defined in **Equation 2**.

We evaluated 10 LLMs, including Gemini-2.5-flash, GPT-4o, Gpt-5-chat, Claude-opus-4.5, Qwen3-235b-a22b-instruct-2507, Qwen3-max, Grok-4-fast-non-reasoning, Kimi-k2-preview, DeepSeek-V3.2 and Mistral-large on the all comparison benchmarks.

Discussion on Non-Monotonicity and Selection Strategy

We acknowledge that the relevance of data samples within the sorted list $L_{\mathcal{D}'_r}$ does not strictly follow a monotonic decreasing trend, which theoretically challenges the standard binary search assumption. A seemingly more rigorous alternative would be a "Greedy Top- K Selection", where we sequentially scan $L_{\mathcal{D}'_r}$ from the beginning and

Table 4: Comparison between a greedy retrieval strategy and SCICUSTOM on healthcare benchmarks. Scores are Spearman correlations with the corresponding expert-curated ground-truth rankings.

Method	Virology	Human aging	Medical genetics	Anatomy	Nutrition
Greedy Search	0.21	0.34	0.16	0.24	0.31
SCICUSTOM	0.55	0.49	0.42	0.62	0.78

Table 5: Ranking consistency analysis of 10 LLMs on 6 chemistry-specific benchmarks. The table compares SCICUSTOM against baseline benchmarks in preserving model rankings, quantified by Spearman correlation coefficients. Higher scores indicate better consistency. Best results are in **bold**.

Method	Analytical chemistry	Inorganic chemistry	Material science	Organic chemistry	Physical chemistry	Technical chemistry
GPT-5	-0.11	0.05	-0.04	0.38	0.24	-0.07
GPT-5-RAG	0.01	-0.02	-0.07	0.41	0.35	-0.02
Embedding	-0.34	-0.59	-0.39	0.11	-0.73	-0.41
Embedding-filtering	0.23	0.04	0.15	0.37	0.04	0.11
SCICUSTOM	0.86	0.67	0.42	0.89	0.74	0.86

Table 6: Ranking consistency analysis of 10 LLMs on 6 healthcare-specific benchmarks. The table compares SCICUSTOM against baseline benchmarks in preserving model rankings, quantified by Spearman correlation coefficients. Higher scores indicate better consistency. Best results are in **bold**.

Method	Virology	Human aging	Medical genetics	Anatomy	Nutrition
GPT-5	0.25	0.20	0.09	0.11	0.52
GPT-5-RAG	0.31	0.26	0.02	0.09	0.48
Embedding	0.18	0.21	-0.21	-0.32	0.27
Embedding-filtering	0.24	0.34	0.13	0.26	0.51
SCICUSTOM	0.55	0.49	0.42	0.62	0.78

select the first K_2 samples judged as relevant by the voting ensemble. However, our experiments show that this greedy strategy ($K_2 = 100$) results in a significantly lower correlation with ground-truth benchmarks compared to the proposed binary search method (Table 4). We hypothesize that the top-ranked instances in $L_{\mathcal{D}_i}$, which share the highest semantic and keyword overlap with the requirement, often correspond to "canonical" textbook knowledge or highly frequent scientific facts. These instances are likely exposed during the pre-training of most LLMs (potential data leakage), leading to a performance ceiling effect where most models achieve near-perfect accuracy. Consequently, these "highly relevant" samples lack the discriminative power to differentiate model capabilities. In contrast, our binary search strategy determines a cutoff based on a broader structural trend. This approach inadvertently bypasses the saturation of the top-most samples and captures a more challenging and diverse distribution of data, thereby offering higher discriminative validity and better alignment with expert rankings.

D More Experiment Results

To further isolate the contribution of our ontology-driven architecture, we report Spearman correla-

tion results for two alternative baselines. *GPT-5-RAG* augments the synthetic *GPT-5* generator with open-book retrieval from the full corpus, while *Embedding-filtering* applies the same binary-search cutoff and proxy subset selection used in SCICUSTOM on top of the *Embedding* candidate pool. As shown in Table 5 and Table 6, adding retrieval to *GPT-5* yields only limited gains over the fully synthetic baseline, while *Embedding-filtering* improves over raw *Embedding* but still remains consistently below SCICUSTOM across both chemistry and healthcare tasks. These results show that neither retrieval augmentation alone nor post-retrieval filtering alone is sufficient to match the performance of the full ontology-grounded pipeline. This confirms that the fine-grained structural guidance provided by our ontology-grounded knowledge units is indispensable for accurately capturing complex scientific requirements.

E Details of Human Evaluation

E.1 Human Annotation Guidelines

Objective: You will be presented with 50 multiple-choice questions generated for a specific scientific requirement (e.g., *Technical Chemistry*). For each question, please provide two binary labels (0 or 1) based on the criteria below.

1. Relevance Label (relevant)

Determine whether the question requires specific knowledge of the target requirement.

- **1 (Relevant):** The question is strictly aligned with the requirement.
- **0 (Irrelevant):** The question is off-topic, generic (can be answered by a layperson), or belongs to a distinctly different scientific field.

2. Correctness Label (correct)

Evaluate the scientific accuracy of the MCQs.

- **1 (Correct):** The questions and options are scientifically accurate.
- **0 (Incorrect):** The question is wrong, or the selected option is factually wrong, scientifically flawed, or there is a significantly better/more accurate option available in the choices.

E.2 Human Annotators

The annotation was performed by three Master’s students specializing in AI for Science (AI4Chemistry). Participants were recruited from an university department and were compensated at an hourly rate exceeding the local minimum wage. For the case study, we engaged a Ph.D. researcher specializing in AI and Chemistry to assess the generated benchmark. The expert was tasked with categorizing the knowledge units and keywords in the MCQs into three distinct levels of granularity and relevance: *Strong Relevance*, *Relevant*, and *Weak Relevance*. The expert was compensated at a professional hourly rate commensurate. Prior to the task, all annotators were informed of the data usage scope and provided written informed consent. The study protocol was reviewed and determined exempt by an ethics review board.

F Automatic Evaluation of Distractor Quality

To validate the quality of LLM-generated distractors, we evaluate 500 generated MCQs with Claude-Opus-4.6 as an independent judge. Each MCQ is scored on a 1–10 scale along two dimensions: **Heuristic Transparency Score (HTS)**, which measures whether distractors can be eliminated by superficial linguistic cues, and **Scientific Plausibility Score (SPS)**, which measures how scientifically relevant and challenging the distractors are. The generated MCQs obtain an average HTS of 8.7 and

SPS of 7.9, indicating low heuristic leakage and strong near-miss quality. We additionally construct a context-free baseline where the judge only sees the options and a simplified question, without the supporting scientific context; accuracy drops to 22%, close to random guess. These results support that our MCQ transformation pipeline produces distractors that are both scientifically plausible and difficult to solve through shallow heuristics alone.

G Prompt Template Inventory

Granularity classification of scientific terms

System: You are a helpful AI assistant.

User: Determine whether the given term is a suitable scientific subfield and answer with one of the following categories:

(moderate): The term is appropriately specific for a scientific subfield. It refers to a category that is neither too general nor too specific for its scientific context. Its scale is similar to the scale of chapter names in subject textbooks.

(too coarse): The term is overly broad or vague for a scientific subfield. It encompasses a wide range of concepts that could be divided into smaller, more specific subfields.

(too fine): The term is overly specific and pertains to a very narrow aspect of a scientific subfield. It may be too detailed to serve as a broader category within the discipline.

Answer at the beginning, explain later. The examples are as follows:

Example 1:

Input: term: anatomical entity

Output: (moderate); Explanation: Anatomical entity refers to a category that is sufficiently specific for many biological subfields but not too narrow.

Example 2:

Input: term: nuclear structure

Output: (moderate); Explanation: Nuclear structure is a well-defined category in molecular biology, specific but not too narrow.

Example 3:

Input: term: electronic file status

Output: (too fine); Explanation: The term refers to a very specific technical concept that is too detailed to be considered a scientific subfield.

Example 4:

Input: term: b-lymphocyte

Output: (too fine); Explanation: While important, the term refers to a very specific type of cell, not a broad enough category to encompass a subfield.

Example 5:

Input: term: continuant

Output: (too coarse); Explanation: The term is too general and could refer to a wide variety of objects

or concepts, making it too broad for a specific subfield.

Example 6:

Input: term: *occurrent*

Output: (*too coarse*); Explanation: *Occurrent is overly vague and applies to many concepts, making it too broad for a scientific subfield.*

Input: term: {term}.

Query generation

System: You are a helpful AI assistant.

User:

Low Complexity: You are a {persona}. Generate a user query containing the following keywords: {keywords}. Do not introduce other scientific entities or topics. Only return the query.

High Complexity: You are a {persona}. Generate a user query containing the following keywords: {keywords}. Do not introduce other scientific entities or topics. Make the query long and complex. Only return the query.

Scientific Personas:

- Astrophysicist
- Marine Biologist
- AI Researcher
- Molecular Geneticist
- Quantum Physicist
- Environmental Chemist
- Neuroscientist
- Ecologist
- Bioinformatician
- Pharmacologist
- Geologist
- Biomedical Engineer
- Mathematical Modeler
- Virologist
- Behavioral Psychologist
- Data Scientist
- Theoretical Chemist
- Climate Scientist
- Structural Biologist
- Robotics Engineer

Benchmark requirements and descriptions

Chemistry benchmarks:

- *Analytical chemistry:* Generate questions that test knowledge and reasoning in analytical chemistry. The questions should assess understanding of how experimental analytical signals (e.g., NMR, IR, UV-Vis, mass spectra, chromatographic behavior, titration curves) relate to molecular structure, composition, concentration, or purity. Focus on conceptual interpretation and chemical reasoning rather than numerical data processing or instrument-specific operating procedures.
- *Inorganic chemistry:* Generate questions that test core knowledge and reasoning in inorganic chemistry. The questions should focus on electronic structure, oxidation states, coordination geometry,

ligand field effects, symmetry, and periodic trends in inorganic systems. Emphasize conceptual understanding of structure–property relationships rather than memorization of isolated facts.

- *Material science:* Generate questions that evaluate understanding in materials science. The questions should assess how atomic or microstructural features (e.g., crystal structure, defects, phases, interfaces) determine macroscopic properties such as mechanical strength, electrical conductivity, or thermal behavior. Focus on structure–property reasoning rather than detailed synthesis protocols.
- *Organic chemistry:* The organic chemistry benchmark assesses a wide range of skills on reasoning about chemical structures and reaction pathways, such as Reaction Mechanism Identification, Product Prediction, NMR Signal Prediction, Number of Isomers, Polymer Chemistry, Nomenclature Conversion and Organic Reactivity.
- *Physical chemistry:* Generate questions that test conceptual understanding in physical chemistry. The questions should assess reasoning about thermodynamics, kinetics, equilibrium, and molecular-level physical principles. Emphasize qualitative reasoning about trends and relationships rather than explicit numerical calculation.
- *Technical chemistry:* Generate questions that assess knowledge in technical and industrial chemistry. The questions should focus on chemical processes at scale, such as reactor behavior, process optimization, safety considerations, and material or energy efficiency. Emphasize reasoning about system-level behavior rather than detailed engineering design.

Healthcare benchmarks:

- *Virology:* Generate questions that test conceptual understanding in virology. The questions should assess knowledge of viral structure, replication cycles, genome organization, and interactions with host cells and immune systems. Avoid clinical treatment guidelines or laboratory diagnostic protocols.
- *Human aging:* Generate questions that probe understanding of biological mechanisms of human aging. The questions should focus on molecular, cellular, and systemic processes associated with aging, such as genomic stability, cellular senescence, metabolic regulation, and tissue-level decline. Emphasize mechanistic reasoning rather than epidemiological statistics.
- *Medical genetics:* Generate questions that test reasoning in medical genetics. The questions should assess understanding of inheritance patterns, genotype–phenotype relationships, penetrance, and genetic variation. Focus on conceptual genetic reasoning rather than clinical decision-making.
- *Anatomy:* Generate questions that evaluate knowledge of human anatomy. The questions should focus on the identification, spatial relationships, and functional roles of anatomical structures. Avoid surgical procedures or pathological conditions.

- *Nutrition*: Generate questions that assess understanding of nutritional science. The questions should focus on the biological roles of macro- and micronutrients, their involvement in metabolism, and the physiological consequences of deficiency or imbalance. Emphasize mechanistic understanding over dietary recommendations.

Voting-based relevant tag selection

System: You are an expert in **{domain}**. Your task is to map a benchmark description to the most relevant technical tags

User: Task:

Given a **{domain}** benchmark description, identify and rank the most relevant tags from a candidate list.

Benchmark Description: **{description}**

Candidate Tags (sorted by frequency, lowest frequency first; lower frequency usually indicates higher specificity): **{Tag list}**.

Ranking Principles: Rank tags from highest to lowest relevance to the benchmark, following these rules:

Relevance First:

A tag is relevant if it directly reflects the core concepts, tasks, data modalities, or evaluation focus of the benchmark. Irrelevant or weakly related tags should not be selected.

Specificity as a Tie-breaker:

If multiple tags are similarly relevant, rank the more specific and narrowly scoped tag higher. Prefer concrete technical terms (e.g., “Histone Acetylation Prediction”) over broader categories (e.g., “Epigenetics”).

Avoid Overly Generic Tags:

High-level or generic tags (e.g., “biological process”, “chemical entity”) should only be selected if no more specific alternative applies.

Frequency Awareness:

When relevance and specificity are comparable, prefer lower-frequency tags, as they tend to be more precise.

Output Requirements:

Return a single list of tags, sorted from most to least relevant. For efficiency, return only the top 100 tags (or fewer if fewer are relevant). Do not include explanations, scores, or extra text—output the ranked list only.

Benchmark generation

System: You are an expert in **{domain}** and tasked with constructing a high-quality benchmark to assess the domain-specific knowledge abilities of large language models. Please return the benchmark in a JSON format.

User: Your task is to generate exactly **{K}**

single-choice questions in the domain of **{domain}**. Detailed description of this domain: **{description}**
The questions should:

1. Focus on core concepts, expert-level knowledge, and non-trivial reasoning in this domain.
2. Avoid trivial definitions, purely factual memorization, or overly ambiguous questions.
3. Include a mix of:
 - Conceptual understanding
 - Mechanism or principle-based reasoning
 - Application or scenario-based reasoning
4. Be answerable without external tools, but not solvable by surface-level pattern matching.

Question format:

1. Each question must have 4–5 options.
2. Options should be concise and mutually exclusive.
3. Each question have only one correct answers.

Output format (STRICT):

Return only a JSON array of length **{K}**.

Each element must have the following structure:

```
{
  "query": "<question text with options labeled A, B, C, D (and E if applicable)>",
  "answer": "<correct option label>"
}
```

MCQ transformation

System: You are an expert in **{domain}** and tasked with curating a rigorous benchmark to evaluate the capabilities of Large Language Models. Please return the processed entry in a JSON format.

User: Your task is to convert the following raw problem content into a standardized single-choice question suitable for LLM evaluation.

Raw problem: **{input_content}**

Conversion Guidelines:

1. Format Adaptation:
 - If the input is already a multiple-choice question: Preserve the original stem and options exactly. Ensure the formatting aligns with the output requirements.
 - If the input is not a multiple-choice question: Convert it into a single-choice question by generating 3–4 incorrect options (distractors).
2. Distractor Engineering:
 - Avoid trivial errors, logical fallacies that are easily filtered, or clearly unrelated concepts.
3. Fidelity & Difficulty:
 - Strict adherence to the factual truth and reasoning logic of the original content is required.
 - Do not simplify the problem complexity. The resulting MCQ must maintain the same discriminative power as the original input.
4. Exclusivity: Ensure there is exactly one indisputably correct option.

Question format:

1. The final output must contain 4–5 options (A, B, C, D, [E]).
2. Options should be concise and mutually exclusive.

Output format (STRICT):

Return only a single JSON object.
The object must have the following structure:

```
{  
  "query": "<question stem followed by options  
labeled A, B, C, D (and E if applicable), separated  
by newlines>",  
  "answer": "<correct option label, e.g., 'A'>"  
}
```

H LLMs Usage

We adhere to the ACL Code of Ethics. We use large language models solely for polishing writing. All scientific contributions remain entirely our own.