



Perovskite-LLM: Knowledge-Enhanced Large Language Models for Perovskite Solar Cell Research

Xiang Liu^{1,*} Penglei Sun^{1,*} Shuyan Chen^{1,2} Longhan Zhang^{1,2}
Peijie Dong¹ Huajie You¹ Yongqi Zhang¹ Chang Yan^{1,2}
Xiaowen Chu^{1,2,†} Tong-yi Zhang^{1,2,†}

¹The Hong Kong University of Science and Technology (Guangzhou)

²Guangzhou Municipal Key Laboratory of Materials Informatics

Abstract

The rapid advancement of perovskite solar cells (PSCs) has led to an exponential growth in research publications, creating an urgent need for efficient knowledge management and reasoning systems in this domain. We present a comprehensive knowledge-enhanced system for PSCs that integrates three key components. First, we develop **Perovskite-KG**, a domain-specific knowledge graph constructed from 1,517 research papers, containing 23,789 entities and 22,272 relationships. Second, we create two complementary datasets: Perovskite-Chat, comprising 55,101 high-quality question-answer pairs generated through a novel multi-agent framework, and Perovskite-Reasoning, containing 2,217 carefully curated materials science problems. Third, we introduce two specialized large language models: **Perovskite-Chat-LLM** is used for domain-specific knowledge assistance, and **Perovskite-Reasoning-LLM** is used for scientific reasoning tasks. Experimental results demonstrate that our system significantly outperforms existing models in domain-specific knowledge retrieval and scientific reasoning tasks, providing researchers with effective tools for literature review, experimental design, and complex problem solving in PSC research.

1 Introduction

Perovskite solar cells (PSCs) have emerged as one of the most promising next-generation photovoltaic technologies, achieving remarkable progress with power conversion efficiencies (PCEs) exceeding 27.0% within just over a decade (National Renewable Energy Laboratory, 2025; Snaith, 2018; Correa-Baena et al., 2017; Wu et al., 2021; Ang et al., 2022; Sathaye et al., 2011; Bogdanov et al., 2019). The rapid development of PSCs has generated an exponential growth in research publications,

making it increasingly challenging for researchers to efficiently access and utilize the vast amount of knowledge in this field. This challenge is particularly acute given the complex interplay between material composition, fabrication processes, and device structure that characterizes PSC research.

Traditional approaches to scientific knowledge management, such as literature reviews and databases, while valuable, are limited in their ability to capture the intricate relationships between different aspects of PSC research (Yang et al., 2024b; Han et al., 2025). Furthermore, existing artificial intelligence systems in materials science typically focus on specific prediction tasks or general scientific knowledge, lacking the specialized capability to handle the unique characteristics of the research of perovskite solar cells and more general materials science research (Han et al., 2025; Seifrid et al., 2022; Mishra et al., 2024; Song et al., 2023a; Schilling-Wilhelmi et al., 2025; Miret and Krishnan, 2024). This gap highlights the need for an integrated system that can both systematically organize domain knowledge and provide intelligent assistance to researchers.

To address these challenges, we present a comprehensive knowledge-enhanced system specifically designed for the perovskite solar cell domain, consisting of three key components. First, we develop **Perovskite-KG**, a domain-specific knowledge graph constructed from 1,517 research articles, containing 23,789 entities and 22,272 relationships between manufacturing processes, parameters, and performance metrics. Second, we create a multi-agent framework for generating high-quality instruction-tuning data, which not only reduces annotation costs but also ensures high reliability and low hallucination through the synergy of multiple specialized agents and expert guidance. This framework generates two complementary datasets: (1) **Perovskite-Chat**, an instruction-tuning dataset comprising 55,101 high-quality question-answer

*Equal Contribution.

†Corresponding Author.

pairs generated from 2,214 high-impact papers using a novel multi-agent framework, and (2) **Perovskite-Reasoning**, a collection of 2,217 carefully curated materials science problems designed to enhance scientific reasoning capabilities. Third, we introduce two specialized large language models: **Perovskite-Chat-LLM** for domain-specific knowledge assistance and **Perovskite-Reasoning-LLM** for tackling complex materials science reasoning tasks. Our work makes the following key contributions:

❶ **Construction of a comprehensive perovskite knowledge graph.** We construct the first comprehensive knowledge graph for perovskite solar cells, organizing domain knowledge into a structured format that captures the relationships between materials, processes, and device performance.

❷ **Multi-agent framework for high-quality data generation.** We propose an effective multi-agent framework for generating high-quality instruction-tuning data, resulting in two specialized datasets: a diverse domain-specific dataset covering seven research categories and a focused reasoning dataset for enhancing scientific problem solving capabilities.

❸ **Specialized LLMs for perovskite research.** We develop and evaluate two specialized large language models for perovskite solar cells that demonstrate superior performance compared to baseline models: one optimized for domain-specific queries and another for scientific reasoning tasks.

❹ **Extensive experimental validation.** We provide extensive experimental results showing the effectiveness of our integrated system in supporting various research tasks, from literature review to experimental design and complex problem solving in materials science.

2 Related Work

2.1 LLM in Materials Science

The convergence of language modeling and computational materials science has unlocked transformative potential for rapid discovery. Recent breakthroughs in domain-specific architectures (e.g. hierarchical attention mechanisms (Kononova et al., 2021) and multimodal fusion networks (Swain and Cole, 2016)) have addressed critical challenges in crystal structure prediction (Walker et al., 2021) and phase diagram analysis (Trewartha et al., 2022). As evidenced by the Materials Genome Initiative benchmarks (Tshitoyan et al., 2019), three main

research thrusts have emerged: (1) structured information extraction from heterogeneous corpora, (2) knowledge graph embeddings for composition-property relationships, and (3) neurosymbolic reasoning for synthesis pathway optimization (Seifrid et al., 2022; Mishra et al., 2024; Song et al., 2023a; Schilling-Wilhelmi et al., 2025; Miret and Krishnan, 2024).

Building upon these foundations, knowledge-enhanced systems have achieved state-of-the-art performance through two complementary paradigms: graph-based approaches employing heterogeneous graph neural networks (HGNNs) now attain 89.7% accuracy on multi-hop material property queries (An et al., 2024), while agent-based frameworks demonstrate 18.7% improvement in autonomous experimental design through chain-of-thought prompting (Zhang et al., 2024a; Song et al., 2023b; Gupta et al., 2022; Hira et al., 2024).

The field’s maturation is further evidenced by systematic resource development: (i) The SciQAG framework (Wan et al., 2024) introduces a novel curriculum learning paradigm for generating 120K domain-specific QA pairs, reducing expert annotation requirements by 78%; (ii) Standardized evaluation now spans chemical synthesis (ChemLLM-Bench’s reaction yield prediction task (Guo et al., 2023)), biomedical applications (MultiMedQA’s toxicity prediction challenge (Singhal et al., 2023)), and cross-domain reasoning (SciEval’s materials-device co-design track (Sun et al., 2023)).

2.2 Knowledge Graph in Materials Science

Domain-specific knowledge graphs have evolved into structured semantic frameworks that systematically consolidate heterogeneous multi-source data through machine-readable representations, allowing cross-domain knowledge integration to accelerate discovery pipelines (Pan et al., 2024c; Song et al., 2024; Zhu et al., 2022; Bordes et al., 2013; Galkin et al.). In materials informatics, current implementations manifest two distinct paradigms: literature-derived systems exemplified by MatKG (Venugopal and Olivetti, 2024) and DISCOMAT (Gupta et al., 2023), which employ NLP and graph techniques to extract material compositions from textual sources, while empirical architectures represented by MatSciKB (Zhang et al., 2024b), Propnet (Mrdjenovich et al., 2020), MekG (Statt et al., 2023), and MOF-KG (An et al., 2022) focus on encoding experimental provenance and computational models through graph-based

representations of material lineages. However, these approaches face the challenges that manual curation processes face with resource burdens, while existing extraction methods exhibit limited granularity in resolving complex synthesis-process-property relationships from unstructured text. To address these limitations, we propose an LLM-driven framework specifically optimized for perovskite materials research that features a hybrid architecture that synergizes domain ontologies with self-supervised relationship extraction, augmented by automated quality control pipelines that enforce materials science constraints.

2.3 Multi-agent systems

The landscape of AI system architectures encompasses two distinct paradigms: multi-agent systems and autonomous agents (Zhuge et al., 2023; Hong et al., 2024a; Zhang et al., 2024c; Wang et al., 2023; Su et al., 2025). While autonomous agents rely on independent decision-making capabilities, multi-agent systems excel through structured collaboration between specialized components. The latter approach offers practical advantages by building upon established expertise rather than requiring complex behavioral modeling.

Research in multi-agent frameworks has evolved along two primary trajectories. The first focuses on domain-agnostic systems that leverage collective intelligence for general problem-solving (Wei et al., 2022; Diao et al., 2024; Wang et al., 2022; Madaan et al., 2023; Wang et al., 2024). The second pathway explores domain-specific applications, with notable implementations in: Code generation and debugging (Hong et al., 2024b; Ridnik et al., 2024; Zhong et al., 2024a), Data analytics (Xie et al., 2024; Ye et al., 2024; Li et al., 2024; Zhou et al., 2023b), Mathematical reasoning (Zhong et al., 2024b; Xu et al., 2024; Pan et al., 2024b), Knowledge retrieval (Nori et al., 2023; Zhou et al., 2024)

Despite significant progress in identifying effective agent configurations for specific use cases, the field still faces the challenge of developing systematic approaches for new domains. This highlights the importance of research into automated methods for framework design and optimization.

In this section, we collect 1,517 paper in perovskite domain to build Perovskite-KG and design the automatic knowledge graph construction pipeline including three stages document filtering, knowledge extracting and knowledge graph organi-

zation (Mrdjenovich et al., 2020), as shown in the Appendix A.

2.4 Reasoning alignment

Recent advances in parameter-efficient alignment have witnessed multiple research teams pursue different methodologies to align the performance of o1 (OpenAI, 2024). Contemporary approaches bifurcate along two technical axes: (1) reinforcement learning paradigms exemplified by DeepSeek-R1’s adversarial preference optimization (Guo et al., 2025) and K1.5’s multi-objective reward shaping (Team et al., 2025), versus (2) supervised fine-tuning strategies employing distilled datasets at scale ($\geq 10^4$ examples) as demonstrated in (Team, 2025; Xu et al., 2025; Labs, 2025). In particular, S1 (Muennighoff et al., 2025) and LIMO (Ye et al., 2025) operationalize the Superficial Alignment Hypothesis (Zhou et al., 2023a) through curriculum-based sparse fine-tuning, achieving comparable reasoning capabilities with merely 1,000-2,000 carefully curated examples, a 92% reduction in annotation costs relative to conventional SFT approaches.

3 Perovskite-KG

• **Document Filtering.** Drawing upon expert knowledge, we have developed the schema for perovskite materials. This schema, shown in the Appendix Table 6, integrates three ontologies $\{o_i \mid o_i \in \text{schema}\}$: fabrication, parameters, and performance. The fabrication ontology encompasses the procedures and conditions required to synthesize perovskite materials. The parameters ontology defines the ingredients, structural components, and other compositional aspects of the device. The performance ontology is concerned with the efficiency and functional characteristics of perovskite devices. Each ontology o_i is further divided into sub-ontologies $so_i^{(j)}$, where $o_i = \bigcup_{j=1}^{n_i} so_i^{(j)}$ and n_i represents the number of sub-ontologies within o_i . Each sub-ontology $so_i^{(j)}$ provides a domain-specific description, denoted as $d_i^{(j)}$, along with a corresponding data type, denoted as $t_i^{(j)}$, that is relevant to its particular scope.

For each sub-ontology $[so_i^{(j)}, d_i^{(j)}, t_i^{(j)}]$ (e.g., "Coating Parameter" - "Details about the coating method used in the material deposition process" - "Float"), we create the prompts to query documents $D = \{D_k \mid k = 1, \dots, m\}$ using a large language model. These prompts facilitate the extraction of relevant information for each sub-ontology. The

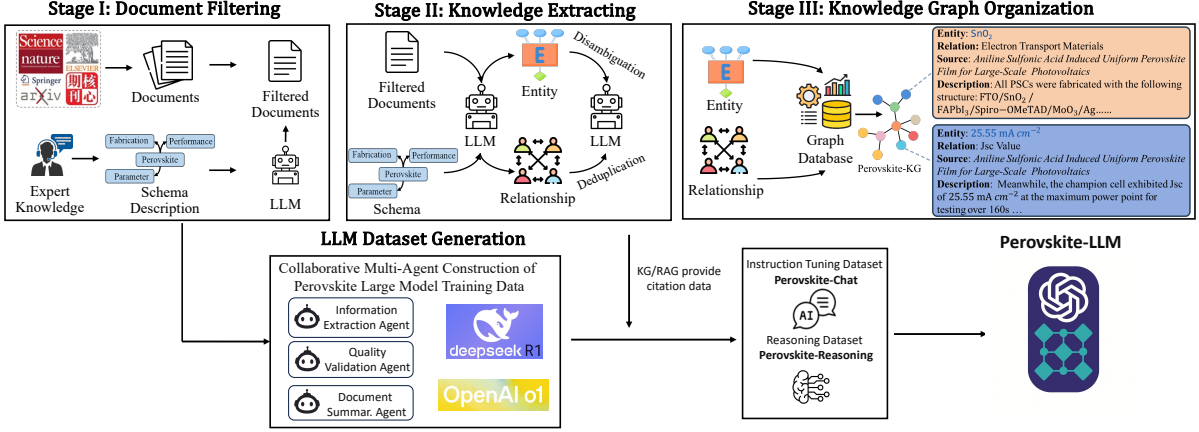


Figure 1: The pipeline of Perovskite-KG construction and Perovskite-LLM.

output $D_{\text{filtered}}^{(i,j)}$ is defined as:

$$D_{\text{filtered}}^{(i,j)} = \{D_k \in D \mid so_i^{(j)} \subset D_k\}, \quad (1)$$

where $D_{\text{filtered}}^{(i,j)}$ represents the set of filtered documents containing pertinent details for sub-ontology $so_i^{(j)}$ across the collection. This approach ensures a systematic and efficient retrieval of targeted information for each sub-ontology.

• **Knowledge Extracting.** We employ a prompt function, denoted as $f_{\text{prompt}}(\cdot)$, to transform the sub-ontology $[so_i^{(j)}, d_i^{(j)}, t_i^{(j)}]$ into a document prompt, represented as $f_{\text{prompt}}(so_i^{(j)}, d_i^{(j)}, t_i^{(j)})$. To extract the potential domain knowledge K , we utilize a pre-trained large language model (LLM), expressed as $\text{LLM}(\cdot; \theta)$, under a zero-shot setting where the parameters θ remain fixed. The whole pipeline can be formulated as below:

$$K = \underset{D_{\text{filtered}}^{(i,j)}}{\text{search}} \text{LLM}(f_{\text{prompt}}(so_i^{(j)}, d_i^{(j)}, t_i^{(j)}); \theta), \quad (2)$$

where the search function $\text{search}(\cdot)$ may involve an argmax operation to identify the highest-scoring output or a sampling approach to generate outputs according to the probability distribution specified by the adopted $\text{LLM}(\cdot; \theta)$.

After extracting knowledge, we conduct quality control procedures to ensure accuracy and reliability. These procedures include entity disambiguation and relationship deduplication. Entity disambiguation in a knowledge graph aims to resolve ambiguity by identifying the unique entity that corresponds to an ambiguous mention, denoted as e_{mention} , within a subgraph. The objective is to determine a distinct entity e^* that accurately represents e_{mention} . Relationship deduplication involves

identifying and merging redundant relations in the knowledge graph. For instance, given two relations $r_i = (e_1, r, e_2)$ and $r_j = (e'_1, r', e'_2)$, if they convey the same semantic meaning—that is, if (e_1, e_2) and (e'_1, e'_2) refer to identical entities and the relations r and r' are equivalent.

• **Knowledge Graph Organization.** We construct the Perovskite Knowledge Graph (Perovskite-KG) using a graph database. The Perovskite-KG consists of 23,789 entities and 22,272 relationships. By incorporating citation relationships between papers, we enable our LLM to provide references for its responses, enhancing credibility and reducing hallucination.

4 Instruction Tuning Dataset Generation

In this section, we collect 2,214 top-level publications in the perovskite domain and design the instruction tuning dataset including question answering and multiple choice questions, containing 55,101 instances around 4.4 million tokens, named **Perovskite-Chat**. Our experiments show that our perovskite instruction tuning dataset can effectively improve the performance of LLMs on perovskite-related tasks.

Figure 1 illustrates this multi-agent framework for the generation of instruction tuning dataset. The process begins with expert guidance and academic literature from various sources (including Science, Nature, Elsevier, Springer, arXiv, and others) as input. Expert guidance is provided by the domain expert focused on 7 research categories, 21 research questions. Table 1 further expands this classification by presenting 21 specific research questions (Q1-Q21) organized within these seven categories; more details can be found in Appendix C.1. These inputs feed into a multi-agent system: (1) an In-

Category	Rationale
Device Structure	Fundamental aspects focusing on high-efficiency (>25% PCE) device architecture and fabrication processes (Q1-Q3)
Perf. Enhancement	Analysis of problem-solving approaches and strategic choices in high-performance devices (Q4-Q5)
Metrics	Key performance indicators (V_{OC} , FF, J_{SC}) and their optimization methods (Q6-Q9)
Stability	Critical stability aspects addressing main degradation pathways: moisture, thermal, and light stability (Q10-Q12)
Defect & Recom.	Fundamental mechanisms affecting device efficiency through defect passivation and recombination control (Q13-Q14)
Interface	Interface engineering and charge transport optimization (Q15-Q17)
Materials	Comprehensive analysis of functional materials and their characteristics in different device components (Q18-Q21)

Table 1: Classification of Research Questions in Perovskite Solar Cell Studies

Evaluation Dataset For better evaluation, we design the evaluation dataset including multiple choice questions and question answering in the perovskite domain. The evaluation dataset also extract from the top level publications in perovskite domain with our multi-agent framework and extral expert double check. The evaluation dataset contains 1,103 question answering named **Perovskite QA** and 1,103 multiple choice questions named **Perovskite MCQ**.

For question answering, we set the Rouge-L score and the LLM-as-a-Judge (Zheng et al., 2023) score as the evaluation metric. In our experiments, we find that both metrics can effectively measure the quality of question answering and consistency with each other.

For multiple choice questions, we set the accuracy as the evaluation metric. Using LLaMA-3.1-8B-Instruct (Dubey et al., 2024) as the baseline model, the difficulty level of each question is determined by its zero shot performance. Specifically, if LLaMA-3.1-8B-Instruct can correctly answer a question in a zero-shot setting (without any task-specific training or prompt engineering), we classify it as an "easy" question. Conversely, questions that LLaMA-3.1-8B-Instruct fails to answer correctly are categorized as "Hard". This classification method resulted in 823 easy questions and 280 hard questions in our evaluation dataset, providing a balanced assessment of model capabilities across different difficulty levels.

To evaluate the performance of **Perovskite-**

Reasoning, we incorporated Minerva (Lewkowycz et al., 2022) and GPQA Diamond (Rein et al., 2023) as a benchmark. These contain undergraduate and PhD level science questions from biology, chemistry, and physics.

5 Perovskite-LLM

5.1 Experiment Design

In this section, we conduct the instruction tuning experiments on the **Perovskite-Chat** and **Perovskite-Reasoning** dataset. We select the LLaMA-3.1-8B-Instruct (Dubey et al., 2024) and Qwen-2.5-7B-Instruct (Yang et al., 2024a) as the baseline model, and **Perovskite-Chat-LLM** and **Perovskite-Reasoning-LLM** are fine-tuned version of LLaMA-3.1-8B-Instruct and Qwen-2.5-7B-Instruct with **Perovskite-Chat** and **Perovskite-Reasoning** dataset. For the training methodology, we employed a full fine-tuning approach. While Parameter-Efficient Fine-Tuning (PEFT) methods like LoRA (Hu et al., 2021) are prominent for their efficiency, some studies suggest they can exhibit weaker convergence on downstream tasks compared to full fine-tuning (Pan et al., 2024a). As our computational resources were sufficient, we opted for full fine-tuning to maximize model performance and adaptation to our specialized datasets. We leave the exploration of PEFT methods for future work, especially when scaling to larger models.

For the training process, we used the full parameter fine-tuning method to fine-tune the Perovskite-

LLM. The experiment is carried out on the A800 GPU server, with flash attention (Dao, 2024) and mixed precision training for efficient training. For more details on the training process, please refer to Appendix D.

For the evaluation process, we use the perplexity (PPL), Rouge-L score, LLM-Judge, and Human-Judge score to evaluate the performance on the Perovskite QA benchmark, the accuracy to evaluate the performance on the Perovskite MCQ benchmark, and the pass@1 rate on Minerva and GPQA benchmarks to evaluate the performance on the Perovskite-Reasoning-LLM. All experiments are conducted with zero-shot settings and three times to get the average results.

5.2 Results and analysis

Perovskite-Chat-LLM Table 3 presents the evaluation results of various models on the Perovskite QA dataset. The baseline models include GPT-3.5-Turbo, GPT-4o-Mini, GPT-4o, LLaMA-3.1-8B. Among these, Perovskite-Chat-LLM demonstrates superior performance across all metrics, achieving a perplexity of 2.97, a Rouge-L score of 41.25, and an LLM-Judge score of 2.97. This represents a significant improvement over the baseline LLaMA-3.1-8B model. The GPT family of models, while competitive in terms of LLM-Judge scores, showed lower performance in Rouge-L scores compared to Perovskite-Chat-LLM, with GPT-4o achieving 11.36 for Rouge-L and 1.41 for LLM-Judge. With the case study in Figure 3, we can see that Perovskite-Chat-LLM can generate more accurate and consistent answers compared to other models, and ChatGPT only can offer a general and nonspecific answer which leads to low performance on the Rouge-L and LLM-Judge metrics.

To further validate the effectiveness and credibility of our model, we incorporated human expert evaluation into our assessment methodology. Table 2 presents the results of this human expert evaluation alongside the automatic metrics. Notably, Perovskite-Chat-LLM achieved the highest scores in both LLM-Judge (2.97) and Human Expert (3.12) assessments, outperforming baseline models such as GPT-4o and LLaMA-3.1-8B. This consistency across both human and LLM-Judge evaluations demonstrates the robustness and reliability of our model, as well as its ability to generate responses that are not only accurate but also align with expert expectations in the field.

Model	LLM-Judge \uparrow	Human Expert \uparrow
GPT-4o	1.41	1.73
LLaMA-3.1-8B	1.28	1.17
Perovskite-Chat-LLM	2.97	3.12

Table 2: Comparison of automatic and human expert evaluation scores on the Perovskite QA dataset.

Model	Perovskite QA		
	PPL \downarrow	Rouge-L \uparrow	LLM-Judge \uparrow
GPT-3.5-Turbo	-	11.24	1.24
GPT-4o-Mini	-	11.90	1.34
GPT-4o	-	11.36	1.41
LLaMA-3.1-8B	6.77	13.18	1.28
LLaMA-3.1-70B	4.98	17.38	1.80
Qwen-2.5-7B	6.23	11.22	1.39
Qwen-2.5-72B	5.12	10.17	1.31
Perovskite-Chat-LLM	2.97	41.25	2.97
+Perovskite-KG	2.80	43.10	3.15

Table 3: Performance of Perovskite-Chat-LLM on Perovskite QA

Table 4 presents the evaluation results of various models on the Perovskite MCQ dataset, categorized into Easy, Hard, and All difficulty levels. Among these, GPT 4o achieves the highest overall score of 84.68, with scores of 91.37 for Easy and 65.00 for Hard questions. Perovskite-Chat-LLM, highlighted in red, shows strong performance with a score of 67.86 on Hard questions, the highest in this category, and an overall score of 81.96.

Model	Perovskite MCQ		
	Easy	Hard	All \uparrow
GPT-3.5-Turbo	86.63	49.29	77.15
GPT-4o-Mini	89.79	61.79	82.68
GPT-4o	91.37	65.00	84.68
LLaMA-3.1-8B	100.00	0.00	74.21
LLaMA-3.1-70B	93.44	66.43	86.58
Qwen-2.5-7B	92.22	55.36	82.86
Qwen-2.5-72B	93.07	64.29	85.77
Perovskite-Chat-LLM	95.50	62.86	87.22
+Perovskite-KG	97.32	65.10	89.21

Table 4: Performance of Perovskite-Chat-LLM on Perovskite MCQ. The LLaMA-3.1-8B baseline model’s performance defines Easy/Hard question categories.

Perovskite-Reasoning-LLM Table 5 presents the evaluation results of Perovskite-Reasoning-LLM compared against various baseline models on the GPQA and Minerva benchmarks. In the 7B-scale model category, our Perovskite-Reasoning-LLM achieves state-of-the-art performance with remarkable data efficiency over 7B models, requiring only 2K training examples compared to 800K for R1-Qwen2.5-7B and 114K for OpenThinker-

7B. Our model achieves 43.95 on GPQA and 44.49 on Minerva. Compared to 32B models, while our GPQA performance shows room for improvement (suggesting GPQA’s sensitivity to model size), our Minerva score (44.49) is competitive with larger models like LIMO-32B (44.90) and approaches S1-32B (47.79). This shows that our efficient training approach can achieve strong performance on STEM reasoning tasks even with a smaller model architecture.

Model	# ex	GPQA \uparrow	Minerva \uparrow	Avg \uparrow
API Models				
o1	-	77.30	-	-
o1-preview	-	73.30	47.10	60.20
o1-mini	-	60.00	-	-
Deepseek-R1	-	71.50	-	-
32B				
Qwen2.5-32B-Instruct	-	48.00	41.20	44.60
QwQ-32B-preview	-	65.10	39.00	52.05
LIMO-32B*	0.8K	66.70	44.90	55.80
S1-32B*	1K	59.60	47.79	53.69
7B				
LLaMA-3.1-8B Instruct	-	30.40	13.97	22.18
R1-Qwen2.5-7B*	800K	44.49	25.25	34.87
R1-LLaMA3-8B*	800K	19.19	30.51	24.85
OpenThinker-7B*	114K	42.90	41.10	42.00
Perovskite-R-LLM	2K	43.95	44.49	44.22

Table 5: We evaluate the performance of Perovskite-Reasoning-LLM on the GPQA and Minerva benchmarks. * indicates the results are from our evaluation. # ex = number of examples used for fine-tuning. # Perovskite-R-LLM : Perovskite-Reasoning-LLM

5.3 Integrate Perovskite-KG with Perovskite-Chat-LLM

To ensure the Perovskite-Chat-LLM possesses current and specialized knowledge, we employ retrieval-augmented generation (RAG) to dynamically integrate information from our Perovskite-KG. A key advantage of this approach is the continuous and automatic updating of the LLM’s knowledge base, which is directly facilitated by the automated construction methodology of the Perovskite-KG itself. The technical implementation of this integration leverages the LightRAG framework (Guo et al., 2024). The performance improvements resulting from the integration of Perovskite-KG with Perovskite-Chat-LLM (“+Perovskite-KG”) are presented in Table 3 and Table 4. When enhanced by the Perovskite-KG, the Perovskite-Chat-LLM demonstrates improved scores in both Perovskite QA (2.97 \rightarrow 3.15) and Perovskite MCQ (87.22 \rightarrow 89.21). Representative case studies illustrating these benefits are detailed in Figure 4 and Figure 10. The Perovskite-Chat-LLM, when augmented with

this knowledge integration, demonstrates an enhanced capability to access up-to-date information and deliver fine-grained, domain-specific responses compared to its baseline version lacking such integration. For instance, when presented with a query such as, “Can you tell me what is CuO? And give me 5 usage in perovskite?”, the enhanced model provides the requested details, with each stated usage accompanied by a citation from the underlying knowledge source.

5.4 Case Study

To illustrate the practical advantages of Perovskite-LLM over existing models, Figure 3 presents a comparative analysis of the responses of Perovskite-Chat-LLM and ChatGPT to a question about the fabrication of high-efficiency perovskite solar cells. The responses demonstrate a clear distinction in the depth and specificity of knowledge provided by each model. For instance, Perovskite-Chat-LLM specifies precise conditions for the HTL preparation (150°C for 15 minutes) and details a two-step spin-coating procedure (1,000 rpm for 10 seconds, followed by 6,000 rpm for 30 seconds) with specific process modifications like anisole addition. This level of detail reflects Perovskite-Chat-LLM’s specialized training on domain-specific literature and its ability to provide practical, implementation-ready guidance for researchers in the field. For more details, please refer to Appendix E.

6 Conclusion

In this work, we present a comprehensive knowledge-enhanced system for perovskite solar cell research, integrating three key components: (1) Perovskite-KG, a domain-specific knowledge graph containing 23,789 entities and 22,272 relationships; (2) multi-agent system for generating two complementary datasets, Perovskite-Chat and Perovskite-Reasoning, designed for domain-specific knowledge assistance and scientific reasoning respectively; and (3) two specialized large language models that demonstrate superior performance in both knowledge retrieval and reasoning tasks. Our experimental results show significant improvements over existing models, with Perovskite-Chat-LLM achieving state-of-the-art performance on domain-specific tasks and Perovskite-Reasoning-LLM showing competitive performance on scientific reasoning benchmarks despite using substantially fewer training exam-

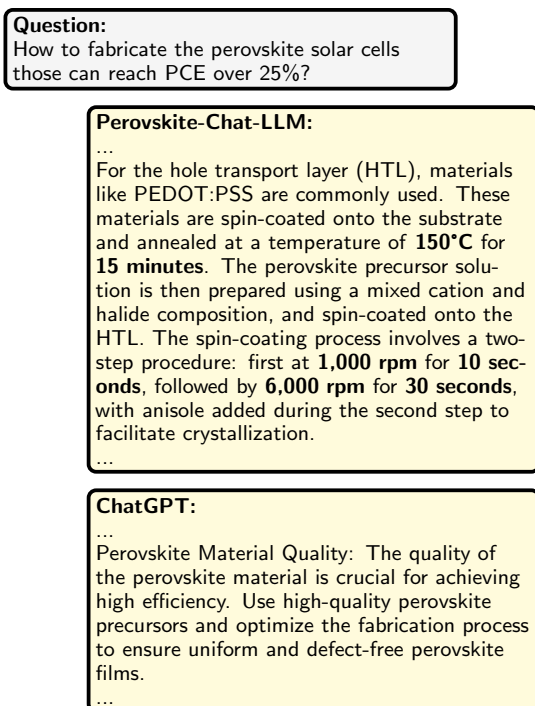


Figure 3: Comparison of responses between Perovskite-LLM and ChatGPT: Perovskite-LLM provides detailed operational steps with specific parameters, while ChatGPT only offers general conceptual guidance.

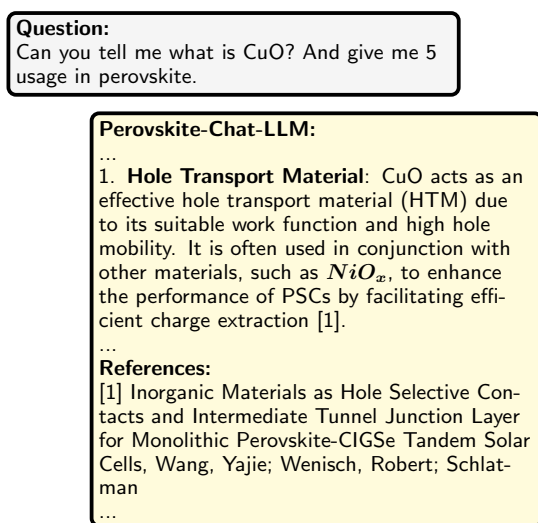


Figure 4: A case study of Perovskite-Chat-LLM's ability to provide detailed and accurate information with references.

ples. The system provides researchers with effective tools to review the literature, design experiments, and solve complex problems in PSC research. Our work demonstrates the potential of LLMs to accelerate innovation and discovery in materials science by enabling more efficient knowledge access and reasoning capabilities.

7 Limitations

Despite the promising results, our current system has several limitations that warrant future investigation:

- **Knowledge Coverage:** While our knowledge graph covers a substantial portion of the PSC literature, it may not capture all emerging research directions and novel experimental techniques. Future work should focus on developing mechanisms for automatic knowledge base expansion and updates to maintain its relevance.
- **Model Size Trade-offs:** While our 7B-parameter models achieve competitive performance, there might be certain complex tasks that benefit from larger model architectures, suggesting a potential trade-off between efficiency and capability. Future work will investigate model compression techniques and more efficient architectures.

To address these limitations, our future work will focus on three main directions: (1) developing a continuous knowledge integration framework that can automatically update the knowledge base with new research findings, (2) enhancing the reasoning capabilities through advanced model architectures and training strategies, and (3) improving the system's practical utility through better validation mechanisms and more efficient model designs. These improvements will make the system more robust, up-to-date, and accessible to the broader research community.

Acknowledgments

This work was supported by the Advanced Materials-National Science and Technology Major Project (Grant No. 2025ZD0620100), Guangdong Basic and Applied Basic Research Foundation 2025A1515010304, Guangzhou Science and Technology Planning Project 2025A03J4491

References

Yuan An, Jane Greenberg, Alex Kalinowski, Xintong Zhao, Xiaohua Hu, Fernando J. Uribe-Romo, Kyle Langlois, Jacob Furst, and Diego A. Gómez-Gualdrón. 2024. [Knowledge graph question answering for materials science \(kgqa4mat\): Developing natural language interface for metal-organic frameworks knowledge graph \(mof-kg\) using llm](#). *Preprint*, arXiv:2309.11361.

- Yuan An, Jane Greenberg, Xintong Zhao, Xiaohua Hu, Scott McClellan, Alex Kalinowski, Fernando J Uribe-Romo, Kyle Langlois, Jacob Furst, Diego A Gómez-Gualdrón, et al. 2022. Building open knowledge graph for metal-organic frameworks (mof-kg): Challenges and case studies. In *28TH ACM. SIGKDD. CONFERENCE. ON KNOWLEDGE DISCOVERY. AND DATA MINING*.
- Tze-Zhang Ang, Mohamed Salem, Mohamad Kamarol, Himadry Shekhar Das, Mohammad Alhuyi Nazari, and Natarajan Prabakaran. 2022. A comprehensive study of renewable energy sources: Classifications, challenges and suggestions. *Energy Strategy Reviews*, 43:100939.
- Dmitrii Bogdanov, Javier Farfan, Kristina Sadovskaia, Arman Aghahosseini, Michael Child, Ashish Gulagi, Ayobami Solomon Oyewo, Larissa de Souza Noel Simas Barbosa, and Christian Breyer. 2019. Radical transformation pathway towards sustainable electricity via evolutionary steps. *Nature communications*, 10(1):1–16.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26.
- Juan-Pablo Correa-Baena, Michael Saliba, Tonio Buonassisi, Michael Grätzel, Antonio Abate, Wolfgang Tress, and Anders Hagfeldt. 2017. Promises and challenges of perovskite solar cells. *Science*, 358(6364):739–744.
- Tri Dao. 2024. FlashAttention-2: Faster attention with better parallelism and work partitioning. In *International Conference on Learning Representations (ICLR)*.
- Shizhe Diao, Pengcheng Wang, Yong Lin, Rui Pan, Xiang Liu, and Tong Zhang. 2024. [Active prompting with chain-of-thought for large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1330–1350, Bangkok, Thailand. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Mikhail Galkin, Xinyu Yuan, Hesham Mostafa, Jian Tang, and Zhaocheng Zhu. Towards foundation models for knowledge graph reasoning. In *The Twelfth International Conference on Learning Representations*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Taicheng Guo, Kehan Guo, Bozhao Nan, Zhenwen Liang, Zhichun Guo, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. 2023. [What can large language models do in chemistry? a comprehensive benchmark on eight tasks](#). *Preprint, arXiv:2305.18365*.
- Zirui Guo, Lianghao Xia, Yanhua Yu, Tu Ao, and Chao Huang. 2024. Lightrag: Simple and fast retrieval-augmented generation. *arXiv preprint arXiv:2410.05779*.
- Tanishq Gupta, Mohd Zaki, Devanshi Khatsuriya, Kausik Hira, NM Anoop Krishnan, et al. 2023. Discomat: Distantly supervised composition extraction from tables in materials science articles. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13465–13483.
- Tanishq Gupta, Mohd Zaki, NM Anoop Krishnan, and Mausam. 2022. Matscibert: A materials domain language model for text mining and information extraction. *npj Computational Materials*, 8(1):102.
- Jiye Han, Keonwoo Park, Shaun Tan, Yana Vaynzof, Jingjing Xue, Eric Wei-Guang Diao, Mounqi G Bawendi, Jin-Wook Lee, and Il Jeon. 2025. Perovskite solar cells. *Nature Reviews Methods Primers*, 5(1):3.
- Kausik Hira, Mohd Zaki, Dhruvil Sheth, NM Anoop Krishnan, et al. 2024. Reconstructing the materials tetrahedron: challenges in materials information extraction. *Digital Discovery*, 3(5):1021–1037.
- Sirui Hong, Yizhang Lin, Bang Liu, Bangbang Liu, Binhao Wu, Danyang Li, Jiaqi Chen, Jiayi Zhang, Jinlin Wang, Li Zhang, Lingyao Zhang, Min Yang, Mingchen Zhuge, Taicheng Guo, Tuo Zhou, Wei Tao, Wenyi Wang, Xiangru Tang, Xiangtao Lu, Xiawu Zheng, Xinbing Liang, Yaying Fei, Yuheng Cheng, Zongze Xu, and Chenglin Wu. 2024a. Data interpreter: An LLM agent for data science. *CoRR*, abs/2402.18679.
- Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. 2024b. Metagpt: Meta programming for A multi-agent collaborative framework. In *ICLR*. OpenReview.net.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Chen, Alex Smola, and Robert Korbak. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Olga Kononova, Tanjin He, Haoyan Huo, Amalie Trevartha, Elsa A Olivetti, and Gerbrand Ceder. 2021. Opportunities and challenges of text mining in materials research. *IScience*, 24(3).

- Bespoke Labs. 2025. [Bespoke-stratos: The unreasonable effectiveness of reasoning distillation](#). Accessed: 2025-01-22.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. 2022. Solving quantitative reasoning problems with language models. *Advances in Neural Information Processing Systems*, 35:3843–3857.
- Boyan Li, Yuyu Luo, Chengliang Chai, Guoliang Li, and Nan Tang. 2024. The dawn of natural language to sql: Are we fully ready? *arXiv preprint arXiv:2406.01265*.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhunoye, Yiming Yang, et al. 2023. Self-refine: Iterative refinement with self-feedback. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Santiago Miret and Nandan M Krishnan. 2024. Are llms ready for real-world materials discovery? *arXiv preprint arXiv:2402.05200*.
- Vaibhav Mishra, Somaditya Singh, Dhruv Ahlawat, Mohd Zaki, Vaibhav Bihani, Hargun Singh Grover, Biswajit Mishra, Santiago Miret, NM Krishnan, et al. 2024. Foundational large language models for materials research. *arXiv preprint arXiv:2412.09560*.
- David Mrdjenovich, Matthew K Horton, Joseph H Montoya, Christian M Legaspi, Shyam Dwaraknath, Vahe Tshitoyan, Anubhav Jain, and Kristin A Persson. 2020. Propnet: a knowledge graph for materials science. *Matter*, 2(2):464–480.
- Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. 2025. s1: Simple test-time scaling. *arXiv preprint arXiv:2501.19393*.
- National Renewable Energy Laboratory. 2025. [Best research-cell efficiency chart](#). Photovoltaic research, National Renewable Energy Laboratory (NREL).
- Harsha Nori, Yin Tat Lee, Sheng Zhang, Dean Carignan, Richard Edgar, Nicolò Fusi, Nicholas King, Jonathan Larson, Yuanzhi Li, Weishung Liu, Renqian Luo, Scott Mayer McKinney, Robert Osazuwa Ness, Hoi-fung Poon, Tao Qin, Naoto Usuyama, Chris White, and Eric Horvitz. 2023. Can generalist foundation models outcompete special-purpose tuning? case study in medicine. *CoRR*, abs/2311.16452.
- OpenAI. 2024. [Learning to reason with llms](#).
- Rui Pan, Xiang Liu, Shizhe Diao, Renjie Pi, Jipeng Zhang, Chi Han, and Tong Zhang. 2024a. Lisa: Layerwise importance sampling for memory-efficient large language model fine-tuning. *Advances in Neural Information Processing Systems*, 37:57018–57049.
- Rui Pan, Shuo Xing, Shizhe Diao, Wenhe Sun, Xiang Liu, KaShun Shum, Jipeng Zhang, Renjie Pi, and Tong Zhang. 2024b. [Plum: Prompt learning using metaheuristics](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 2177–2197, Bangkok, Thailand. Association for Computational Linguistics.
- Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jipu Wang, and Xindong Wu. 2024c. Unifying large language models and knowledge graphs: A roadmap. *IEEE Transactions on Knowledge and Data Engineering*.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2023. Gpqa: A graduate-level google-proof q&a benchmark. *arXiv preprint arXiv:2311.12022*.
- Tal Ridnik, Dedy Kredo, and Itamar Friedman. 2024. Code generation with alphacodium: From prompt engineering to flow engineering. *CoRR*, abs/2401.08500.
- Jayant Sathaye, Oswaldo Lucon, Atiq Rahman, John M Christensen, Fatima Denton, Junichi Fujino, Garvin Heath, Monirul Mirza, Hugh Rudnick, August Schlaepfer, et al. 2011. Renewable energy in the context of sustainable development. In *IPCC special report on renewable energy sources and climate change mitigation*. Cambridge University Press.
- Mara Schilling-Wilhelmi, Martiño Ríos-García, Sherjeel Shabih, María Victoria Gil, Santiago Miret, Christoph T Koch, José A Márquez, and Kevin Maik Jablonka. 2025. From text to insight: large language models for chemical data extraction. *Chemical Society Reviews*.
- Martin Seifrid, Robert Pollice, Andres Aguilar-Granda, Zamyra Morgan Chan, Kazuhiro Hotta, Cher Tian Ser, Jenya Vestfrid, Tony C Wu, and Alan Aspuru-Guzik. 2022. Autonomous chemical experiments: Challenges and perspectives on establishing a self-driving lab. *Accounts of Chemical Research*, 55(17):2454–2466.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Abubakr Babiker, Nathanael Schärli, Aakanksha Chowdhery, Philip Mansfield, Dina Demner-Fushman, Blaise Agüera y Arcas, Dale Webster, Greg S. Corrado, Yossi Matias, Katherine Chou, Juraj Gottweis, Nenad Tomasev, Yun Liu, Alvin Rajkomar, Joelle Barral, Christopher Semturs, Alan Karthikesalingam, and Vivek Natarajan. 2023. [Large language models encode clinical knowledge](#). *Nature*, 620(7972):172–180.
- Henry J Snaith. 2018. Present status and future prospects of perovskite photovoltaics. *Nature materials*, 17(5):372–376.

- Yaoxian Song, Penglei Sun, Haoyu Liu, Zhixu Li, Wei Song, Yanghua Xiao, and Xiaofang Zhou. 2024. Scene-driven multimodal knowledge graph construction for embodied ai. *IEEE Transactions on Knowledge and Data Engineering*.
- Yu Song, Santiago Miret, and Bang Liu. 2023a. Matscinnlp: Evaluating scientific language models on materials science language tasks using text-to-schema modeling. *arXiv preprint arXiv:2305.08264*.
- Yu Song, Santiago Miret, Huan Zhang, and Bang Liu. 2023b. Honeybee: Progressive instruction finetuning of large language models for materials science. *arXiv preprint arXiv:2310.08511*.
- Michael J Statt, Brian A Rohr, Dan Guevarra, Santosh K Suram, John M Gregoire, et al. 2023. The materials experiment knowledge graph. *Digital Discovery*, 2(4):909–914.
- Jiamin Su, Yibo Yan, Zhuoran Gao, Han Zhang, Xiang Liu, and Xuming Hu. 2025. Cafes: A collaborative multi-agent framework for multi-granular multimodal essay scoring. *arXiv preprint arXiv:2505.13965*.
- Liangtai Sun, Yang Han, Zihan Zhao, Da Ma, Zhennan Shen, Baocai Chen, Lu Chen, and Kai Yu. 2023. Scieval: A multi-level large language model evaluation benchmark for scientific research. *arXiv preprint arXiv:2308.13149*.
- Matthew C Swain and Jacqueline M Cole. 2016. Chemdataextractor: a toolkit for automated extraction of chemical information from the scientific literature. *Journal of chemical information and modeling*, 56(10):1894–1904.
- Kimi Team, Angang Du, Bofei Gao, Bower Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, Chuning Tang, Congcong Wang, Dehao Zhang, Enming Yuan, Enzhe Lu, Fengxiang Tang, Flood Sung, Guangda Wei, Guokun Lai, Haiqing Guo, Han Zhu, Hao Ding, Hao Hu, Hao Yang, Hao Zhang, Haotian Yao, Haotian Zhao, Haoyu Lu, Haoze Li, Haozhen Yu, Hongcheng Gao, Huabin Zheng, Huan Yuan, Jia Chen, Jianhang Guo, Jianlin Su, Jianzhou Wang, Jie Zhao, Jin Zhang, Jingyuan Liu, Junjie Yan, Junyan Wu, Lidong Shi, Ling Ye, Longhui Yu, Mengnan Dong, Neo Zhang, Ningchen Ma, Qiwei Pan, Qucheng Gong, Shaowei Liu, Shengling Ma, Shupeng Wei, Sihan Cao, Siying Huang, Tao Jiang, Weihao Gao, Weimin Xiong, Weiran He, Weixiao Huang, Wenhao Wu, Wenyang He, Xianghui Wei, Xianqing Jia, Xingzhe Wu, Xinran Xu, Xinxing Zu, Xinyu Zhou, Xuehai Pan, Y. Charles, Yang Li, Yangyang Hu, Yangyang Liu, Yanru Chen, Yejie Wang, Yibo Liu, Yidao Qin, Yifeng Liu, Ying Yang, Yiping Bao, Yulun Du, Yuxin Wu, Yuzhi Wang, Zaida Zhou, Zhaoji Wang, Zhaowei Li, Zhen Zhu, Zheng Zhang, Zhexu Wang, Zhilin Yang, Zhiqi Huang, Zihao Huang, Ziyao Xu, and Zonghan Yang. 2025. [Kimi k1.5: Scaling reinforcement learning with llms](#). *Preprint*, arXiv:2501.12599.
- NovaSky Team. 2025. [Sky-t1: Fully open-source reasoning model with o1-preview performance in \\$450 budget](#). Accessed: 2025-01-09.
- Amalie Trewartha, Nicholas Walker, Haoyan Huo, Sanghoon Lee, Kevin Cruse, John Dagdelen, Alexander Dunn, Kristin A. Persson, Gerbrand Ceder, and Anubhav Jain. 2022. [Quantifying the advantage of domain-specific pre-training on named entity recognition tasks in materials science](#). *Patterns*.
- Vahe Tshitoyan, John Dagdelen, Leigh Weston, Alexander Dunn, Ziqin Rong, Olga Kononova, Kristin A. Persson, Gerbrand Ceder, and Anubhav Jain. 2019. [Unsupervised word embeddings capture latent knowledge from materials science literature](#). *Nature*.
- Vineeth Venugopal and Elsa Olivetti. 2024. Matkg: An autonomously generated knowledge graph in material science. *Scientific Data*, 11(1):217.
- Nicholas Walker, Amalie Trewartha, Haoyan Huo, Sanghoon Lee, Kevin Cruse, John Dagdelen, Alexander Dunn, Kristin Persson, Gerbrand Ceder, and Anubhav Jain. 2021. The impact of domain-specific pre-training on named entity recognition tasks in materials science. *Available at SSRN 3950755*.
- Yuwei Wan, Yixuan Liu, Aswathy Ajith, Clara Grazian, Bram Hoex, Wenjie Zhang, Chunyu Kit, Tong Xie, and Ian Foster. 2024. [Sciqaq: A framework for auto-generated science question answering dataset with fine-grained evaluation](#). *Preprint*, arXiv:2405.09939.
- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*.
- Zhenhailong Wang, Shaoguang Mao, Wenshan Wu, Tao Ge, Furu Wei, and Heng Ji. 2024. Unleashing the emergent cognitive synergy in large language models: A task-solving agent through multi-persona self-collaboration. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 257–279.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Tianhao Wu, Zhenzhen Qin, Yanbo Wang, Yongzhen Wu, Wei Chen, Shufang Zhang, Molang Cai, Songyuan Dai, Jing Zhang, Jian Liu, et al. 2021. The main progress of perovskite solar cells in 2020–2021. *Nano-Micro Letters*, 13:1–18.

- Yupeng Xie, Yuyu Luo, Guoliang Li, and Nan Tang. 2024. Haichart: Human and AI paired visualization system. *Proc. VLDB Endow.*, 17(11):3178–3191.
- Haotian Xu, Xing Wu, Weinong Wang, Zhongzhi Li, Da Zheng, Boyuan Chen, Yi Hu, Shijia Kang, Jiaming Ji, Yingying Zhang, Zhijiang Guo, Yaodong Yang, Muhan Zhang, and Debing Zhang. 2025. [Redstar: Does scaling long-cot data unlock better slow-reasoning systems?](#) *Preprint*, arXiv:2501.11284.
- Yiheng Xu, SU Hongjin, Chen Xing, Boyu Mi, Qian Liu, Weijia Shi, Binyuan Hui, Fan Zhou, Yitao Liu, Tianbao Xie, et al. 2024. Lemur: Harmonizing natural language and code for language agents. In *The Twelfth International Conference on Learning Representations*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024a. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Chuang Yang, Wenjing Hu, Jiale Liu, Chuangzhou Han, Qiaojiao Gao, Anyi Mei, Yinhua Zhou, Fengwan Guo, and Hongwei Han. 2024b. Achievements, challenges, and future prospects for industrialization of perovskite solar cells. *Light: Science & Applications*, 13(1):227.
- Yilin Ye, Jianing Hao, Yihan Hou, Zhan Wang, Shishi Xiao, Yuyu Luo, and Wei Zeng. 2024. Generative AI for visualization: State of the art and future directions. *Vis. Informatics*, 8(1):43–66.
- Yixin Ye, Zhen Huang, Yang Xiao, Ethan Chern, Shijie Xia, and Pengfei Liu. 2025. Limo: Less is more for reasoning. *arXiv preprint arXiv:2502.03387*.
- Huan Zhang, Yu Song, Ziyu Hou, Santiago Miret, and Bang Liu. 2024a. [HoneyComb: A flexible LLM-based agent system for materials science](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3369–3382, Miami, Florida, USA. Association for Computational Linguistics.
- Huan Zhang, Yu Song, Ziyu Hou, Santiago Miret, and Bang Liu. 2024b. Honeycomb: A flexible llm-based agent system for materials science. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3369–3382.
- Jiayi Zhang, Chuang Zhao, Yihan Zhao, Zhaoyang Yu, Ming He, and Jianping Fan. 2024c. Mobileexperts: A dynamic tool-enabled agent team in mobile devices. *CoRR*, abs/2407.03913.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623.
- Li Zhong, Zilong Wang, and Jingbo Shang. 2024a. Debug like a human: A large language model debugger via verifying runtime execution step by step. In *ACL (Findings)*, pages 851–870. Association for Computational Linguistics.
- Qihuang Zhong, Kang Wang, Ziyang Xu, Juhua Liu, Liang Ding, Bo Du, and Dacheng Tao. 2024b. Achieving ζ 97% on gsm8k: Deeply understanding the problems makes llms perfect reasoners. *arXiv preprint arXiv:2404.14963*.
- Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. 2024. Language agent tree search unifies reasoning, acting, and planning in language models. In *Forty-first International Conference on Machine Learning*.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. 2023a. [Lima: Less is more for alignment](#). *Preprint*, arXiv:2305.11206.
- Xuanhe Zhou, Guoliang Li, and Zhiyuan Liu. 2023b. Llm as dba. *arXiv preprint arXiv:2308.05481*.
- Xiangru Zhu, Zhixu Li, Xiaodan Wang, Xueyao Jiang, Penglei Sun, Xuwu Wang, Yanghua Xiao, and Nicholas Jing Yuan. 2022. Multi-modal knowledge graph construction and application: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 36(2):715–735.
- Mingchen Zhuge, Haozhe Liu, Francesco Faccio, Dylan R Ashley, Róbert Csordás, Anand Gopalakrishnan, Abdullah Hamdi, Hasan Abed Al Kader Hammoud, Vincent Herrmann, Kazuki Irie, et al. 2023. Mindstorms in natural language-based societies of mind. *arXiv preprint arXiv:2305.17066*.

A Schema in Perovskite-KG

Table 6 presents a comprehensive schema for the Perovskite-KG, organized into three main ontological categories: Fabrication, Parameters, and Performance. The Fabrication ontology encompasses process-related attributes such as coating parameters, methods, and annealing conditions. The Parameters ontology covers structural and compositional aspects including solvents, device architecture, and additives. The Performance ontology captures various stability metrics and efficiency parameters like thermal stability, light stability, and power conversion efficiency. Each category is further detailed with specific data types and examples to ensure precise knowledge representation. This structured schema enables systematic organization and retrieval of perovskite solar cell information while maintaining data consistency across the knowledge graph.

B Prompts

The system employs four specialized agents, each with carefully designed prompts to perform specific tasks in the perovskite solar cell knowledge processing pipeline:

1. **Information Extraction Agent** (Table 7): Processes research papers using a structured set of 20 predefined questions across seven key categories, including device structure, performance enhancement, stability, and materials. The agent returns answers in a standardized JSON format, marking unavailable information as "Not mentioned" to maintain data quality.

2. **Verification Agent** (Table 8): Validates extracted information by comparing it with source texts, focusing on maintaining accuracy of technical details like numerical values and material names. The agent provides both corrected content and justification for any modifications made.

3. **Organization Agent** (Table 9): Synthesizes verified information from multiple sources into coherent, topic-focused responses. This agent ensures that complex technical information is presented in a logical and accessible manner.

4. **LLM-Judge** (Table 10): Evaluates response quality across four key criteria: accuracy, completeness, relevance, and clarity. Using a 1-5 scoring system, this agent provides detailed assessments and explanations for each criterion, along with an overall evaluation summary.

For LLM using in Perovskite-KG and Perovskite-Chat, we mainly using GPT-4o and GPT-4o-Mini for extraction and verification.

C Instruction Tuning Dataset

C.1 Dataset Statistics

The research questions in perovskite solar cell studies are systematically categorized in Tables 1 and 11. Table 1 provides a high-level overview of seven major research categories, including Device Structure and Fabrication, Performance Enhancement Strategies, Performance Metrics Improvement, Stability Improvements, Defect and Recombination Management, Interface and Extraction Layer Enhancements, and Materials Used in Perovskite Solar Cells. Each category is accompanied by a detailed rationale explaining its scope and relevance. Table 11 further expands this classification by presenting 21 specific research questions (Q1-Q21) organized within these seven categories. The questions cover a wide range of technical aspects, from device architecture optimization and performance enhancement strategies to material characteristics and stability improvements. Each research question is paired with its corresponding technical focus, providing a comprehensive framework for understanding the key areas of investigation in high-performance perovskite solar cell research.

Table 12 shows the distribution of research categories in perovskite solar cells. Device Structure and Fabrication dominates the field, accounting for 24,198 entries (43.8% of total). Performance Enhancement Strategies represents the second largest category with 11,233 entries (20.3%), followed by Defect and Recombination Management with 7,209 entries (13.0%). Stability Improvements, a crucial aspect of perovskite solar cell development, comprises 5,399 entries (9.8%), while Performance Metrics Improvement accounts for 4,527 entries (8.2%). Materials Used in Perovskite Solar Cells and Interface and Extraction Layer Enhancements represent smaller but significant portions of the research focus, with 1,586 (2.9%) and 952 (1.7%) entries respectively.

Figure 5 shows the length distribution patterns across categories in our perovskite instruction tuning dataset. All categories display similar characteristics, with sequence lengths predominantly ranging from 100 to 500 tokens and a median length of 400 tokens. This analysis informs our model design decisions, particularly regarding sequence length

Ontology	Sub-Category	Data Type	Description	Example
Fabrication	Coating Parameter	Float	The specifics of the coating method used in the material deposition process.	5000 rpm, 100 μ l
	Method	String	Different fabrication techniques, involving variations in material deposition.	spin coating
	Annealing Parameter	Float	Refers to the heating conditions applied to the perovskite, which are essential for crystallization and stability.	120°C, 10min
Parameters	Solvent	String	the liquid medium used to dissolve precursors, helping to form a uniform perovskite layer	Dimethylformamide (DMF)
	Device Structure	Patterned String	The architecture of the device (e.g., layer order, material interfaces)	ITO/SAM/perovskite /C60/BCP/Cu
	Additive	String	Any additional materials or chemicals	potassium ions
Performance	Thermal Stability	String	The material’s ability to withstand heat without degrading	>98% of initial efficiency of >24% after 1,500 hours of continuous maximum power point tracking
	Light Stability	String	How resistant the material is to prolonged exposure to light.	>92% of initial performance for 1,200 hours under the damp-heat test (85°C and 85% relative humidity)
	Moisture Stability	String	The material’s resilience against humidity or water exposure.	Initial PCE of control, target-1 and target-2 devices is 21.73%, 24.42% and 24.11%, respectively. Degraded to 78% of initial PCE after 1,500 hours at 55 \pm 5°C
	Fill Factor Value	Float	A measure of the device’s maximum power output.	0.88
	Open-Circuit Voltage Value	Float	The maximum voltage the device can produce under open-circuit conditions.	1.2 V
	Short-Circuit Current Value	Float	The current density when the circuit is closed.	25 mA/cm ²
	Power Conversion Efficiency Value	Float	The efficiency with which the device converts sunlight into electricity.	25 %

Table 6: Schema in Perovskite-KG.

limitations and tokenization strategies. The distributions are independently normalized for each category, enabling direct pattern comparison despite varying sample sizes.

Figure 6 illustrates the length distribution patterns across different categories in our perovskite instruction tuning dataset.

D Instruction Tuning

Table 13 presents the training hyperparameters for Perovskite-LLM.

E More Case Study

F License

GPQA (Rein et al., 2023) and Minerva (Lewkowycz et al., 2022) are under MIT license.




Figure 5: Distribution of prompt and response lengths across different categories in our dataset. The y-axis represents density ($e-3$), and the x-axis shows the word count in logarithmic scale. Each category's distribution is independently normalized.

Table 7: Prompts for Information Extraction Agent.

Prompts for Information Extraction Agent:

Answer the following questions based on the provided text.

```
{
  "Device Structure and Fabrication": [
    "Q1: Summarize the device structures or configurations of the perovskite solar cells those can reach PCE over 25%.",
    "Q2: How to prepare the perovskite precursor solutions those can reach PCE over 25%?",
    "Q3: How to fabricate the perovskite solar cells those can reach PCE over 25%?"
  ],
  "Performance Enhancement Strategies": [
    "Q4: What are problems solved in literatures that report perovskite solar cells those can reach PCE over 25%?",
    "Q5: What are the reasons to choose the strategies that can enhance performance of the perovskite solar cells in literatures that report perovskite solar cells those can reach PCE over 25%?"
  ],
  "Performance Metrics Improvement": [
    "Q6: How to improve the VOC of perovskite solar cells?",
    "Q7: How to improve the FF of perovskite solar cells?",
    "Q8: How to improve the Jsc of perovskite solar cells?"
  ],
  "Stability Improvements": [
    "Q9: How to improve the moisture stability of perovskite solar cells?",
    "Q10: How to improve the thermal stability of perovskite solar cells?",
    "Q11: How to improve the illumination or light stability of perovskite solar cells?"
  ],
  "Defect and Recombination Management": [
    "Q12: How to passivate or reduce defects/traps of perovskite solar cells?",
    "Q13: How to reduce recombination of perovskite solar cells?"
  ],
  "Interface and Extraction Layer Enhancements": [
    "Q14: How to improve the wettability of the buried interface in perovskite solar cells?",
    "Q15: How to improve the hole extraction ability of HTL in perovskite solar cells?",
    "Q16: How to improve the electron extraction ability of ETL in perovskite solar cells?"
  ],
  "Materials Used in Perovskite Solar Cells": [
    "Q17: What are the HTL materials used in perovskite solar cells and the common features of them?",
    "Q18: What are the ETL materials used in perovskite solar cells and their features?",
    "Q19: What are the hole blocking layer materials in perovskite solar cells and their features?",
    "Q20: What are the passivation materials used in perovskite solar cells and their common features?"
  ]
}
```

Below is the text: {paper_text}

Response: Return a JSON object with the following structure, if the text does not contain the answer, return "Not mentioned":

```
{
  "questions": [
    {
      "question": "Q1",
      "answer": "Answer to Question 1"
    },
    {
      "question": "Q2",
      "answer": "Not mentioned"
    },
    ...
  ]
}
```

Table 8: Prompts for Verification Agent.

Prompts for Verification Agent:

You need to verify the accuracy of the extracted information from a perovskite paper. Compare the extracted data with the original text to ensure consistency and correctness. Highlight any discrepancies and fix them. Moreover, maintain the original meaning of the text and the original information, such as numbers and material names.

Input:
 Paragraph {Section_name}:{Text_of_the_section}
 Extracted: {Extracted_information}

Output: Verified information with notes on any discrepancies or confirmation of accuracy.
 Please return a JSON object with the following structure only return one item:

```
{
  "verified_info": {
    "fixed_content": "The fixed paragraph",
    "reason": "The reason for the fix"
  }
}
```

Table 9: Prompts for Organization Agent.

Prompts for Organization Agent:

Your task is to organize the verified information from a perovskite paper related to the question: {question}. Below is the information split into paragraphs that answers the question:
 {answers}

Output: The organized and continuous answer to the question.

Return a JSON object with the following structure:

```
{
  "answer": "The organized and continuous answer to the question."
}
```

Table 10: Prompts for LLM-Judge.

Prompts for LLM-Judge:

You are an expert evaluator. Your task is to compare a model's response to the ground truth answer and provide a detailed evaluation.

Model's response:

{model_response}

Ground truth:

{ground_truth}

Please evaluate the model's response based on the following criteria:

1. Accuracy: How factually correct is the model's response compared to the ground truth?
2. Completeness: Does the model's response cover all the key points mentioned in the ground truth?
3. Relevance: How well does the model's response address the implied question or task?
4. Clarity: Is the model's response clear and easy to understand?

For each criterion, provide a score from 1 to 5, where 1 is the lowest and 5 is the highest. Also, provide a brief explanation for each score.

Finally, give an overall score from 1 to 5 and a summary of your evaluation.

Format your response as a JSON object with the following structure:

```
{
  "accuracy": { "score": 0, "explanation": "" },
  "completeness": { "score": 0, "explanation": "" },
  "relevance": { "score": 0, "explanation": "" },
  "clarity": { "score": 0, "explanation": "" },
  "overall": { "score": 0, "summary": "" }
}
```

ID	Research Question	Technical Focus
<i>I. Device Structure and Fabrication</i>		
Q1	Summarize device structures for PCE >25%	Device architecture optimization
Q2	Perovskite precursor solution preparation for PCE >25%	Solution chemistry and processing
Q3	Fabrication methods for PCE >25%	Manufacturing processes
<i>II. Performance Enhancement Strategies</i>		
Q4	Problems solved in high-efficiency (>25%) devices	Critical challenges and solutions
Q5	Rationale for performance enhancement strategies	Strategic approach justification
<i>III. Performance Metrics Improvement</i>		
Q6	V _{OC} improvement methods	Open-circuit voltage optimization
Q7	FF improvement methods	Fill factor enhancement
Q8	J _{SC} improvement methods	Short-circuit current density optimization
Q9	PLQY-iV _{OC} relationship	Photoluminescence quantum yield correlation
<i>IV. Stability Improvements</i>		
Q10	Moisture stability enhancement	Water resistance strategies
Q11	Thermal stability enhancement	Temperature tolerance methods
Q12	Light stability enhancement	Photo-stability improvement
<i>V. Defect and Recombination Management</i>		
Q13	Defect/trap passivation methods	Defect control strategies
Q14	Recombination reduction approaches	Charge recombination suppression
<i>VI. Interface and Extraction Layer Enhancements</i>		
Q15	Buried interface wettability improvement	Interface engineering
Q16	HTL hole extraction enhancement	Hole transport optimization
Q17	ETL electron extraction enhancement	Electron transport optimization
<i>VII. Materials Used in Perovskite Solar Cells</i>		
Q18	HTL materials and features	Hole transport materials
Q19	ETL materials and features	Electron transport materials
Q20	Hole blocking layer materials	Blocking layer characteristics
Q21	Passivation materials and features	Surface passivation materials

Table 11: Systematic Classification of Research Questions in High-Performance Perovskite Solar Cell Studies

Abbreviated Name	Full Name	Count
Perf. Enhancement	Performance Enhancement Strategies	11,233
Stability	Stability Improvements	5,399
Defect & Recom.	Defect and Recombination Management	7,209
Device Structure	Device Structure and Fabrication	24,198
Metrics	Performance Metrics Improvement	4,527
Materials	Materials Used in Perovskite Solar Cells	1,586
Interface	Interface and Extraction Layer Enhancements	952

Table 12: Correspondence between abbreviated and full names of research categories in perovskite solar cells

Parameter	Chat-LLM	Reasoning-LLM
Learning Rate	1e-6	5e-6
Batch Size	8	32
Number of Epochs	3	15
Optimizer	AdamW	AdamW
Learning Rate Scheduler	Cosine	Cosine
Warmup Steps	20	20

Table 13: Training Hyperparameters for Perovskite-Chat-LLM and Perovskite-Reasoning-LLM

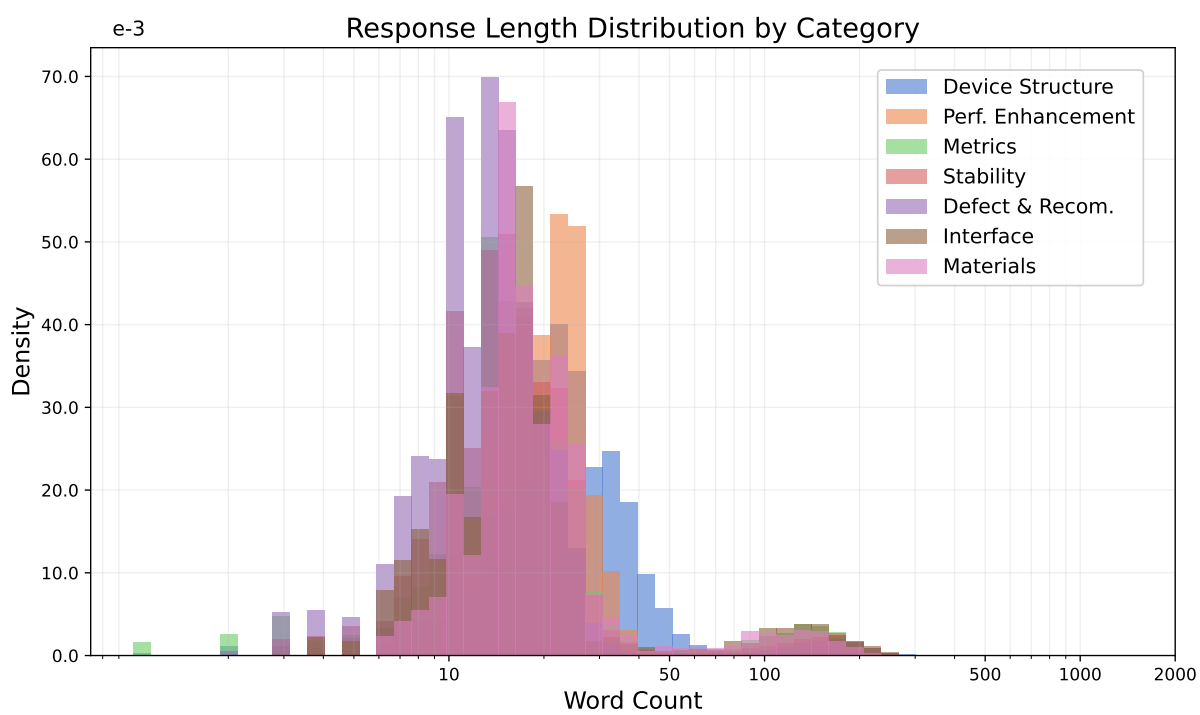



Figure 6: Distribution of prompt and response lengths across different categories in our dataset (top: prompt, bottom: response).

Question:
How to fabricate the perovskite solar cells those can reach PCE over 25%?

Perovskite-Chat-LLM:
To fabricate perovskite solar cells that can achieve a power conversion efficiency (PCE) over 25%, several key steps and considerations are essential. First, the fabrication process should start with a clean and dry substrate, typically glass or flexible substrates, which are cleaned thoroughly using a combination of detergents and solvents. The substrates are then treated with UV ozone to enhance their wetting properties.
For the hole transport layer (HTL), materials like PEDOT:PSS are commonly used. These materials are spin-coated onto the substrate and annealed at a temperature of **150°C for 15 minutes**. The perovskite precursor solution is then prepared, typically using a mixed cation and halide composition, and spin-coated onto the HTL. The spin-coating process involves a two-step procedure: first at **1,000 rpm for 10 seconds**, followed by **6,000 rpm for 30 seconds**, with anisole added during the second step to facilitate crystallization.
After spin-coating, the perovskite films are annealed at **100°C for 10 minutes**. The electron transport layer (ETL) is then deposited, often using a solution of PC61BM in chlorobenzene, which is spin-coated at **1,000 rpm for 30 seconds** and annealed at **100°C for 10 minutes**. Finally, a metal electrode, typically gold, is deposited through thermal evaporation. The entire process is conducted in a nitrogen-filled glovebox to minimize exposure to moisture.

ChatGPT:
To fabricate perovskite solar cells with a power conversion efficiency (PCE) over 25%, several key factors need to be considered and optimized in the fabrication process. Here are some steps that can help in achieving high-efficiency perovskite solar cells:
1. Perovskite Material Quality: The quality of the perovskite material is crucial for achieving high efficiency. Use high-quality perovskite precursors and optimize the fabrication process to ensure uniform and defect-free perovskite films.
2. Optimal Perovskite Layer Thickness: The thickness of the perovskite layer plays a significant role in determining the efficiency of the solar cell. Optimize the thickness of the perovskite layer to maximize light absorption and charge carrier generation.
3. Tuning the Bandgap: By engineering the composition of the perovskite material, you can tune the bandgap to match the solar spectrum for better light absorption and improved efficiency.
4. Interface Engineering: Pay attention to the interfaces between different layers in the solar cell structure. Optimize the interfaces to minimize recombination losses and improve charge carrier transport.
5. Passivation Layers: Use passivation layers to reduce defects and trap states in the perovskite layer, which can improve the overall performance of the solar cell.
6. Electron and Hole Transport Layers: Choose suitable electron and hole transport materials to facilitate efficient charge extraction and transport within the solar cell.

Figure 7: A full conversation between Perovskite-LLM and ChatGPT, Perovskite-LLM provides detailed operational steps with specific parameters, while ChatGPT only offers general conceptual guidance.

Question:
What are the ETL materials used in perovskite solar cells and the features of them?

Perovskite-Chat-LLM:
In perovskite solar cells, various electron transport layer (ETL) materials are utilized, each with distinct features. Common ETL materials include:

1. Tin Oxide (SnO₂): SnO₂ is a popular choice for ETLs due to its high electron mobility, low-temperature process ability, and compatibility with flexible substrates. It is often used in n-i-p structured perovskite solar cells.
2. Zinc Oxide (ZnO): ZnO is another widely used ETL material, known for its high electron mobility and low-temperature processing. It is commonly used in both n-i-p and p-i-n structured perovskite solar cells.
3. Tungsten Oxide (WO₃): WO₃ is utilized as an interlayer between the perovskite layer and the ETL, enhancing charge extraction and improving device stability. It is particularly effective in p-i-n structured devices.
4. Phenyl-C61-Butyric Acid Methyl Ester (PCBM): PCBM is a fullerene derivative used as an ETL in p-i-n structured perovskite solar cells. It offers high electron mobility and is compatible with various perovskite compositions.
5. Tin Oxide Doped with Fluorine (F-SnO₂): This material is used to enhance the performance of SnO₂ ETLs, improving their electron mobility and stability, which leads to higher efficiency in perovskite solar cells.
6. Cesium Lead Halide (CsPbX₃): CsPbX₃ is employed as a wide-bandgap ETL in all-inorganic perovskite solar cells, contributing to improved stability and efficiency.

...

ChatGPT:
ETL (Electron Transport Layer) materials play a crucial role in perovskite solar cells by facilitating the efficient extraction and transport of electrons generated by the perovskite layer. Some commonly used ETL materials in perovskite solar cells include:

1. TiO₂ (Titanium Dioxide): TiO₂ is one of the most widely used ETL materials in perovskite solar cells. It has good electron transport properties, high electron mobility, and is chemically stable. TiO₂ can be easily processed into thin films, making it suitable for large-scale production of solar cells.
2. SnO₂ (Tin Oxide): SnO₂ is another popular ETL material in perovskite solar cells. It has a high electron mobility, good transparency, and can be deposited at low temperatures. SnO₂ films can also be easily doped to improve their conductivity and electron extraction efficiency.

...

Figure 8: A full conversation between Perovskite-LLM and ChatGPT, Perovskite-LLM provides detailed information about the ETL materials and their features, while ChatGPT only offers general information about the ETL materials.

Question:
How to improve the thermal stability of perovskite solar cells?

Perovskite-Chat-LLM:
Improving the thermal stability of perovskite solar cells involves several strategies. One effective approach is the use of thermally stable materials for the charge transporting layer, such as nickel oxide (NiO) and copper thiocyanate (CuSCN), which can enhance the thermal stability of the devices. Additionally, the incorporation of additives like 1,8-diodoctane (DIO) and 1-chloronaphthalene (CN) has been shown to improve the morphology of perovskite films, leading to better thermal stability. Furthermore, the use of thermally stable hole transport materials (HTMs) and electron transport materials (ETMs) is crucial, as they can withstand high temperatures without degrading the perovskite layer. The development of new HTMs, such as poly(3-hexylthiophene) (P3HT), has also contributed to improved thermal stability. Lastly, the introduction of interfacial layers, like graphene oxide, can enhance the thermal stability of the perovskite layer by reducing defects and improving charge transport.

...

ChatGPT:
Improving the thermal stability of perovskite solar cells is crucial for their long-term performance and reliability. Here are some strategies to enhance their thermal stability:

1. Encapsulation: One of the most effective ways to improve the thermal stability of perovskite solar cells is to encapsulate them with protective materials. Encapsulation helps to shield the cells from environmental factors such as moisture, oxygen, and heat, which can degrade the perovskite material.
2. Improved Perovskite Composition: Researchers are constantly working on developing new perovskite compositions that are more stable at high temperatures. By tweaking the chemical composition of the perovskite material, it is possible to enhance its thermal stability.

...

Figure 9: A full conversation between Perovskite-LLM and ChatGPT, Perovskite-LLM provides detailed information about the thermal stability of perovskite solar cells, while ChatGPT only offers general information about the thermal stability of perovskite solar cells.

Question:
Can you tell me what is CuO? And give me 5 usage in perovskite.

Perovskite-Chat-LLM:
CuO refers to copper(II) oxide, a chemical compound that plays a significant role in various applications, particularly in perovskite solar cells (PSCs). Here are five key usages of CuO in perovskite technology:

- 1. Hole Transport Material:** CuO acts as an effective hole transport material (HTM) due to its suitable work function and high hole mobility. It is often used in conjunction with other materials, such as NiO_x , to enhance the performance of PSCs by facilitating efficient charge extraction [1].
- 2. Interface Optimization:** The use of CuO helps in optimizing the interfaces within perovskite solar cells. By depositing CuO films on top of perovskite layers, it can improve the interfacial properties and reduce recombination losses at these critical junctions[2].
- 3. Thermal Stability Enhancement:** Copper oxide (CuO) has been shown to provide thermal stability to perovskite materials, which is crucial for maintaining high efficiency during prolonged operation under elevated temperatures. This property makes it an attractive choice for applications where environmental stability is a concern [3].
- 4. Doping and Passivation:** CuO can be utilized as a dopant in other materials used in PSCs, such as SnO₂, enhancing their electrical properties like conductivity and mobility. Additionally, the doping effects can lead to improved passivation of defects at the perovskite layer, contributing to better performance metrics [4].
- 5. Integration with Other Materials:** In bifacial solar cells, CuO has been integrated with materials such as carbon nanotubes (CNTs) to create high-performance architectures that benefit from synergistic effects. This integration not only enhances charge transport but also improves the overall efficiency of the devices [5].

References:
[1] Inorganic Materials as Hole Selective Contacts and Intermediate Tunnel Junction Layer for Monolithic Perovskite-CIGSe Tandem Solar Cells, Wang, Yajie; Wenisch, Robert; Schlatman
[2] Critical Review of Cu-Based Hole Transport Materials for Perovskite Solar Cells: From Theoretical Insights to Experimental Validation, Sun, Qingde; Sadhu, Anupam; Lie, Stener; Wong, Lydia Helena, Advanced Materials, 2024
[3] Processing-Performance Evolution of Perovskite Solar Cells: From Large Grain Polycrystalline Films to Single Crystals, Haque, Md Azimul; Troughton, Joel; Baran, Derya, Advanced Energy Materials, 2020
[4] Air-Processed Perovskite Solar Cells with $\geq 25\%$ Efficiency and High Stability Enabled by Crystallization Modulation and Holistic Passivation, Shi, Xiaoyu; Liu, Tianxiao; Dou, Yunjie; Hu, Xiaodong; Liu, Yangyang; Wang, Feifei; Wang, Lingyuan; Ren, Zhijun; Chen, Shangshang, Advanced Materials, 2024
[5] Critical Review of Cu-Based Hole Transport Materials for Perovskite Solar Cells: From Theoretical Insights to Experimental Validation, Sun, Qingde; Sadhu, Anupam; Lie, Stener; Wong, Lydia Helena, Advanced Materials, 2024

Figure 10: A case study of Perovskite-Chat-LLM's ability to provide detailed and accurate information with references.