# What Would You Ask When You First Saw $a^2 + b^2 = c^2$? Evaluating LLM on Curiosity-Driven Question Generation

**Anonymous ACL submission**

## Abstract

Large language models (LLMs) are increasingly widely used as critical components of knowledge retrieval systems and agentic systems. These systems can benefit from knowledge-seeking capabilities of LLMs, in other words, curiosity. However, this capability has not been evaluated quantitatively. Toward bridging this gap, we propose an evaluation framework, CDQG (Curiosity-Driven Question Generation). The CDQG task prompts LLMs to generate questions about a statement introducing scientific knowledge, simulating a curious person when facing the statement for the first time. The CDQG dataset contains 1,988 statements including physics, chemistry, and mathematics with distinct levels of difficulty, general knowledge statements, and intentionally erroneous statements. We score the qualities of the questions generated by LLMs along multiple dimensions. These scores are validated by rigorous controlled ablation studies and human evaluations. While large models like GPT-4 and Mistral 8x7b can generate highly coherent and relevant questions, the smaller Phi-2 model is equally or more effective. This indicates that size does not solely determine a model's knowledge acquisition potential. CDQG quantifies a critical model capability, and opens up research opportunities for developing future knowledge retrieval systems driven by LLMs.

## 1 Introduction

Nowadays, large language models (LLMs) trained on internet-scale datasets are capable of storing and processing massive amounts of knowledge. LLMs are used as critical components of knowledge retrieval and processing systems, and the performance of these systems is related to the LLMs' capability to seek knowledge (Krishna et al., 2024; Huang and Huang, 2024; Gao et al., 2024).

However, to the best of our knowledge, this capability has not been evaluated quantitatively. Previous works in the literature assessed the capability to store knowledge (Liu et al., 2024a; Petroni et al., 2019), to be aware of the knowledge (Suzgun et al., 2024; Ferrando et al., 2024) and the capability to use knowledge (Zhu et al., 2024). We take an alternate perspective, assessing the capability of LLMs to *seek* knowledge.

Our setup is inspired by how humans seek knowledge: asking questions out of curiosity. Questioning is a key cognitive skill that underpins learning and knowledge acquisition. By asking questions, humans seek to understand the world around them, explore how things work, and challenge existing beliefs. This act of inquiry not only helps humans learn new information but also sharpens their thinking, promotes critical analysis, and drives innovation. Effective questioning fuels intellectual growth by sparking curiosity, encouraging deeper exploration of subjects, and improving comprehension (Acar et al., 2023). In education, questioning is closely linked to higher-level thinking skills like analysis, synthesis, and evaluation (Kurdi et al., 2020). The complexity and depth of questions asked often reflect a person's grasp and understanding of a topic (Kotov and Zhai, 2010).

Questions also play a crucial role in reasoning (Zelikman et al., 2024; Hao et al., 2023) since asking insightful questions requires logical thinking, clarifying assumptions, identifying knowledge gaps, and exploring alternative viewpoints (Lucas et al., 2024). OpenAI's o1 model uses its own "chain of thought" to engage in structured reasoning (OpenAI, 2024). Thoughtful questions are essential for thorough and logical reasoning (Ashok Kumar et al., 2023). Questioning is equally important for fact-checking. Good questions guide the verification process by identifying gaps, biases, and inconsistencies in the information (Li et al., 2017). Questions like "Does this agree with other sources?" or "Is this consistent with historical data?" lead to careful checking of facts and encourage cross-

referencing across multiple sources. Effective fact-checking requires context and nuance, and good questions can help reveal false or misleading information. Besides reasoning and fact-checking, questioning plays a major role in many other areas (Masterman et al., 2024), like encouraging creativity (Wang et al., 2024a), stimulating discussion, and driving innovation (Si et al., 2024; Ghafarollahi and Buehler, 2024). Thoughtful questions can open doors to new ideas and solutions.

Inspired by human questioning, we propose a framework, CDQG, that evaluates the LLMs' potential for discovering new knowledge. This framework is centered around a curiosity-driven question generation (CDQG) task, where a model is prompted to imagine itself as a human encountering a new statement for the first time, eliciting the most immediate questions that would arise. The questions are then scored along three metrics — relevance, coherence, and diversity — scores with roots in the literature of psychology (Zhao et al., 2023). We use state-of-the-art LLMs to compute these scores. The scores are validated by human judgment as well as rigorous ablation studies. Recent work by (Ke et al., 2024) explores how foundation models can independently gather information, highlighting parallel advancements in our field as we examine LLMs' curiosity-driven questioning.

We collect the CDQG dataset. The CDQG dataset contains 1,101 statements in physics, chemistry, and math, spanning across distinct levels of difficulty. Additionally, the CDQG dataset includes a section of 300 general knowledge statements and a special section of erroneous statements. CDQG challenges the models' critical inquiry skills and facilitates rigorous and generalizable evaluation.

Using the CDQG framework, we evaluate pre-trained language models of varying sizes, ranging from smaller ones like Phi-2 (Mojan Javaheripi, 2023) to larger models like GPT-4 (OpenAI et al., 2024) and Gemini (Gemini Team et al., 2024). While the larger models score high in coherence and relevance, the smaller Phi-2 model scores comparably well (or even better), indicating that the size might not be the only factor for the knowledge acquisition potential.

Our contributions can be summarized as follows:

- We introduce the CDQG framework, a novel approach for evaluating the ability of LLMs to generate questions given new information.
- We establish and validate a set of evaluation metrics, to systematically measure the depth and comprehensiveness of the questions generated by the LLMs.
- We compile the CDQG dataset, which includes varied and challenging content to test the questioning capabilities of LLMs.
- We conduct extensive testing with state-of-the-art LLMs to demonstrate the effectiveness of our framework through an ablation study.
- We highlight the practical applications of our findings in educational technology and AI-driven content creation.

To our knowledge, we are the first to introduce an evaluation framework assessing LLMs' questioning abilities based on knowledge statements. Our research encourages questioning-based evaluations to deepen the understanding of LLMs as critical components of knowledge-processing systems.

## 2 Related Works

### 2.1 Question Generation

Question generation has long been recognized as a critical task in education, with numerous studies underscoring its significance (Elkins et al., 2023; Kurdi et al., 2020). The evolution of this field has seen a progression from early rule-based question generation systems (Yao et al., 2022) to more sophisticated methods employing transformer-based models. Most recently, the application of LLMs represents the latest advancement in this area. The transition from rule-based systems to transformers, and ultimately to LLMs, highlights a shift towards utilizing deep learning techniques that better mimic human-like questioning abilities. This evolution enhances the relevance and quality of the generated questions, and also opens new possibilities for dynamic interactions within educational software (Abbasiantaeb et al., 2024) and conversation systems (Wang et al., 2024b).

### 2.2 Evaluation of Generative Models

In evaluating text generation from LLMs, recent methodologies have expanded beyond traditional metrics to include multifaceted approaches that align more closely with human judgment. GPTScore (Fu et al., 2023) and UniEval (Leiter et al., 2023) utilize the natural language understanding capabilities of LLMs to tailor evaluations to specific criteria, with GPTScore focusing on customized fluency and UniEval using a Boolean question-answering format for multiple quality dimensions. Similarly, CheckEval (Lee

et al., 2024) employs a structured checklist to enhance reliability, while X-Eval (Liu et al., 2024b) dynamically selects evaluation aspects, enhancing adaptability and depth. Further enriching these approaches are frameworks like the zero-shot comparative methodology (Liusie et al., 2024), which performs direct quality judgments, and the Unified Framework (Zhong et al., 2022), which combines traditional and specialized models for the assessment. PlanBench (Valmeekam et al., 2023) explores LLMs' reasoning through various planning tasks, while TIGERSCORE (Jiang et al., 2023) emphasizes explainability in evaluations. These are complemented by strategies that assess LLMs' ability to follow complex instructions (He et al., 2024) and a composite metric system that aggregates individual scores for a holistic view (Verga et al., 2024), enhancing the development and refinement of LLMs across different applications. However, these methodologies primarily center on how LLMs answer questions and perform predefined tasks, with little exploration into how effectively these models can generate meaningful questions themselves. Different from prior works, we focus explicitly on the questioning abilities of LLMs, introducing a new assessment dimension.

### 2.3 Prompt Engineering

Recent advancements in LLM evaluation have focused on optimizing prompting techniques to align more closely with human judgment. Studies show that LLM evaluations are more reproducible than human evaluations and effectively use a five-point Likert scale (Chiang and Lee, 2023). The G-EVAL framework improves evaluation accuracy by leveraging GPT-4 with chain-of-thought prompting (Liu et al., 2023a). Emphasizing prompt engineering, research demonstrates that well-crafted instructions and score aggregation significantly enhance LLM performance (Baswani et al., 2023). Additionally, smaller LLMs, guided by effective prompts, can match larger models' evaluation performance, highlighting the importance of prompt design (Kotonya et al., 2023). Moreover, reference-free evaluation methods show that LLMs can assess text quality through comparative judgments and rationales (Chen et al., 2023). These advancements informed our prompt engineering strategies.

### 2.4 LLMs for Evaluation

Recent studies highlight LLMs' potential to achieve human-level assessment quality in various tasks (Gilardi et al., 2023; Huang et al., 2024). The GEMBA framework, for instance, showcases the effectiveness of LLMs in reference-free machine translation evaluation (Kocmi and Federmann, 2023), while FrugalScore offers a streamlined approach by combining LLM-based metrics with lightweight models for efficient assessment (Kamal Eddine et al., 2022). Literature such as "Is ChatGPT a Good NLG Evaluator?" underscores ChatGPT's strong alignment with human judgments across NLG tasks (Wang et al., 2023). AUTOCALIBRATE enhances LLM-human alignment by iteratively refining evaluation criteria with human feedback (Liu et al., 2023b). Additionally, LLMs have proven effective in delivering relevance judgments with natural language explanations (Faggioli et al., 2023). Evaluations in machine translation and chatbot conversations show LLMs closely align with human ratings (Zheng et al., 2023). Instruction tuning has been shown to improve the correlation between LLM evaluations and human judgments (Xiong et al., 2024), while the development of explainable metrics emphasize the importance of transparency in LLM assessments (Leiter et al., 2024). While numerous methods have been introduced, many still have limitations and lack full robustness. In this paper, we propose a new framework that includes incremental noise addition to demonstrate the robustness of LLM evaluation without relying on human evaluation references.

## 3 CDQG framework

As summarized by Figure 1, this section describes the CDQG framework. CDQG specifically prompts models to ask questions elicited from intrinsic curiosity. CDQG then systematically evaluates these models across three critical performance metrics.

### 3.1 CDQG task

The CDQG task starts with a statement sampled from the CDQG dataset (which we'll explain in detail in Section 3.3). A prompt is constructed to accommodate the distinct instructional formats of multiple models. The prompting approach is centered around personification, where we ask each model to conceptualize itself as a human. This hypothetical human, encountering a statement for the first time and devoid of prior knowledge, is prompted to generate the top five questions that would instinctively arise. This allows us to elicit the models' inquisitive capabilities in a novel and

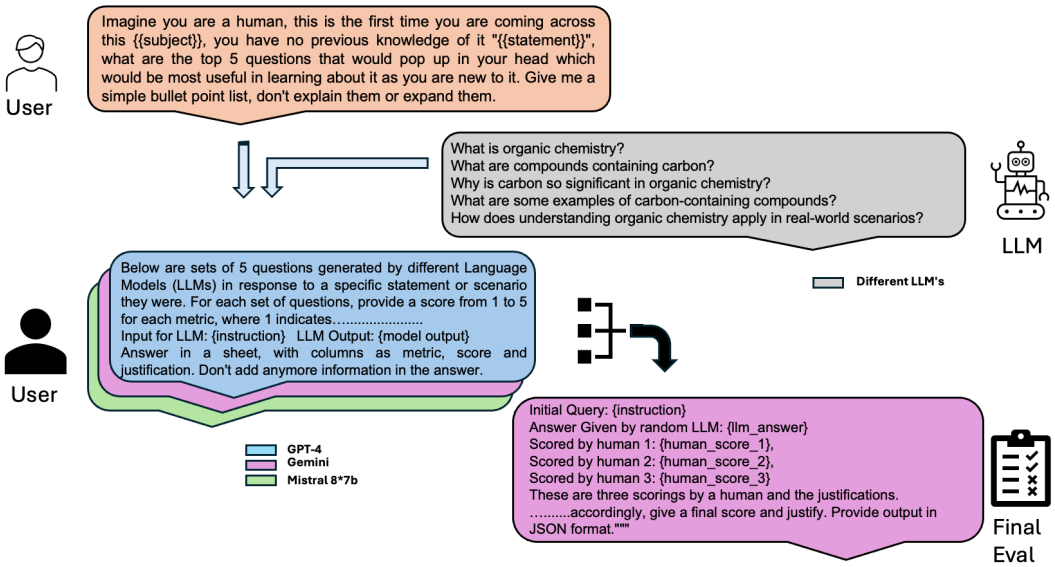


Figure 1: The CDQG framework. The top half shows the CDQG task, and the lower half shows the evaluation method of the generated questions.

controlled environment. The full prompt template is listed in Appendix B.

### 3.2 CDQG evaluation

The questions generated by the models are scored via a multi-dimensional evaluation procedure. The following three scores are computed:

*Relevance:* Relevance assesses how directly each question pertains to the specific details, elements, or concepts presented in the statement or scenario. The relevance criterion checks if questions aim to clarify, expand upon, or directly explore the content of the statement, focusing on the immediate context rather than the topics not directly introduced by the statement.

*Coherence:* Coherence assesses how logically the questions within each set connect to one another and whether they form a coherent line of inquiry that would logically progress a beginner's understanding of the topic. The coherence criterion checks if the sequence of questions or their thematic connection facilitates a structured exploration of the statement.

*Diversity:* Finally, diversity determines the range of aspects covered by the questions in relation to the statement, ensuring that each question brings a new dimension or perspective to understanding the statement. The diversity criterion checks if while maintaining direct relevance, the questions collectively offer a broad exploration of the topic, including but not limited to definitions, implications, applications, or theoretical underpinnings.

We use LLMs to score the generations on the aforementioned three dimensions, following the recent LLM-as-a-judge trend (Li et al., 2024). We select three large language models, GPT-3.5 Turbo, Mistral 8x7b, and Gemini, based on their accessibility, state-of-the-art performance characteristics, and diverse architectural approaches. For each specified metric, we prompt the LLM judge to generate a score in 5-point Likert scale and the corresponding justifications.

Then, we use Gemini as a "metareviewer" that summarizes the three evaluations (score with justification) into one final score, with a brief sentence as metareview. While the metareview sentence is not directly used to compare the models, it helps the Gemini models to provide a fair summary score.

To further ensure the validity of this evaluation protocol, we set up two validation experiments: an automatic noise-injection experiment and a human validation experiment. The details of the two validation studies are described in Section 6.

### 3.3 CQQG dataset

The CDQG dataset facilitates the CDQG evaluation framework. We leverage GPT-4's generative capabilities under human oversight to assemble the dataset incrementally (Xu et al., 2023), selecting statements that span diverse topics and complexity levels. Table 1 shows a breakdown of the dataset's splits and their corresponding sizes. We

| Subject | Split | | | | Total |
|---|---|---|---|---|---|
| | Basic | Intermediate | Advanced | Wrong | |
| Physics | 100 | 101 | 100 | 225 | 526 |
| Chemistry | 161 | 161 | 161 | 181 | 664 |
| Math | 108 | 108 | 101 | 181 | 498 |
| General | | | | | 300 |
| Total | 369 | 370 | 362 | 587 | 1,988 |

Table 1: Splits and sizes of the CDQG dataset.

consider the following desiderata when constructing the CDQG dataset.

**Multiple subjects** We include three subjects: chemistry, physics, and mathematics, to encompass a range of academic scenarios that an LLM may be useful. We additionally include general statements reflecting everyday life scenarios to broaden the coverage of the dataset.

**Distinct difficulty levels** For each of the academic subjects, we split the dataset into distinct difficulty levels, allowing stratified assessments of the LLMs' knowledge-seeking behavior regarding the statements with distinct levels of difficulty. Each level contains approximately the same number of statements to ensure a balanced distribution.

**Wrong statements** A unique feature of our dataset is the inclusion of these intentionally erroneous statements such as "The sum of 5 and 6 is 55", which probe the models' critical questioning abilities. These wrong statements span all three scientific domains, created by subtly modifying accurate statements. This subset tests whether models can identify and question statement veracity and logical consistency, particularly when treating the information as novel. We hypothesize that if a model operates as though it possesses prior knowledge, it will naturally question statement legitimacy. This dataset component serves as a critical test for evaluating models' depth of inquiry and their ability to critically engage with new information.

## 4 Models

We examine models ranging from a wide array of sizes: Llama 7b, Llama 13b, Llama 70b (Touvron et al., 2023), Mistral 8x7b (Jiang et al., 2024), Microsoft Phi-2 2.7b, Gemini, GPT 3.5 Turbo (Brown et al., 2020), and GPT-4. Our selection is based on practical considerations such as open-source availability and ease of access through APIs. Mistral's architecture, designed for handling complex queries, and Phi-2.7b's specialization in Q&A, make them well-suited for CDQG. By choosing models with varying architectures and parameter sizes, we ensure a broad comparison of model capabilities while maintaining accessibility and relevance to the task. The Gemini, GPT-3.5 turbo, and GPT-4 models are accessed using available APIs, and the other models (Llama-2 7b, Llama-2 13b, Llama-2 70b, Mistral 8x7b, Microsoft Phi-2) are accessed using the open-source weights, downloaded from Hugging face.

## 5 Results

Table 2, Figure 2 and Figure 3 illustrate our main results, with the rest in the Appendix.

### 5.1 Performance by model

**GPT-4:** Dominates in almost all metrics and subjects, especially in advanced tasks. This superior performance can be attributed to its extensive training on a diverse dataset, which equips it with a broad knowledge base and sophisticated reasoning capabilities.

**Mistral 8x7b:** Frequently matches or exceeds GPT-4, showing exceptional strength in Chemistry and Maths. Its use of a sparse mixture-of-experts architecture allows it to efficiently manage specific query types, demonstrating the benefits of mixture-of-experts architecture.

**Phi-2:** Phi-2's performance is particularly noteworthy. Despite its smaller scale of 2.7 billion parameters, Phi-2 consistently produces relevant and coherent questions at basic to intermediate task levels. This model benefits significantly from high-quality, curated training data that emphasizes "textbook-quality" content (Mojan Javaheripi, 2023), enhancing its capability in logical reasoning and common-sense understanding. Additionally, Phi-2's architecture leverages a scaled knowledge transfer (Mojan Javaheripi, 2023) from its predecessor, Phi-1.5, which improves its performance on benchmark tests. These factors make Phi-2 an exceptional model within the specified tasks, demonstrating that well-planned training and design can yield high performance, challenging the prevailing notion that larger models are inherently superior.

**Llama2 Models:** These models consistently score lower, indicating possible limitations in their training or architectural adaptability. This underperformance highlights that large models do not

| | Relevance | | Coherence | | Diversity | |
|---|---|---|---|---|---|---|
| Dataset | Highest | Lowest | Highest | Lowest | Highest | Lowest |
| Physics | | | | | | |
| - Basic | GPT-4 | Llama2-7b | Phi-2 | Llama2-70b | GPT-4 | Llama2-7b |
| - Intermediate | GPT-4 | Gemini | GPT-4 | Llama2-7b | Phi-2 | Gemini |
| - Advanced | GPT-4 | Llama2-70b | GPT-4 | Llama2-7b | Gemini | Llama2-7b |
| - Wrong | GPT-3.5 | Mistral 8x7b | GPT-3.5 | Llama2-70b | GPT-3.5 | GPT-4 |
| Chemistry | | | | | | |
| - Basic | GPT-4 | Llama2-7b | GPT-4 | Llama2-7b | Mistral 8x7b | Llama2-7b |
| - Intermediate | GPT-4 | Llama2-7b | Mistral 8x7b | Llama2-7b | Mistral 8x7b | Llama2-70b |
| - Advanced | Mistral 8x7b | Llama2-7b | GPT-4 | Llama2-7b | Mistral 8x7b | Llama2-70b |
| - Wrong | Gemini | Llama2-7b | Gemini | Phi-2 | Gemini | Phi-2 |
| Maths | | | | | | |
| - Basic | Phi-2 | Llama2-7b | Phi-2 | Llama2-7b | GPT-4 | GPT-3.5 |
| - Intermediate | Mistral 8x7b | Llama2-7b | GPT-4 | Llama2-7b | GPT-4 | Gemini |
| - Advanced | GPT-4 | Llama2-7b | GPT-4 | Llama2-7b | Mistral 8x7b | Llama2-7b |
| - Wrong | Mistral 8x7b | Llama2-13b | Phi-2 | GPT-3.5 | GPT-4 | Gemini |
| General | GPT-4 | Llama2-7b | GPT-4 | Llama2-7b | GPT-4 | Llama2-7b |

Table 2: Models with the highest and the lowest scores across datasets and expertise levels.

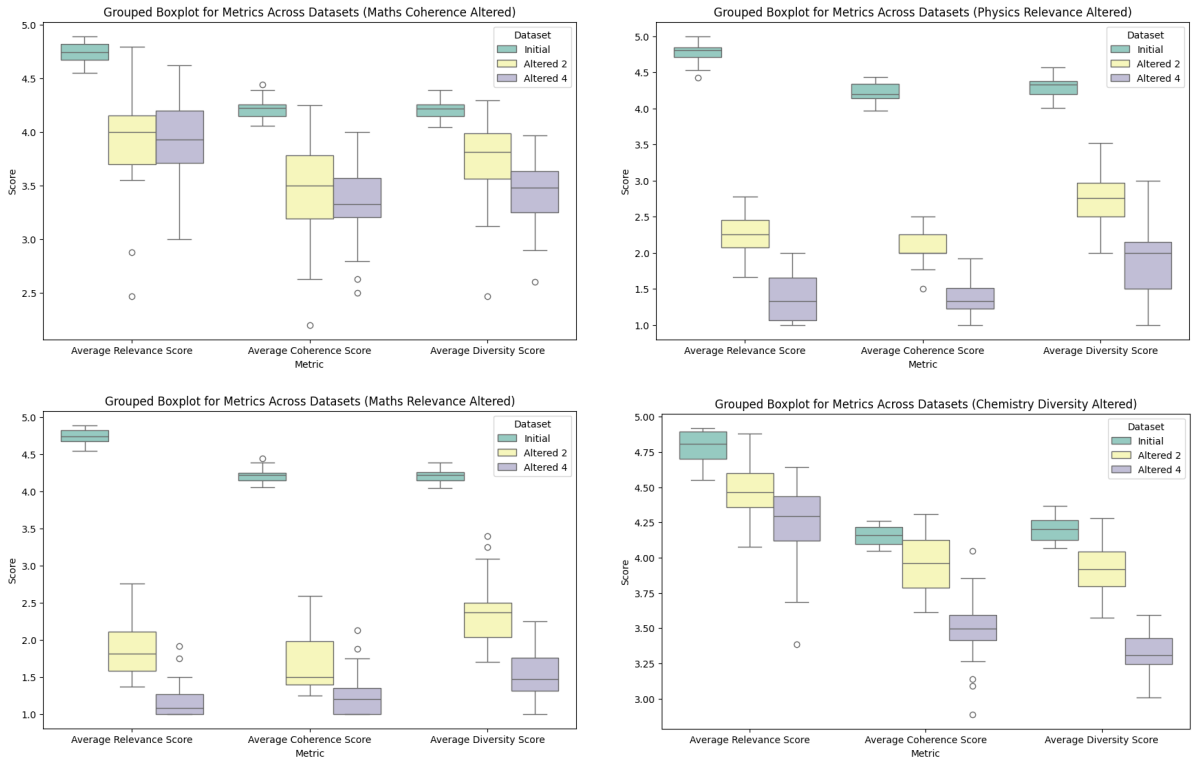


Figure 2: Boxplots showing average relevance, coherence, and diversity scores across altered datasets in Mathematics, Physics, and Chemistry. Each includes "Initial", "2-Altered", and "4-Altered" dataset versions.

universally guarantee superior performance. This could be attributed to their primary optimization for other types of tasks, such as chat and dialogue scenarios, rather than the specific demands of generating novel and contextually deep questions in academic subjects. The models are trained on a new mix of publicly available online data, ensuring a broad knowledge base. However, their performance variability in CDQG tasks suggests that while they have strong general capabilities, they may require further tuning to excel specifically in the academic question generation domain.

**Gemini:** While generally showing lower overall performance, it excels in diversity, perhaps due to its multi-modal training. This suggests that it can generate more varied and creative outputs, which purely text-based models may not achieve.

**Insights and implications** While larger models like GPT-4 generally offered robust overall perfor-

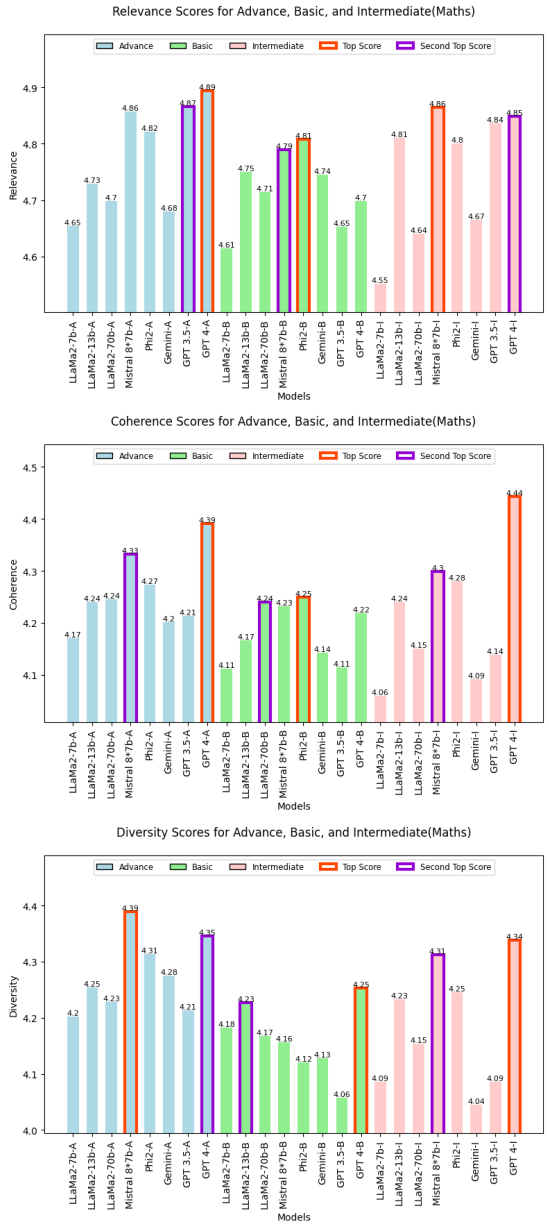


Figure 3: **Metric scores on Maths:** The set of bar charts provides a multidimensional analysis of various models, evaluated by three key performance metrics — Relevance (top), Coherence (middle), and Diversity (bottom). Each chart contrasts the scores across Advanced, Basic, and Intermediate expertise levels for maths, with distinct colors signifying the respective categories. Highlighted bars denote the top and second-highest scoring models within each metric, offering a visual synopsis of comparisons.

mance, smaller or specialized models like Phi-2 and Mistral 8x7b performed exceptionally well. This challenges the conventional notion that bigger is inherently better (Hoffmann et al., 2022), suggesting a nuanced approach to model selection based on specific task requirements.

### 5.2 Questioning the wrong statements

We expect to see the models doubt the credibility of the statements that are intentionally erroneous. While models generally follow the instructions by asking questions, their responses include questioning the credibility of dubious statements with probing questions like "*Are there any exceptions to this rule?*" While all the models do this, but how often they challenge a statement's truth varies. The models like Mistral, LLama 70b, and GPT-4 frequently ask this question in about 250 out of 600 cases the most. In contrast, GPT 3.5 and Llama 7b ask it less often, only about 100 to 150 times the least.

## 6 Ensuring the validity of CDQG

We validate the CDQG evaluation through an ablation study that incrementally add noise, as well as a human validation.

### 6.1 Noise-addition ablation

**Setup** For each entry in the output dataset containing five generated questions, we create two derivative entries by deliberately introducing disturbances. The first variant modifies two questions (**2 Altered**), while the second alters four questions (**4 Altered**). We execute this noise addition using GPT-4 (See Appendix B for the prompt template) and verify that exactly 2 or 4 questions are modified in each respective variant, ensuring the noise addition diminishes question quality. This process yields six new datasets corresponding to each evaluation metric, divided between the two and four modified question scenarios. When we reintroduce these altered datasets to our evaluation process, we expect to observe a decline in scores across all metrics proportional to the added noise. This anticipated degradation aims to demonstrate an inverse correlation between LLM-generated content integrity and noise level. This approach validates our hypothesis that LLMs can effectively differentiate between high-quality (signal) and compromised (noise) data inputs. By showing that introduced inaccuracies result in predictable evaluation score decreases, we employ a logical framework similar to mathematical proof by contradiction. This method demonstrates LLMs' effectiveness in judging relevance, coherence, and diversity.

**Results** The added noise significantly impacts all metrics, showing consistent decline as noise increases from no alterations to **2 Altered** and **4 Altered**, though the magnitude varies across metrics.

The relevance metric exhibits the most pronounced trend, with scores dropping from approximately 4.8 to 2.2 to 1.0. This demonstrates that question precision and topic relevance are most sensitive to noise-induced disruptions. The coherence metric shows a less significant decrease, as alterations to individual questions do not always disrupt the overall logical flow and order. Diversity presents unique challenges, as effectively reducing this metric requires deep subject matter understanding and awareness of topic interconnections. As shown in Figure 2 and Figure 8, while our modifications decrease the diversity score, the reduction is less pronounced than the relevance and coherence, reflecting the complexity of perturbing question diversity while maintaining topical consistency.

Our analysis confirms that noise introduction leads to degradation in LLM performance across all metrics. This validates our hypothesis that LLMs effectively differentiate between high-quality and noise-compromised content, and supports the robustness of our evaluation framework.

### 6.2 Human evaluation

**Setup** The human study starts by selecting a subset of the data. To maintain a manageable workload, we select questions from the first 10 statements of three models (out of eight) for two subject areas (out of four). We include all three variations for each subject area to ensure the generalizability of validation results. The selected 1,320 statements for human evaluation represent approximately 19.6% of our sub-dataset's 6,708 statements. A PhD student manually rated the questions on relevance, coherence, and diversity.

**Results** To analyze the agreement between human and LLM evaluations, we employ Cohen's kappa with linear weights (Doewes et al., 2023). This approach accounts for the ordinal nature of the rating scale and appropriately weights the proximity of agreement on scores, reducing penalties for minor discrepancies between evaluators. The resulting agreement scores demonstrate strong correlations: 0.736 for relevance, 0.698 for coherence, and 0.697 for diversity, indicating robust alignment between LLM and human evaluations.

## 7 Discussion

**Questioning for better LM agents** The ability to raise curiosity-driven questions is crucial for agentic systems that involve knowledge. Current technologies like tree-of-thought (Yao et al., 2024), maieutic prompting (Jung et al., 2022) and Reflexion (Shinn et al., 2023) incorporate functions resembling self-questioning. With improved questioning capabilities, future LM-based agents can better recognize low-quality information and reason about it, eventually being more robust against misinformation. A particularly useful use case for LM agents involves the external memory. Questioning equips the LM agents to inspect and potentially fix the errors within the memory.

**Questioning for scientific discovery** Curiosity-driven questioning has always been a critical step in scientific discovery. Human scientists raise questions along many steps of the endeavor of discovery. Questions like "Why can't an alternative method work here?" and "Why can't an alternative theory explain the data?" are the initial steps toward novel scientific discoveries.

**Questioning in human-machine collaborations** Language models have shown capabilities to elicit human preference (Li et al., 2023). As LMs appear more widely used in chatbots and other human-machine interaction systems, questioning becomes an increasingly important function that improves personalization. Questions can allow the models to clarify the human users' unspoken thoughts and intentions, improving the overall quality of communication (Wadhwa et al., 2024; Wu et al., 2024).

## 8 Conclusion

We propose CDQG and start the exploration for assessing an important capability of LLMs: the potential to seek knowledge driven by curiosity. The CDQG framework includes a task that elicits curiosity-driven questions, a dataset covering statements with varying levels of difficulty and supporting stratified studies, and an LLM-based evaluation setting which is validated by both noise-addition ablation and human evaluations. We find that across various subject domains, LLMs exhibit a strong capability to formulate relevant and coherent questions, underscoring their potential to engage in meaningful inquiry. The automated questioning setting has broad potential applications to improve the performance and usability of knowledge systems.

## 9 Limitations

While this study introduces an innovative framework for evaluating the questioning capabilities of LLMs, it primarily utilizes three metrics: relevance, coherence, and diversity. Though robust, these metrics may not capture the full depth of human-like questioning, such as emotional intelligence, knowledge acquisition, factual reasoning, etc. Future research could explore the integration of metrics that assess these human-centric qualities to better mimic real-world applications. Additionally, the evaluations are performed within a controlled academic setting, which might not fully reflect the complexities of natural environments where LLMs typically operate. Extending the evaluation to more dynamic settings or incorporating unstructured, real-world conversation data could enhance the applicability of the findings. Moreover, while our noise addition ablation study strengthens the assessment of model robustness, exploring more varied disturbances could provide a richer understanding of how LLMs perform under realistic and unpredictable conditions.

## References


Zahra Abbasiantaeb, Yifei Yuan, Evangelos Kanoulas, and Mohammad Aliannejadi. 2024. Let the LLMs Talk: Simulating Human-to-Human Conversational QA via Zero-Shot LLM-to-LLM Interactions. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pages 8–17, Merida Mexico. ACM.

Selcuk Acar, Kelly Berthiaume, and Rebecca Johnson. 2023. What kind of questions do creative people ask? *Journal of Creativity*, 33(3):100062.

Nischal Ashok Kumar, Nigel Fernandez, Zichao Wang, and Andrew Lan. 2023. Improving Reading Comprehension Question Generation with Data Augmentation and Overgenerate-and-rank. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 247–259, Toronto, Canada. Association for Computational Linguistics.

Pavan Baswani, Ananya Mukherjee, and Manish Shrivastava. 2023. LTRC_iiith's 2023 Submission for Prompting Large Language Models as Explainable Metrics Task. In *Proceedings of the 4th Workshop on Evaluation and Comparison of NLP Systems*, pages 156–163, Bali, Indonesia. Association for Computational Linguistics.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

Yi Chen, Rui Wang, Haiyun Jiang, Shuming Shi, and Ruifeng Xu. 2023. Exploring the Use of Large Language Models for Reference-Free Text Quality Evaluation: An Empirical Study. In *Findings of the Association for Computational Linguistics: IJCNLP-AACL 2023 (Findings)*, pages 361–374, Nusa Dua, Bali. Association for Computational Linguistics.

Cheng-Han Chiang and Hung-yi Lee. 2023. Can Large Language Models Be an Alternative to Human Evaluations? *arXiv preprint*. ArXiv:2305.01937 [cs].

Afrizal Doewes, Nughthoh Kurdhi, and Akrati Saxena. 2023. Evaluating quadratic weighted kappa as the standard performance metric for automated essay scoring. In *Proceedings of the 16th International Conference on Educational Data Mining*, pages 103–113. International Educational Data Mining Society (IEDMS). 16th International Conference on Educational Data Mining, EDM 2023, EDM 2023 ; Conference date: 11-07-2023 Through 14-07-2023.

Sabina Elkins, Ekaterina Kochmar, Iulian Serban, and Jackie C. K. Cheung. 2023. How Useful Are Educational Questions Generated by Large Language Models? In *Artificial Intelligence in Education. Posters and Late Breaking Results, Workshops and Tutorials, Industry and Innovation Tracks, Practitioners, Doctoral Consortium and Blue Sky*, pages 536–542, Cham. Springer Nature Switzerland.

Guglielmo Faggioli, Laura Dietz, Charles L. A. Clarke, Gianluca Demartini, Matthias Hagen, Claudia Hauff, Noriko Kando, Evangelos Kanoulas, Martin Potthast, Benno Stein, and Henning Wachsmuth. 2023. Perspectives on Large Language Models for Relevance Judgment. In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval*, pages 39–50, Taipei Taiwan. ACM.

Javier Ferrando, Oscar Obeso, Senthooran Rajamanoharan, and Neel Nanda. 2024. Do I Know This Entity? Knowledge Awareness and Hallucinations in Language Models. *arXiv preprint*.

Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. GPTScore: Evaluate as You Desire. *arXiv preprint*. ArXiv:2302.04166 [cs].

Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. Retrieval-augmented generation for large language models: A survey. *Preprint*, arXiv:2312.10997.

Gemini Team, Rohan Anil, and et al. 2024. Gemini: A Family of Highly Capable Multimodal Models. *arXiv preprint*. ArXiv:2312.11805 [cs].

Alireza Ghafarollahi and Markus J. Buehler. 2024. Sciagents: Automating scientific discovery through multi-agent intelligent graph reasoning. *Preprint*, arXiv:2409.05556.

Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. Chatgpt outperforms crowd workers for text-annotation tasks.

Shibo Hao, Yi Gu, Haodi Ma, Joshua Hong, Zhen Wang, Daisy Wang, and Zhiting Hu. 2023. Reasoning with Language Model is Planning with World Model. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8154–8173, Singapore. Association for Computational Linguistics.

Qianyu He, Jie Zeng, Wenhao Huang, Lina Chen, Jin Xiao, Qianxi He, Xunzhe Zhou, Jiaqing Liang, and Yanghua Xiao. 2024. Can Large Language Models Understand Real-World Complex Instructions? *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(16):18188–18196.

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.

Fan Huang, Haewoon Kwak, Kunwoo Park, and Jisun An. 2024. ChatGPT rates natural language explanation quality like humans: But on which scales? In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3111–3132, Torino, Italia. ELRA and ICCL.

Yizheng Huang and Jimmy Huang. 2024. A survey on retrieval-augmented text generation for large language models. *Preprint*, arXiv:2404.10981.

Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, Lélio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Théophile Gervet, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2024. Mixtral of Experts. *arXiv preprint*. ArXiv:2401.04088 [cs].

Dongfu Jiang, Yishan Li, Ge Zhang, Wenhao Huang, Bill Yuchen Lin, and Wenhu Chen. 2023. TIGERScore: Towards Building Explainable Metric for All Text Generation Tasks. *arXiv preprint*. ArXiv:2310.00752 [cs].

Jaehun Jung, Lianhui Qin, Sean Welleck, Faeze Brahman, Chandra Bhagavatula, Ronan Le Bras, and Yejin Choi. 2022. Maieutic prompting: Logically consistent reasoning with recursive explanations. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1266–1279, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Moussa Kamal Eddine, Guokan Shang, Antoine Tixier, and Michalis Vazirgiannis. 2022. FrugalScore: Learning Cheaper, Lighter and Faster Evaluation Metrics for Automatic Text Generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1305–1318, Dublin, Ireland. Association for Computational Linguistics.

Nan Rosemary Ke, Danny P. Sawyer, Hubert Soyer, Martin Engelcke, David P Reichert, Drew A. Hudson, John Reid, Alexander Lerchner, Danilo Jimenez Rezende, Timothy P Lillicrap, Michael Mozer, and Jane X Wang. 2024. Can foundation models actively gather information in interactive environments to test hypotheses? *Preprint*, arXiv:2412.06438.

Tom Kocmi and Christian Federmann. 2023. Large Language Models Are State-of-the-Art Evaluators of Translation Quality. In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 193–203, Tampere, Finland. European Association for Machine Translation.

Neema Kotonya, Saran Krishnasamy, Joel Tetreault, and Alejandro Jaimes. 2023. Little Giants: Exploring the Potential of Small LLMs as Evaluation Metrics in Summarization in the Eval4NLP 2023 Shared Task. In *Proceedings of the 4th Workshop on Evaluation and Comparison of NLP Systems*, pages 202–218, Bali, Indonesia. Association for Computational Linguistics.

Alexander Kotov and ChengXiang Zhai. 2010. Towards natural question guided search. In *Proceedings of the 19th international conference on World wide web*, pages 541–550, Raleigh North Carolina USA. ACM.

Satyapriya Krishna, Kalpesh Krishna, Anhad Mohananey, Steven Schwarcz, Adam Stambler, Shyam Upadhyay, and Manaal Faruqui. 2024. Fact, fetch, and reason: A unified evaluation of retrieval-augmented generation. *Preprint*, arXiv:2409.12941.

Ghader Kurdi, Jared Leo, Bijan Parsia, Uli Sattler, and Salam Al-Emari. 2020. A Systematic Review of Automatic Question Generation for Educational Purposes. *International Journal of Artificial Intelligence in Education*, 30(1):121–204.

Yukyung Lee, Joonghoon Kim, Jaehee Kim, Hyowon Cho, and Pilsung Kang. 2024. CheckEval: Robust Evaluation Framework using Large Language Model via Checklist. *arXiv preprint*. ArXiv:2403.18771 [cs].

Christoph Leiter, Piyawat Lertvittayakumjorn, Marina Fomicheva, Wei Zhao, Yang Gao, and Steffen Eger. 2024. Towards Explainable Evaluation Metrics for Machine Translation. *Journal of Machine Learning Research*, 25(75):1–49.

Christoph Leiter, Juri Opitz, Daniel Deutsch, Yang Gao, Rotem Dror, and Steffen Eger. 2023. The Eval4NLP 2023 Shared Task on Prompting Large Language Models as Explainable Metrics. In *Proceedings of the 4th Workshop on Evaluation and Comparison of NLP Systems*, pages 117–138, Bali, Indonesia. Association for Computational Linguistics.

Belinda Z. Li, Alex Tamkin, Noah Goodman, and Jacob Andreas. 2023. Eliciting human preferences with language models. *Preprint*, arXiv:2310.11589.

Dawei Li, Bohan Jiang, Liangjie Huang, Alimohammad Beigi, Chengshuai Zhao, Zhen Tan, Amrita Bhattacharjee, Yuxuan Jiang, Canyu Chen, Tianhao Wu, Kai Shu, Lu Cheng, and Huan Liu. 2024. From generation to judgment: Opportunities and challenges of llm-as-a-judge. *Preprint*, arXiv:2411.16594.

Jiwei Li, Alexander H. Miller, Sumit Chopra, Marc'Aurelio Ranzato, and Jason Weston. 2017. Learning through Dialogue Interactions by Asking Questions. *arXiv preprint*. ArXiv:1612.04936 [cs].

Kevin Liu, Stephen Casper, Dylan Hadfield-Menell, and Jacob Andreas. 2024a. Cognitive Dissonance: Why Do Language Model Outputs Disagree with Internal Representations of Truthfulness? In *EMNLP*.

Minqian Liu, Ying Shen, Zhiyang Xu, Yixin Cao, Eunah Cho, Vaibhav Kumar, Reza Ghanadan, and Lifu Huang. 2024b. X-Eval: Generalizable Multi-aspect Text Evaluation via Augmented Instruction Tuning with Auxiliary Evaluation Aspects. *arXiv preprint*. ArXiv:2311.08788 [cs].

Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023a. G-Eval: NLG Evaluation using Gpt-4 with Better Human Alignment. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.

Yuxuan Liu, Tianchi Yang, Shaohan Huang, Zihan Zhang, Haizhen Huang, Furu Wei, Weiwei Deng, Feng Sun, and Qi Zhang. 2023b. Calibrating LLM-Based Evaluator. *arXiv preprint*. ArXiv:2309.13308 [cs].

Adian Liusie, Potsawee Manakul, and Mark Gales. 2024. LLM Comparative Assessment: Zero-shot NLG Evaluation through Pairwise Comparisons using Large Language Models. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 139–151, St. Julian's, Malta. Association for Computational Linguistics.

Evan Lucas, Kelly S. Steelman, Leo C. Ureel, and Charles Wallace. 2024. For those who don't know (how) to ask: Building a dataset of technology questions for digital newcomers. *arXiv preprint*. ArXiv:2403.18125 [cs].

Tula Masterman, Sandi Besen, Mason Sawtell, and Alex Chao. 2024. The Landscape of Emerging AI Agent Architectures for Reasoning, Planning, and Tool Calling: A Survey. *arXiv preprint*. ArXiv:2404.11584 [cs].

Sébastien Bubeck Mojan Javaheripi. 2023. Phi-2: The surprising power of small language models.

OpenAI. 2024. Learning to reason with large language models. `https://openai.com/index/learning-to-reason-with-llms/`. Accessed: 2024-09-13.

OpenAI, Josh Achiam, Steven Adler, and et al. 2024. GPT-4 Technical Report. *arXiv preprint*. ArXiv:2303.08774 [cs].

Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language Models as Knowledge Bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.

Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. *Preprint*, arXiv:2303.11366.

Chenglei Si, Diyi Yang, and Tatsunori Hashimoto. 2024. Can llms generate novel research ideas? a large-scale human study with 100+ nlp researchers. *Preprint*, arXiv:2409.04109.

Mirac Suzgun, Tayfun Gur, Federico Bianchi, Daniel E. Ho, Thomas Icard, Dan Jurafsky, and James Zou. 2024. Belief in the Machine: Investigating Epistemological Blind Spots of Language Models. *arXiv preprint*.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross

Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. *arXiv preprint*. ArXiv:2307.09288 [cs].

Karthik Valmeekam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2023. PlanBench: An Extensible Benchmark for Evaluating Large Language Models on Planning and Reasoning about Change. *Advances in Neural Information Processing Systems*, 36:38975–38987.

Pat Verga, Sebastian Hofstatter, Sophia Althammer, Yixuan Su, Aleksandra Piktus, Arkady Arkhangorodsky, Minjie Xu, Naomi White, and Patrick Lewis. 2024. Replacing Judges with Juries: Evaluating LLM Generations with a Panel of Diverse Models. *arXiv preprint*. ArXiv:2404.18796 [cs].

Manya Wadhwa, Xinyu Zhao, Junyi Jessy Li, and Greg Durrett. 2024. Learning to refine with fine-grained natural language feedback. *Preprint*, arXiv:2407.02397.

Haonan Wang, James Zou, Michael Mozer, Anirudh Goyal, Alex Lamb, Linjun Zhang, Weijie J. Su, Zhun Deng, Michael Qizhe Xie, Hannah Brown, and Kenji Kawaguchi. 2024a. Can AI Be as Creative as Humans? *arXiv preprint*. ArXiv:2401.01623 [cs].

Jiaan Wang, Yunlong Liang, Fandong Meng, Zengkui Sun, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. 2023. Is ChatGPT a Good NLG Evaluator? A Preliminary Study. In *Proceedings of the 4th New Frontiers in Summarization Workshop*, pages 1–11, Hybrid. Association for Computational Linguistics.

Wenxuan Wang, Juluan Shi, Chaozheng Wang, Cheryl Lee, Youliang Yuan, Jen tse Huang, and Michael R. Lyu. 2024b. Learning to Ask: When LLMs Meet Unclear Instruction. *Preprint*, arXiv:2409.00557.

Yating Wu, Ritika Mangla, Alexandros G. Dimakis, Greg Durrett, and Junyi Jessy Li. 2024. Which questions should i answer? salience prediction of inquisitive questions. *Preprint*, arXiv:2404.10917.

Siheng Xiong, Ali Payani, Ramana Kompella, and Faramarz Fekri. 2024. Large Language Models Can Learn Temporal Reasoning. *arXiv preprint*. ArXiv:2401.06853 [cs].

Wenda Xu, Danqing Wang, Liangming Pan, Zhenqiao Song, Markus Freitag, William Wang, and Lei Li. 2023. INSTRUCTSCORE: Towards Explainable Text Generation Evaluation with Automatic Feedback. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5967–5994, Singapore. Association for Computational Linguistics.

Bingsheng Yao, Dakuo Wang, Tongshuang Wu, Zheng Zhang, Toby Li, Mo Yu, and Ying Xu. 2022. It is AI's Turn to Ask Humans a Question: Question-Answer Pair Generation for Children's Story Books. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 731–744, Dublin, Ireland. Association for Computational Linguistics.

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2024. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems*, 36.

Eric Zelikman, Georges Harik, Yijia Shao, Varuna Jayasiri, Nick Haber, and Noah D. Goodman. 2024. Quiet-STaR: Language Models Can Teach Themselves to Think Before Speaking. *arXiv preprint*. ArXiv:2403.09629 [cs].

Wei Zhao, Michael Strube, and Steffen Eger. 2023. DiscoScore: Evaluating Text Generation with BERT and Discourse Coherence. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3865–3883, Dubrovnik, Croatia. Association for Computational Linguistics.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. *Advances in Neural Information Processing Systems*, 36:46595–46623.

Ming Zhong, Yang Liu, Da Yin, Yuning Mao, Yizhu Jiao, Pengfei Liu, Chenguang Zhu, Heng Ji, and Jiawei Han. 2022. Towards a Unified Multi-Dimensional Evaluator for Text Generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2023–2038, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Yuqi Zhu, Xiaohan Wang, Jing Chen, Shuofei Qiao, Yixin Ou, Yunzhi Yao, Shumin Deng, Huajun Chen, and Ningyu Zhang. 2024. LLMs for Knowledge Graph Construction and Reasoning: Recent Capabilities and Future Opportunities. *arXiv preprint*.

# A Appendix

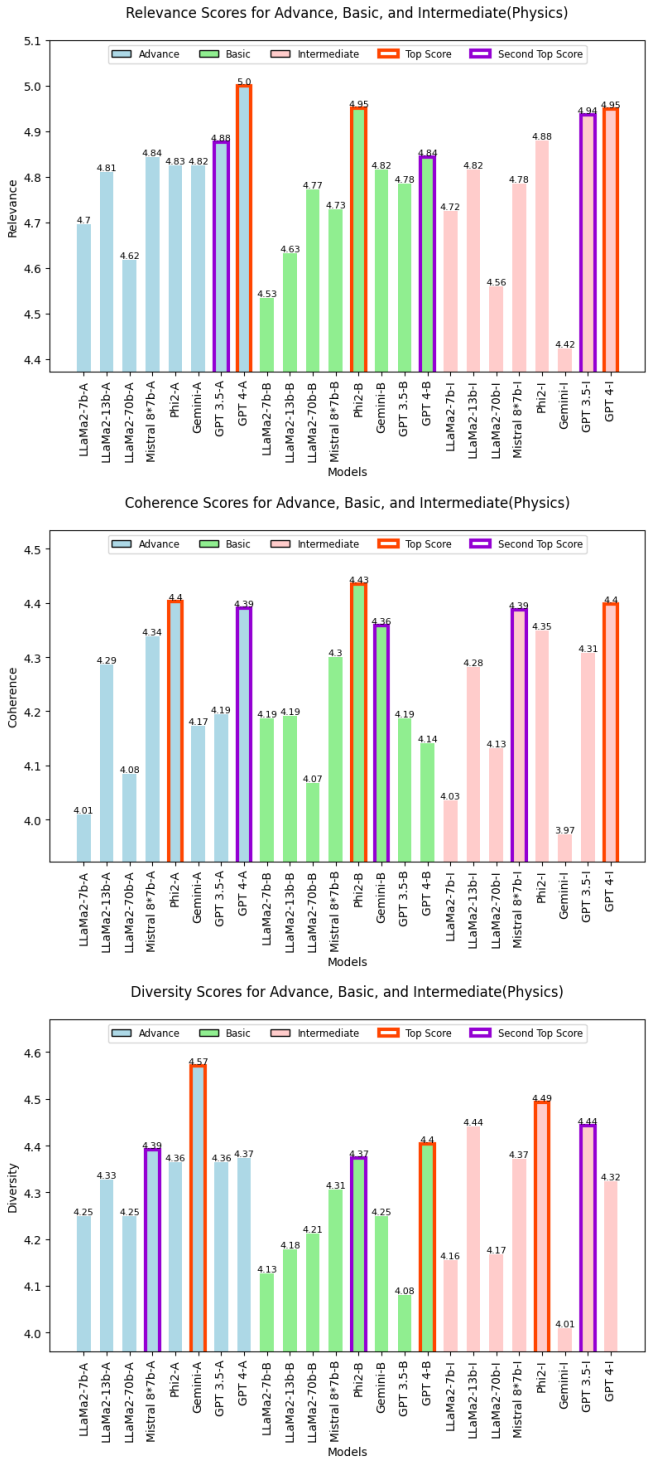


Figure 4: **Metric scores on Physics:** The set of bar charts provides a multidimensional analysis of various models, evaluated by three key performance metrics — Relevance, Coherence , and Diversity. Each chart contrasts the scores across Advanced , Basic , and Intermediate expertise levels for maths, with distinct colors signifying the respective categories. Highlighted bars denote the top and second-highest scoring models within each metric, offering a visual synopsis of comparisons.

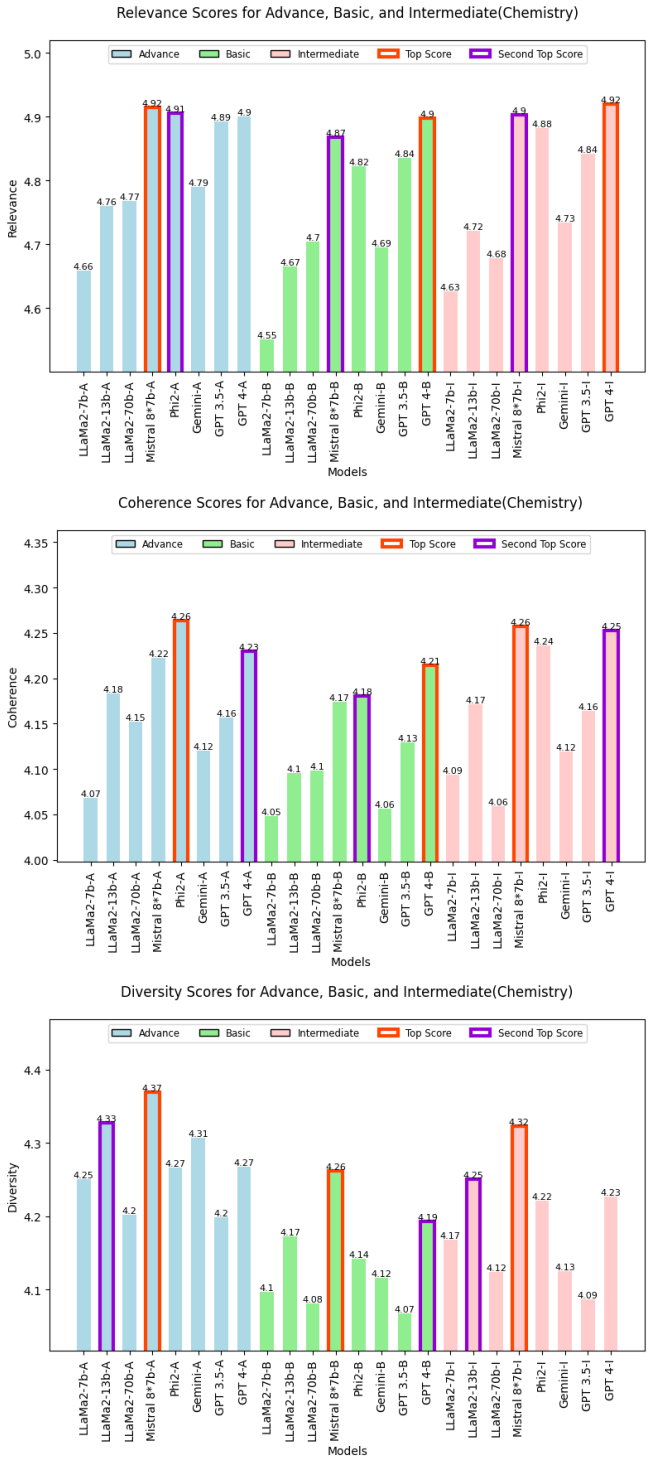


Figure 5: **Metric scores on Chemistry:** The set of bar charts provides a multidimensional analysis of various models, evaluated by three key performance metrics — Relevance, Coherence , and Diversity. Each chart contrasts the scores across Advanced , Basic , and Intermediate expertise levels for Chemistry, with distinct colors signifying the respective categories. Highlighted bars denote the top and second-highest scoring models within each metric, offering a visual synopsis of comparisons.

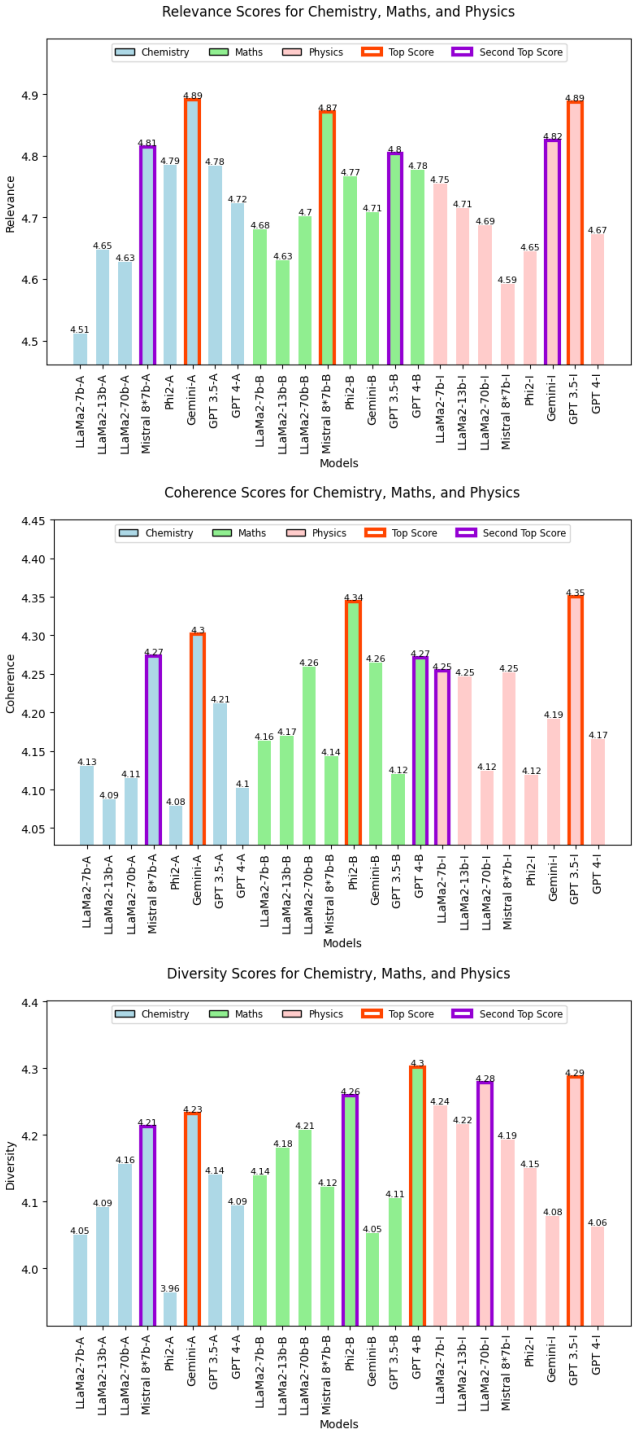


Figure 6: **Metric scores on Wrong Statements:** The set of bar charts provides a multidimensional analysis of various models, evaluated by three key performance metrics — Relevance, Coherence , and Diversity. Each chart contrasts the scores across Chemistry , Maths , and Physics subjects , with distinct colors signifying the respective categories. Highlighted bars denote the top and second-highest scoring models within each metric, offering a visual synopsis of comparisons.

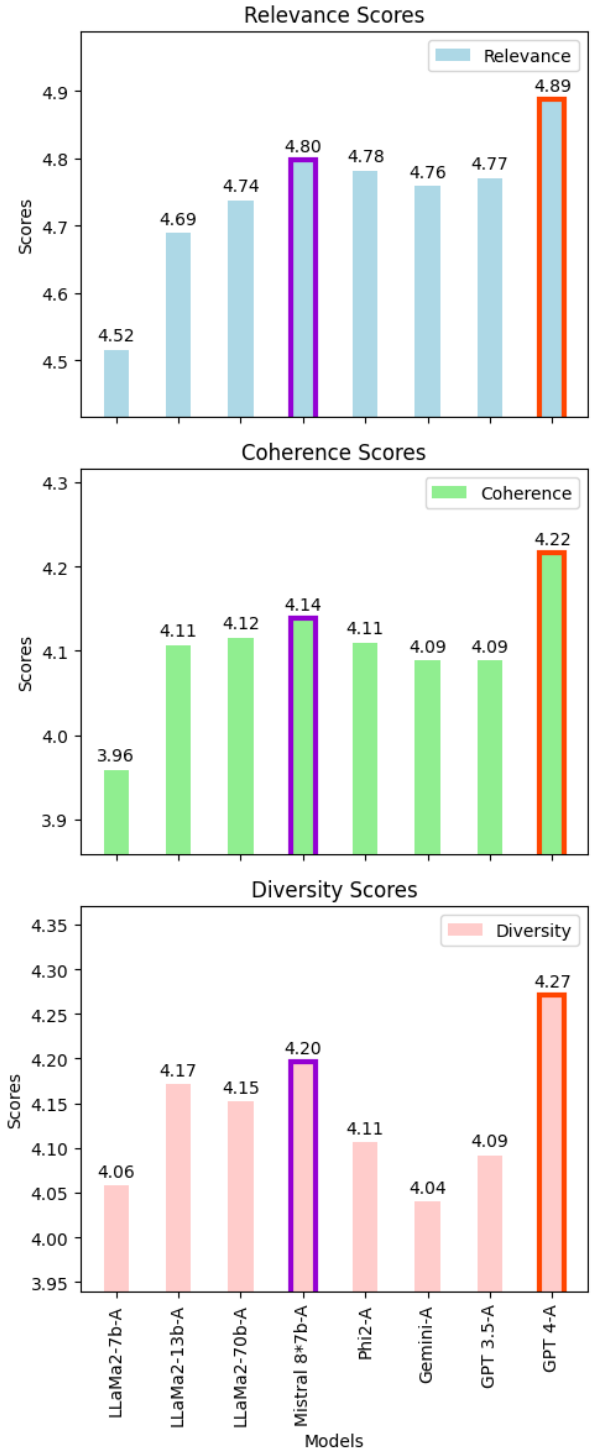


Figure 7: **Metric scores on General Statements:** The set of bar charts provides a multidimensional analysis of various models, evaluated by three key performance metrics — Relevance, Coherence , and Diversity. Each chart contrasts the scores across Relevance , Coherence , and Diversity . Highlighted bars denote the top and second-highest scoring models within each metric, offering a visual synopsis of comparisons.

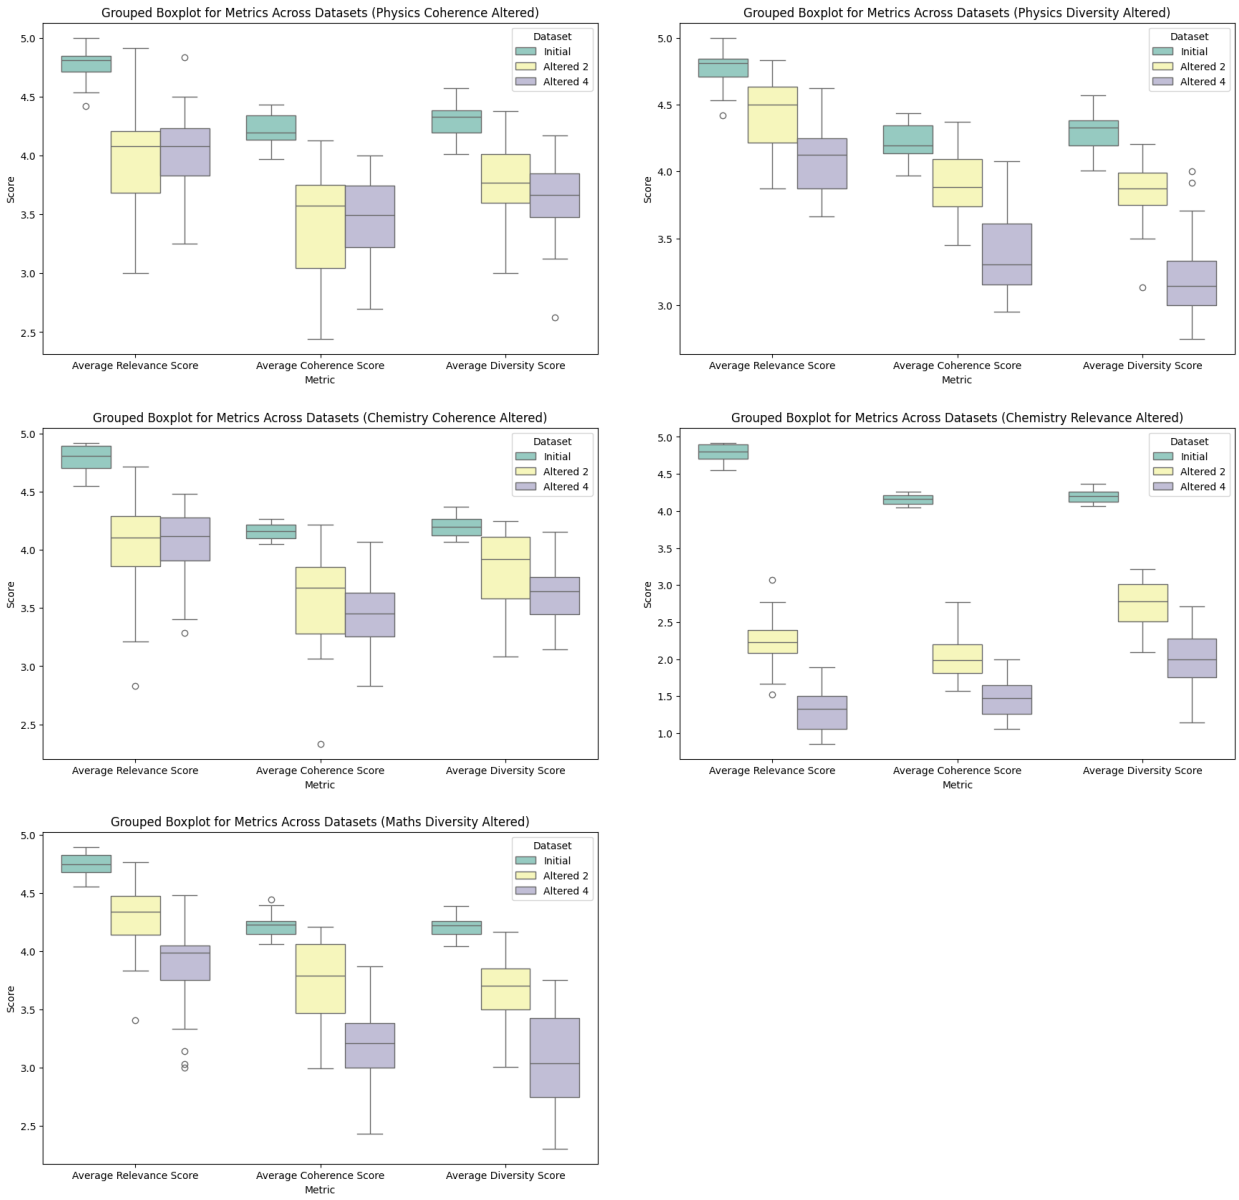


Figure 8: Boxplots showing average relevance, coherence, and diversity scores across altered datasets in Physics (first row), Chemistry (second row), and Maths (third row). Each includes 'Initial', 'Altered 2', and 'Altered 4' dataset versions.

## B List of prompt templates

**Prompt 1: Curiosity-Driven Question Generation**

Imagine you are a human encountering this *{subject}* for the first time: "*{scenario}*". List the top 5 questions that would come to your mind, useful for learning about it as you are new to it. Provide your questions in a simple bullet point list.

**Prompt 3: Combining Scoring and Justification using Gemini**

Initial Query: {instruction}
Answer Given by LLM: {llm_answer}
Scores by humans: Human 1: {human_score_1}, Human 2: {human_score_2}, Human 3: {human_score_3}.

These are three scorings by a human and the justifications. Now, consider all the scorings and their justifications and give final scores for relevance, coherence, and diversity. Don't just take the average of scores or support one scorer; instead, read the justifications and, accordingly, give a final score and justify. Provide output in JSON format.

**Prompt 2: Evaluation Task**

Below are sets of 5 questions generated by different Language Models (LLMs) in response to a specific statement or scenario they were presented with for the first time. Your task is to evaluate these questions based on the following three metrics: Coherence, Relevance, and Diversity. Each set of questions is aimed at uncovering and understanding the elements and concepts within the given statement.

**Criteria for each metric:**

- **Relevance:** Assess how directly each question pertains to the specific details, elements, or concepts presented in the statement or scenario. Questions should aim to clarify, expand upon, or directly explore the content of the statement, focusing on the immediate context rather than peripheral or advanced topics not directly introduced by the statement.
- **Coherence:** Evaluate how logically the questions within each set connect to one another and whether they form a coherent line of inquiry that would logically progress a beginner's understanding of the topic. Consider if the sequence of questions or their thematic connection facilitates a structured exploration of the statement.
- **Diversity:** Determine the range of aspects covered by the questions in relation to the statement, ensuring that each question brings a new dimension or perspective to understanding the statement. While maintaining direct relevance, the questions should collectively offer a broad exploration of the topic, including but not limited to definitions, implications, applications, or theoretical underpinnings.

For each set of questions, provide a score from 1 to 5 for each metric, where 1 indicates that the questions poorly meet the criteria and 5 indicates excellent adherence to the criteria. Additionally, provide brief justifications for your scores, highlighting strengths and areas for improvement in relation to the three metrics.

Your evaluation will help determine which LLM produced the most effective set of questions for fostering an understanding of the given statement or scenario, balancing direct relevance to the statement, logical coherence in inquiry, and diversity in exploration.

**Input for LLM:** {*instruction*}
**LLM Output:** {*model_output*}

**Prompt 4: Alteration Prompt**
Initial Query to random LLM: {instruction} and the Output given by that LLM: {model_output},
Given a set of questions related to a specific statement provided by an LLM, modify exactly 4 questions for each metric to intentionally introduce noise. The objective is to decrease the values of three specified metrics: relevance, coherence, and diversity, in relation to the original statement.

**For Relevance:** Alter 4 random questions to make them less directly connected to the main topic of the statement. The goal is to subtly shift focus without completely diverging into unrelated topics.

**For Coherence:** Revise the sequence or content of 4 random questions to break the logical flow of inquiry. Adjustments should make the progression less structured and more challenging to follow, thus impacting the coherence of the set.

**For Diversity:** Change or add 4 random questions to concentrate more narrowly on similar aspects or repeat themes. This reduces the range of explored topics, affecting the overall diversity of the question set.

After making these modifications, specify the number of questions you altered for each metric and provide the altered list of questions. Your output should demonstrate the impact of introduced noise on the measurement of each metric.

**Required Output Format**

Your response should be structured in JSON format, comprising three sections corresponding to the metrics: Relevance, Coherence, and Diversity. Each section must detail the number of questions modified ('changed') and include the revised list of questions after changes ('questions'). Avoid including explanations or content beyond this structured format.

## C Model Configuration Details

**Gemini Settings:** The Gemini model was configured with a low temperature setting of 0.1 to ensure predictable and consistent outputs. The top_p and top_k parameters were both set to 1, constraining the model to the most likely outcomes. The maximum output tokens were limited to 400 to balance detail with computational efficiency. Safety settings were established to minimize the risk of generating harmful content, with no blocks applied across categories such as harassment, hate speech, sexually explicit content, and dangerous content.

**Mistral Model Setup:** The Mistral model utilized a tokenizer and model settings specifically tailored for instruction-based tasks. This setup included using the AutoTokenizer and AutoModelForCausalLM from a pretrained snapshot, equipped with BitsAndBytesConfig for efficient quantization. The configuration ensured operations were optimized for 4-bit quantization and the compute dtype set to float16, enhancing the model's performance while reducing memory usage. The text-generation pipeline was adjusted with a temperature of 0.1 and a repetition penalty of 1.1 to generate more coherent and less repetitive text, with a limit of 128 new tokens per generation instance.

**Llama Model Configurations:** For the Llama models, including, Llama 7b, Llama 13b and Llama 70b, configurations were similarly tailored to enhance performance and efficiency. Both models used quantization settings conducive to low-memory consumption while maintaining computational precision. These settings were crucial for managing the large parameter size inherent to these models. Each model's generation pipeline was configured to produce full-text outputs with controlled temperature settings and repetition penalties to ensure relevance and diversity in the generated text.

**Phi2 Model Configuration:** The Phi2 model from Microsoft was set up with advanced quantization techniques to support efficient processing. The model and tokenizer were loaded from a specific snapshot with settings that enabled high-performance text generation. The generation settings included a controlled temperature for predictability, a sampling strategy to introduce variety, and a repetition penalty to avoid redundant content, making it well-suited for generating diverse and engaging text.

**Compute Resources:** For models accessed via API, computations were performed using CPU resources. In contrast, models retrieved from HuggingFace were run on a single NVIDIA GPU setup equipped with 48GB of RAM. Notably, all models utilized in this study were quantized versions, optimizing computational efficiency and resource usage.