

LFQA-HP-1M: A Large-Scale Human Preference Dataset for Long-Form Question Answering

Rafid Ishrak Jahan Fahmid Shahriar Iqbal Sagnik Ray Choudhury

University of North Texas, Denton, Texas, USA

{RafidIshrakJahan, FahmidShahriarIqbal}@my.unt.edu

Sagnik.Raychoudhury@unt.edu

Abstract

Long-form question answering (LFQA) demands nuanced evaluation of multi-sentence explanatory responses, yet existing metrics often fail to reflect human judgment. We present LFQA-HP-1M, a large-scale dataset comprising 1.3M human pairwise preference annotations for LFQA. We propose nine rubrics for answer quality evaluation, and show that simple linear models based on these features perform comparably to state-of-the-art LLM evaluators. We further examine transitivity consistency, positional bias, and verbosity biases in LLM evaluators and demonstrate their vulnerability to adversarial perturbations. Overall, this work provides one of the largest public LFQA preference datasets and a rubric-driven framework for transparent and reliable evaluation.

Keywords: Long-Form Question Answering (LFQA); Human Preference Modeling; LLM-as-a-judge

1. Introduction

A key strength of modern LLMs is their ability to generate detailed, multi-sentence answers to complex, real-world questions. This field of Long Form Question Answering (LFQA) has advanced substantially over the past years, accompanied by the creation of numerous datasets to evaluate model performance. These include resources where annotators explicitly *write* questions and answers tailored to particular tasks (ShARC (Saeidi et al., 2018), NarrativeQA (Kočiský et al., 2018), QASPER (Dasigi et al., 2021)), as well as datasets assembled from Q&A content hosted on public platforms like Stack-Overflow or Reddit (ELI-5 (Fan et al., 2019)). Evaluation on these benchmarks generally relies on reference-based metrics that assess the quality of generated answers by comparing system outputs with one or more human-written reference responses. Such evaluation methods include: a) lexical overlap metrics such as ROUGE (Lin, 2004) and BLEU (Papineni et al., 2002) that are insensitive to paraphrasing and deeper semantic equivalence; and b) approaches based on contextual token embeddings such as BERTScore (Zhang et al., 2020), as well as learned scoring models like BLEURT (Sellam et al., 2020) and BARTScore (Yuan et al., 2021) that estimate semantic similarity. However, these reference-based metrics perform poorly under domain adaptation as they rely on fixed, human-written references and general-domain semantics.

An increasing number of studies focus on *reference-free* evaluation, where proposed methods or models assign direct numerical scores to generated text, without comparing it with reference text. This is commonly known as LLMs-as-judge (Li et al., 2024) when the evaluator models are LLMs. Examples include UniEval (Zhong et al., 2022),

GPTScore (Fu et al., 2024), and the Prometheus family of models (Kim et al., 2024). In many applications, especially LFQA, it is more practical for a model-based evaluator to produce pairwise preference judgments between two candidate answers, as this is generally less demanding. A robust framework for such comparative evaluation helps users make more informed and reliable choices when selecting LLMs for their tasks.

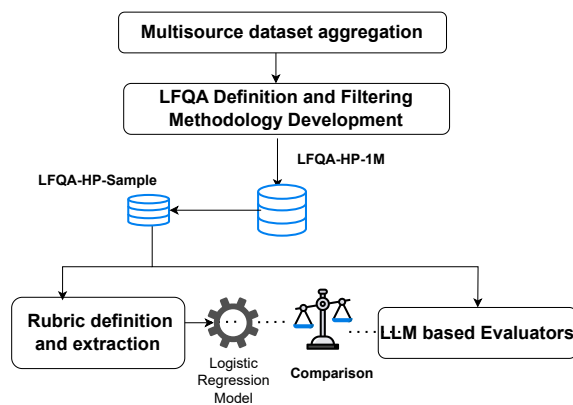


Figure 1: Workflow for our framework: from multi-source dataset aggregation and LFQA filtering to rubric-based evaluation and model comparison.

We develop a framework¹ (Fig. 1) to advance the field of comparative reference-free evaluation of LFQA. Our contributions are as follows:

- **LFQA-HP-1M: A large-scale preference judgment dataset² for LFQA.** We construct a large-scale LFQA dataset comprising 1.3 million human

¹ <https://github.com/nlpatunt/lfqa-eval-lrec-2026>

² <https://huggingface.co/datasets/nlpatunt/LFQA-HP-1M>

preference annotations, where each question is paired with a) two candidate answers and b) a human judgment indicating which answer is preferred. The dataset is curated from existing large-scale preference judgment datasets available on the web: SHP-2 (Ethayarajh et al., 2022), LFQA Eval (Xu et al., 2023), and ChatArena (Zheng et al., 2023). These datasets contain human annotations comparing the quality of multiple long responses to a given prompt. However, not all of these prompts can be classified as questions. To address this, we first develop a rigorous definition of what constitutes a long-form question and operationalize it to construct a high-precision prompt-based filtering pipeline that systematically removes non-long-form question (non-LFQ) utterances (§3).

- **Rubric development for answer quality evaluation.** We propose rubrics for long-form answer evaluation to address the limitations of traditional lexical overlap and semantic metrics. These rubrics cover nine fine-grained aspects, including completeness, accuracy, grammar, and factuality (§4).
- **Model evaluation & discussions.** We benchmark state-of-the-art closed-source LLMs’ ability to mimic human preference judgment on a carefully sampled subset of our data. We compare that with interpretable linear models that use the rubrics as features. We also include discussions about multidimensional adversarial robustness for the LLMs. Specifically, we test for models’ a) transitive consistency and propensity toward verbosity and position bias in LFQA preference evaluation, and b) robustness for semantics preserving adversarial perturbations to the answers. At each stage of the framework, multiple challenges are systematically addressed to maintain an optimal balance among model performance, computational efficiency, and overall cost (§5).

Our framework unifies large-scale human preference data, interpretable rubric-based scoring, and systematic robustness analyses, enabling transparent, reproducible, and theoretically grounded reference-free evaluation in LFQA. We publicly release the LFQA-HP-1M dataset and all of our code to facilitate future work into LFQA evaluation.

2. Related Work

LFQA progress has been hindered by limited dataset availability. Fan et al. (2019) introduced ELI5, a 270K-pair dataset of open-ended “wh” and “how” questions from Reddit’s ExplainLikeIamFive community, establishing LFQA as a distinct research area. In parallel, Nakano et al. (2022) developed WebGPT, a dataset of factually grounded

long answers generated via web search and reinforcement learning from human feedback, but its questions remained largely ELI5-style, limiting domain diversity.

Beyond dataset scarcity, prior LFQA studies have offered inconsistent definitions of the task, which our work addresses. Fan et al. (2019) described LFQA as producing elaborate, in-depth answers to open-ended questions, while Krishna et al. (2021) defined it as retrieving and generating paragraph-length responses. Su et al. (2022) similarly emphasized generating detailed explanations, whereas Stelmakh et al. (2022) viewed LFQA as answering questions that require such explanations. Xu et al. (2023) highlighted the generation of long, complex answers using LLMs, and Sachdeva et al. (2025) framed LFQA as requiring thorough responses to complex queries. These variations reveal the lack of a unified definition, complicating automated LFQA identification.

In addition to limited datasets and inconsistent definitions, LFQA evaluation remains challenging. Rosset et al. (2021) offered one of the first formalizations of human preference judgment in this setting, emphasizing mathematical consistency (e.g., transitivity, monotonicity) in pairwise comparisons. However, their framework prioritized abstract consistency over rubric-level interpretability, offering little guidance for selecting one answer over another. Other early studies relied on lexical metrics such as ROUGE, BERTScore (Zhang et al., 2020), and BARTScore (Yuan et al., 2021), yet Xu et al. (2023) showed that these metrics correlate weakly with human judgments. Their work provided the first comprehensive empirical comparison of automatic metrics against expert preferences, revealing that standard measures fail to capture human notions of quality. Although they introduced a multi-dimensional evaluation framework informed by expert feedback, the study did not propose a new scoring system or practical evaluation tool.

Recent studies have explored using LLMs to evaluate long-form answers. D’Souza et al. (2025) proposed YESciEVAL, a framework for scientific question answering that decomposes responses into multiple reasoning dimensions and leverages LLMs for multi-aspect scoring aligned with expert annotations. However, its scope is confined to scientific domains, limiting its generalizability to open-domain LFQA. Prometheus (Kim et al., 2024) extends this idea to open-ended text generation, using fine-tuning on large human-preference datasets to achieve strong agreement with human judgments across multiple evaluation metrics. Focusing specifically on LLM-as-a-judge paradigms, JudgeBench (Tan et al., 2024) introduced the first standardized benchmark to assess LLM judgment quality, testing for consistency, bias, and alignment with human

preferences. It also identified systematic biases such as length and positional bias, though it primarily targets bias detection rather than intrinsic answer quality assessment.

Collectively, prior research underscores three enduring challenges in LFQA research: the scarcity of large-scale human preference datasets; the poor interpretability of existing automatic evaluation metrics; and the lack of understanding of the factors that drive human preference judgments. This paper addresses these gaps.

3. Preference Judgment Dataset Development

Constructing a large-scale, robust, and reliable dataset for LFQA demands a careful balance between domain diversity, annotation reliability, and scalability. This section explains the dataset-building process, the refinement of LFQA definitions, and the development of cost-effective LFQA filtering methods.

3.1. Multi-Source Dataset Aggregation

We initiated data construction by aggregating human preference judgment data from three heterogeneous sources, including both conversational and knowledge-intensive sources:

1. The **Chatbot Arena Conversation Dataset** consists of 33K crowdsourced conversations with pairwise human preference annotations and 3.3K expert-level human annotations. It contains a wide range of conversational topics and model-generated answers, making it a rich source for further analysis to capture human preference from LLM-generated responses.
2. The **LFQA Eval: LFQA Evaluation Dataset**. In contrast to typical conversational datasets, the LFQA_Eval dataset focuses on questions that require reasoning and detailed, multi-sentence responses. It contains approximately 17K long-form questions with pairwise human preference.
3. The large-scale **SHP-2: Stanford Human Preferences Dataset v2** is collected from community-driven discussion portals such as Reddit and Stack Exchange. The dataset was primarily built on the assumption that if comment A was written after comment B but still has a higher score, then A is ostensibly preferred to B. The dataset contains approximately 900K prompts and 4.8M responses, and pairwise human judgments across 129 subject areas.

3.2. LFQA Definition

Both SHP-2 and ChatBot Arena contain prompts that can not necessarily be characterized as long-

form questions. To filter them out, we need to rigorously define long-form questions. The first challenge is to differentiate a question from other types of text utterances, as the syntax alone is insufficient. At an abstract level, a question is an utterance whose illocutionary force is to elicit a response that fills an information gap in the speaker's mental state, and is not necessarily characterized by computational cues such as starting with a "wh" word or an overt interrogative intent ("could you please explain.."). For example, utterances such as "it would help to know your opinion on X" are valid questions that are not represented in most general question answering (extractive or MCQ) datasets.

One can distinguish between a short and a long-form question based on the length of a possible answer, but that alone cannot be the key determining factor, as a question can admit "correct" answers of varying lengths. We propose a rigorous and operational definition of long-form questions (LFQ) developed through an iterative annotation process involving five annotators excluding the authors: two Ph.D. students, two master's students, and one undergraduate CS student with high English proficiency.

LFQA Definition

A LFQA must express a single, well-defined information need that requires a detailed, multi-sentence answer involving exposition, explanation, reasoning, exploration, or description of the process. These questions are usually complex or open-ended, often asking "why" or "how" about a process, reason, or concept in an objective manner, which means answers should be grounded in facts, reasoning, or conceptual explanation.

It must not:

- Combine multiple distinct or loosely related sub-questions
- Request personal advice, express individual preferences, seek personal opinions, and recommendations (e.g., product, place, city, food, restaurant, flight, school, course)
- Be answerable with a single fact, a yes/no response, only code generation, direct formula calculation, or creative content (e.g., poem, joke, story, music, song lyrics, image, or audio)

In the first round, 100 questions from the **Chatbot Arena Conversation Dataset** were annotated as LFQ or not, with labels based on the majority

vote. Inter-annotator agreement, calculated using Gwet’s AC1 statistic³, was 0.65, indicating substantial agreement. Then the LFQA definition (3.2) was refined (see Definition 3.2) using annotator feedback (see Appendix A.1 for the initial LFQA definition). We conducted another annotation round with the updated definition using longer questions from SHP-2 with a Gwet’s AC1 score of 0.55, indicating moderate agreement. This implies that the annotators are consistent in identifying LFQA, even under stress testing with challenging and borderline examples.

3.3. Definition to Prompt-based Filtering

To evaluate the effectiveness of LFQA definitions, we designed two prompts: one with the initial definition and another with the refined definition. The performance of these prompts was tested using three state-of-the-art LLMs: GPT-4o, Llama-4, and Gemini-2.5, and evaluated using standard metrics, including precision, recall, and F1-score. The ground truth was defined as the majority vote from human annotations.

As shown in Table 1, the prompt derived from the updated definition improved performance across all evaluation metrics. Therefore, we selected the updated prompt for large-scale data filtering.

Sample Set	Model	Precision	Recall	F1 Score
First 100 initial definition	llama-4	0.7500	0.2553	0.3810
	gpt-4o	0.9091	0.2128	0.3448
	gemini-2.5	0.8750	0.2979	0.4444
Second 100 initial definition	llama-4	0.6905	0.6558	0.6528
	gpt-4o	0.7133	0.6485	0.6382
	gemini-2.5	0.7313	0.7346	0.7293
Second 100 updated definition	llama-4	0.9167	0.5893	0.7174
	gpt-4o	0.9459	0.6250	0.7527
	gemini-2.5	0.8140	0.6250	0.7071

Table 1: LLMs evaluation metrics (Precision, Recall, F1) on the first and second 100 samples, with both initial and updated definition.

3.4. Scaling

While GPT-4o achieved the highest precision, recall, and F1 scores with the updated definition prompt, its performance gains over Llama-4 were marginal. Given the dataset’s size ($\approx 900K$ data points), cost was a critical consideration. As we use Openrouter APIs (OpenRouter, 2024), Llama-4 Scout costs \$0.08 per million input tokens and \$0.30 per million output tokens, compared to GPT-4o’s \$5.0 and \$15.0, respectively, about 60 \times higher for inputs and 50 \times for outputs. Thus, at this scale,

³It has been shown to provide more stable estimates than Cohen’s/Fleiss’ κ under class imbalance (Gwet, 2008)

the cost difference between the two models becomes substantial, making Llama-4 a more practical choice despite the slight performance trade-off.

However, even with Llama-4, total processing costs remained substantial. To mitigate this, we first experimented with batching multiple data points but observed performance degradation as the batch size increased, an LLM limitation noted in prior work (Lin et al., 2024). We examined whether the model’s internal confidence, as reflected by the `log_score` probability, could serve as a filtering criterion. When we retained only predictions with `log_score = 0` (i.e., 100% confidence), 78% of our test samples met the condition, yielding an F1 score of 0.7174. These results indicated that incorporating `log_score` filtering can improve overall performance. We further explored combining batching with `log_score` filtering to assess the feasibility of cost-efficient processing. In Figure 2, we plot the F1-score for the positive class against cost. It can be seen that Llama-4, when batching 10 data points and considering only the highest-confidence judgment, achieves the highest performance while maintaining low cost.

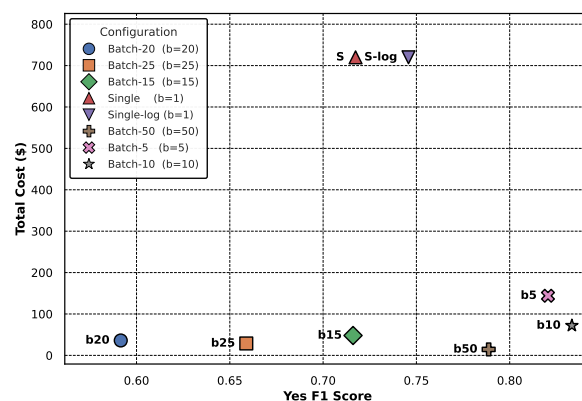


Figure 2: Effect of batch size on filtering performance for Llama-4. Batch sizes above 10 provide marginal cost gains but a large performance drop.

We further evaluated our filtering methodology on the ELI5 subreddit dataset, which is considered one of the earliest large-scale LFQA resources (Fan et al., 2019). The SHP-2 version of ELI5 contains 24,331 questions. Applying our method, 20,048 questions yielded a `log_score` of 0, among which 19,224 were identified as LFQA. Considering only the model’s 100% confidence responses, the methodology achieved 95.89% accuracy on the ELI5 dataset.

Our final multi-source dataset comprised approximately 900K unique prompts (from SHP-2, ChatbotArena, and LFQA Eval). After processing them with Llama-4 using a batch size of 10 and mapping back to their corresponding answer pairs, we ob-

tained a 1.3M-instance LFQA pairwise-judgment dataset (LFQA-HP-1M) consisting of 300K unique questions.

4. Answer Quality Evaluation Rubrics

Evaluating LFQA is challenging without standardized metrics. LLM-as-a-judge methods are promising but lack interpretability. This section examines whether a rubric-based evaluation aligns well with human preferences in LFQA.

4.1. Tools Selection for Rubric Extraction

We constructed a rubric-based evaluation schema grounded in prior LFQA and text-evaluation research (Xu et al., 2023; Krishna et al., 2021; Xian et al., 2025; Hashemi et al., 2024; Li et al., 2025) that emphasizes multi-faced evaluation of LFQA. Building on these insights, we employed nine fine-grained rubrics to evaluate the LFQA (see Tab. 2), ensuring a comprehensive, explainable understanding of answer quality.

4.1.1. Rubric Value Computation

The rubrics were designed so that most could be computed using established, non-LLM-based techniques. However, for dimensions such as relevance, coherence, and fluency, the use of LLM-based evaluators is widely accepted in the summarization literature. To operationalize these aspects, we employ G-EVAL, a widely used LLM-driven evaluation framework that relies on carefully constructed prompts to generate reliable scores for these metrics (Liu et al., 2023). However, the original G-EVAL paper did not provide results for the completeness rubric. We created a G-EVAL style prompt to measure completeness on the YESciEval (D’Souza et al., 2025) benchmark. The prompt achieved approximately 70% accuracy on the BioASQ subset, validating its reliability and justifying its use in our experimental setup. For factuality, we leveraged the Veriscore framework, which evaluates the accuracy of verifiable claims in the long text generation (Song et al., 2023). We also optimized the Veriscore framework by replacing GPT-4o with GPT-4o nano, as both model exhibits a high degree of consistency.

Specificity was quantified using the SubjectiveQA-SPECIFIC model by Pardawala et al. (2024), which offers a generalizable framework for analyzing question-answer pairs using six-dimensional features for measuring subjectivity and informational specificity. Conciseness was computed via a lexical density-based metric implemented with the spaCy model “en_core_web_sm” (Honnibal et al., 2020). Grammar was evaluated using LanguageTool (Naber, 2003). Finally, the

use of examples was detected through simple prompting to Llama-4, whose results aligned with human judgment. Together, these tools provide a multi-aspect evaluation pipeline to examine LFQA.

4.1.2. Score Normalization and Rescaling

Each selected tool outputs a score corresponding to its designated evaluation rubric (Tab. 2). Continuous scores that are not between 1-5 (e.g., Conciseness) are rescaled to a 1–5 range using min–max normalization. The use of examples is a binary indicator: 0 = absent, 1 = present. Similarly, the SubjectiveQA-SPECIFIC model outputs three ordinal specificity labels (Negative, Neutral, Positive). Following standard psychometric conventions (Stevens, 1946; Bock, 1972; McCullagh, 1980), these ordinal categories are treated as partitions of a continuous latent specificity variable, and the logits are converted into a continuous score, which is then rescaled via min–max normalization to a 1–5 range.

4.2. Rubrics Predict Human Preference

We trained a binary logistic regression model on pairwise judgment data to identify which rubrics most strongly influence human preference decisions. Each question includes two candidate answers with their corresponding rubric scores, allowing us to examine the predictive power of inter-answer differences across rubrics. To construct a balanced evaluation set, we randomly sampled 5,074 instances from the full dataset, which we refer to as LFQA-HP-Sample. The sample was designed to ensure broad coverage across domains, data sources, and annotation types, comprising 750 instances from Chatbot Arena, 750 from LFQA-Eval (balanced between expert and non-expert judgments), 2,039 from SHP-2 Reddit, and 1,536 from SHP-2 StackExchange. The dataset spans 129 distinct domains, with at least 20 questions per domain where available. It was partitioned into training (70%), development (15%), and test (15%) subsets while preserving proportional representation of sources, domains, and annotator expertise. This curated dataset serves as the foundation for rubric-level feature analysis, logistic regression modeling, and benchmarking the performance and robustness of state-of-the-art LLM evaluators.

The LR model weights are presented in Table 3 indicate that a) coherence and completeness are the strongest positive indicators for human preference; b) conciseness, use of example, fluency, factuality, and specificity have a moderate influence. Interestingly, we observed two negative weights for relevance and grammar rubrics, suggesting that these two features may be constant across both candidate answers. Also, relevance received the

Rubrics	Definition	Score range
Specificity	measures the precision and depth in explanation over general statements.	0/1/2, CT
Grammar	examines grammatical correctness, adhering to the formal writing style.	1-5, CO
Fluency	measures how clearly the answer communicates.	1-3, Co
Completeness	checks whether the answer addresses all the queries asked by the question.	1-5, CT
Coherence	captures how logically and cohesively the answer is organized.	1-5, CO
Relevance	evaluates whether the answer directly addresses the given question.	1-5, CO
Conciseness	measures how succinct yet sufficiently informative an answer is without redundancy.	0-1, CO
Use of example	checks the presence of a concrete example supporting a relevant explanation.	0/1, CT
Factuality	checks claims made in the answer are accurate, verifiable, and free from factual errors or contradictions.	0-1, CO

Table 2: LFQA evaluation rubric definitions and score ranges. CT/CO denotes whether scores are categorical or continuous.

Feature (Rubric Aspect)	Weight
Specificity	0.0420
Grammar	-0.0378
Fluency	0.0919
Completeness	0.4065
Coherence	0.5910
Relevance	-0.3082
Conciseness	0.1747
Use of Examples	0.1208
Factuality	0.0524

Table 3: Learned LR weights indicate contribution toward human preference

highest negative weights, suggesting that overly detailed and verbose answers may receive a high relevance score but were ultimately not preferred by humans.

Overall, the proposed rubrics for long-form answer evaluation provide a comprehensive foundation for quantitative evaluation.

5. Model Evaluation and Discussions

We assess whether state-of-the-art LLMs, in a zero-shot setting, outperform rubric-based linear models at identifying human preferences in LFQA. Specifically, we use three SOTA LLMs: GPT-4o, Gemini-2.5-flash, and Llama-4-Scout. The models are given a question and two answers, and are prompted with instructions to determine the *better* (i.e., the one with higher human preference) one. All models are accessed via their Openrouter APIs ⁴ using the default parameters with a temper-

ature of 0. Each model is prompted three times per data point (a zero temperature does not necessarily guarantee deterministic output (Subonis, 2024; He and Lab, 2025)), and the final judgement was determined by the majority response.

Then we examine verbosity and position bias (Li et al., 2024) in the LLMs, i.e., whether the LLM performance is influenced by a) length of the answers rather than the quality of the content, and b) position of the answers, i.e., whether they appear first or second after the question. An assessment of transitive preference judgment is also conducted to evaluate logical consistency and identify potential intransitive cycles. Finally, we evaluate the robustness of LLMs on an adversarially perturbed version of the test dataset, in which each instance included semantically preserving modifications (Roth et al., 2024) known to affect the predictions of state-of-the-art Small Language Models such as ModernBERT (Warner et al., 2024).

Table 4 shows that the logistic regression model achieves an accuracy and F1-score of 0.68, comparable to the Gemini-2.5 model (0.68) and Llama-4 model (0.68), and only marginally below the best performing GPT-4o (0.69). This result indicates that interpretable linear classifier models, built solely on fine-grained rubrics, can show comparable performance with SOTA LLMs. The top-performing GPT-4o model was also prompted with detailed definitions for the nine rubrics, but that yielded only a minimal performance improvement. Finally, we fine-tuned the ModernBERT-base model on the LFQA-HP-Sample training data, but its performance remained below that of SOTA LLMs. Overall, the tested models achieved below 70% accuracy, a result slightly higher than the findings of Chen et al. (2024), who demonstrated that open-access

⁴<https://openrouter.ai/>

preference-tuned LLMs consistently struggle to exceed 60% ranking accuracy on standard preference benchmarks.

Model	Acc.	Prec.	Rec.	F1
Logistic Regression	0.6826	0.6821	0.6816	0.6817
<i>LLM-as-a-judge</i>				
GPT-4o	0.6943	0.6970	0.6960	0.6942
Gemini-2.5-flash	0.6826	0.6864	0.6847	0.6822
Llama-4-Scout	0.6840	0.6899	0.6868	0.6833
<i>Rubrics</i>				
GPT-4o (R)	0.6979	0.6981	0.6984	0.6979
<i>Fined-Tuned</i>				
ModernBERT-base	0.6635	0.6546	0.6357	0.6450

Table 4: Comparative performance of LLM-based evaluators and logistic regression models

5.1. Preference Transitivity

To evaluate ranking consistency, we assess the transitivity of preference judgments. A preference relation is transitive if there exist at least three distinct answers for a unique question, and follows the relation that if $A \succ B$ and $B \succ C$, then $A \succ C$. In contrast, a violation occurs when a 3-cycle ($A \succ B$, $B \succ C$, $C \succ A$) forms, indicating intransitivity. However, the existence of three answers is necessary but not sufficient for running transitivity violation detection tests (see Table 5).

Choices	Transitive	Non-Transitive	Not Testable
A B	A	A	B
B C	B	B	B
C A	A	C	-

Table 5: Illustration of transitive, cyclic (3-cycle), and non-testable preference cases. In the final column, the model selects option B in two cases, thereby rendering the third option irrelevant for transitivity detection.

Among the tested LLM-based judges, GPT-4o achieves a transitivity violation rate of 4.7%, indicating strong internal consistency. In contrast, Llama-4 achieves a high violation rate of 12.2%, and Gemini-2.5 exhibits the highest rate of 15.3%, indicating inconsistencies in judging LFQA preferences.

5.2. Position and Verbosity Bias

Biases in model evaluation threaten the reliability of the automated LFQA judging system. As our `LFQA-HP-Sample` dataset consists of pairwise human preferences, we first test whether swapping the answer position or order (e.g., Answer 2 vs Answer 1) changes LLMs’ decision in judging pairwise LFQA. Ideally, the standard performance metrics should remain unchanged as the content of each

answer remains the same. However, we observed performance dropped in the table 6 for Gemini-2.5 and Llama-4 after deliberate position swaps of answers. The performance drop or deviation indicates partiality in the LLM-based evaluator, possibly due to token-level attention alignment. On the other hand, GPT-4o demonstrates relatively marginal changes.

Model	Δ Acc.	Δ Prec.	Δ Rec.	Δ F1
GPT-4o	+0.0052	+0.0043	+0.0009	+0.0024
Gemini-2.5-flash	-0.0196	-0.0201	-0.0252	-0.0238
Llama-4-Scout	-0.0168	-0.0190	-0.0228	-0.0208

Table 6: Change (Δ) in performance metrics after introducing positional bias

A similar instability may arise from verbose rather than concise answers. In Table 7, we examined and quantified cases in which LLMs’ judgments deviated from human judgments. The **Human Short** \rightarrow **LLM Long** indicates that the human-selected answer was concise, but the model favored the lengthy answers; conversely, **Human Long** \rightarrow **LLM Short** shows the opposite. From our `LFQA-HP-Sample` test dataset, GPT-4o favored 419 answers which deviated from human judgment, and out of these, 248 cases where GPT-4o selected lengthy responses over human-preferred concise ones. Gemini-2.5 also exhibits similar behavior across 248 of 432 misjudgment cases. These results indicate that GPT-4o and Gemini-2.5 may prefer lengthy answers, which could lead to human unaligned-judgment. Besides, Llama-4 demonstrated nearly balanced performance, suggesting greater content-awareness.

Model	Human Short \rightarrow LLM Long	Human Long \rightarrow LLM Short	Bias Direction
GPT-4o	248	171	Long-biased
Llama-4	232	199	Nearly balanced
Gemini-2.5-flash	248	184	Long-biased

Table 7: Cross-bias instances where model choice conflicted with human-correct answer length.

5.3. Adversarial Perturbations

To further evaluate the reliability of LLM evaluators, we test their robustness on adversarially perturbed datasets generated using the TextAttack library (Morris et al., 2020). These perturbations introduce subtle, semantics-preserving modifications such as synonym swaps and paraphrastic rewrites that should not affect the underlying answer quality or human preference. Assessing model stability under such controlled perturbations helps determine whether evaluators genuinely capture semantic and discourse-level quality or rely on superficial lexical patterns, thereby offering a more rigorous measure

of interpretive robustness in long-form answer judgment. We use both word-level and character-level perturbations:

- **TextFooler** is a black-box adversarial attack (i.e., it does not access model parameters or gradients) that first identifies important words that most influence a model’s prediction and replaces them with contextually appropriate synonyms that preserve the original meaning (Jin et al., 2020).
- **DeepWordBug** is another black-box method that introduces small modifications such as insertions, deletions, swaps, or substitutions of characters in key words such that the readability is preserved, but the model predictions are altered (Gao et al., 2018).

For both attacks, identifying the most *important* words is crucial, which is done relative to a specific model that is referred to as the *victim model*. The objective is to generate semantically consistent perturbations that cause this model to misclassify examples it would otherwise correctly predict. The victim model in these experiments was a ModernBERT-based classifier. The large (8192 tokens) max-sequence length of ModernBERT makes it a natural choice for long-context classification. The model was fine-tuned on the training part of the LFQA-HP-Sample dataset, achieving approximately 65% accuracy on the clean (unperturbed) test split of 1416 instances. Empirical results revealed superior efficacy for character-level perturbations, with DeepWordBug achieving success (i.e., it changed the ModernBERT prediction) on 25% test instances, in contrast to TextFooler’s 15.6% successful instances.

Next, we analyzed whether the modified sets generated by TextFooler and DeepWordBug could affect the performance of the best-performing LLM (GPT-4o) on our LFQA-HP-Sample dataset. As shown in Tab. 8, the GPT-4o model experiences a performance drop in all performance metrics. On the dataset modified by TextFooler, the model’s performance dropped overall by 3%. However, the DeepWordBug-generated perturbed text reduced performance by 9-10% relative to the original. The performance drop suggests that state-of-the-art LLMs are more vulnerable to orthographic perturbations than to semantically equivalent substitutions.

Perturbation	Δ Acc.(%)	Δ Prec.(%)	Δ Rec.(%)	Δ F1(%)
TextFooler	-2.64%	-2.60%	-2.61%	-2.65%
DeepWordBug	-10.67%	-11.05%	-10.89%	-10.69%

Table 8: Relative performance degradation (Δ %) of GPT-4o under adversarial perturbations.

Overall, these results suggest that a learned,

rubric-based statistical logistic model can match the best performing LLMs in predicting human preference for LFQA, while offering explainable insights into linguistic and reasoning dimensions. The analysis also reveals that LLMs’ judgment on long-form question answering can be non-transitive, influenced by answer order, verbosity, and surface-level textual variations.

6. Conclusion and Future Work

This work advances LFQA evaluation through a large-scale dataset and a rubric-informed modeling framework for human-aligned assessment. We introduce LFQA-HP-1M, a 1.3M-pair human preference dataset, by far the largest of its kind, enabling empirical analysis of human judgment in LFQA. Using nine rubric-based features, we find that coherence and completeness most strongly influence human preferences. A lightweight logistic regression model trained on these rubrics matches the performance of state-of-the-art LLM evaluators while remaining interpretable and bias-resistant. Unlike LLM-based evaluators, it is unaffected by positional or verbosity biases. Lastly, our transitivity consistency analysis reveals that LLM-based judges do not always satisfy transitive preference properties. Overall, our study presents a full rubric-driven pipeline for transparent and reliable LFQA evaluation. Future work can extend LFQA evaluation by developing richer rubric representations, such as sub-scores or learned features, and by training LLMs to autonomously design evaluation plans beyond fixed criteria.

7. Ethics Statement

As part of our work, we construct LFQA-HP-1M, a dataset comprising question–answer pairs and model-generated responses, derived from publicly available data. The dataset collection and construction process did not involve any target data collection. No private, personally identifiable, or confidential information was intentionally collected.

The LFQA-HP-1M dataset contains human judgment; however, such judgments inevitably reflect annotator subjectivity and potential biases. These biases may propagate during model training, testing, and validation. We emphasize considering human preference labels as structured assessments rather than absolute truth.

The dataset is released for research purposes only. We discourage the harmful or manipulative use of the released dataset and toolkit.

8. Limitation

The study represented a significant step towards scalable and cost-effective LFQA identification and interpretable evaluation; however, several open challenges remain. Even though the accuracy was slightly better than that of the fine-tuned model on a similar preference task (Chen et al., 2024), there is plenty of scope for improvement. The dataset’s domain generalizability remains a concern, as the LFQA-HP-1M dataset contains only question-answer-style content (excluding dialogue-style questions). Human annotations may contain noise, as even human experts sometimes disagree on the best answer, which can lower model accuracy. Additionally, the lightweight LR model with a limited number of features may miss nuanced indicators of quality that other state-of-the-art models capture. The LR model could be strengthened by considering latent cognitive features, structural hidden features, semantic hidden features, and deep latent features.

References

- R. Darrell Bock. 1972. Estimating item parameters and latent ability when responses are scored in two or more nominal categories. *Psychometrika*, 37(1):29–51.
- Angelica Chen, Sadhika Malladi, Lily H. Zhang, Xinyi Chen, Qiuyi Zhang, Rajesh Ranganath, and Kyunghyun Cho. 2024. Preference learning algorithms do not learn preference rankings. In *Proceedings of the 8th Workshop on Learning from Human Feedback (LHF ’24)*, pages 1–15. Association for Computational Linguistics.
- Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, and Noah A. Smith. 2021. [QASPER: A dataset for question answering on scientific research papers](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, pages 4600–4617, Online. Association for Computational Linguistics.
- Jennifer D’Souza, Hamed Babaei Giglou, and Quentin Münch. 2025. [YESciEval: Robust LLM-as-a-judge for scientific question answering](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13749–13783, Vienna, Austria. Association for Computational Linguistics.
- Kawin Ethayarajh, Yejin Choi, and Swabha Swayamdipta. 2022. Understanding dataset difficulty with \mathcal{V} -usable information. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 5988–6008. PMLR.
- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [ELI5: Long form question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3558–3567, Florence, Italy. Association for Computational Linguistics.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2024. [Gptscore: Evaluate as you desire](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL) – Long Papers*, pages XXX–XXX, Seattle, WA, USA. Association for Computational Linguistics.
- Ji Gao, Jack Lanchantin, Mary Lou Soffa, and Yanjun Qi. 2018. [Black-box generation of adversarial text sequences to evade deep learning classifiers](#). In *Proceedings of the 2018 IEEE Security and Privacy Workshops (SPW)*, pages 50–56. IEEE.
- Kilem L. Gwet. 2008. [Computing inter-rater reliability and its variance in the presence of high agreement](#). *British Journal of Mathematical and Statistical Psychology*, 61(1):29–48.
- Helia Hashemi, Jason Eisner, Corby Rosset, Benjamin Van Durme, and Chris Kedzie. 2024. [LLM-Rubric: A Multidimensional, Calibrated Approach to Automated Evaluation of Natural Language Texts](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13806–13834, Bangkok, Thailand. Association for Computational Linguistics.
- Horace He and Thinking Machines Lab. 2025. [Defeating nondeterminism in llm inference](#). *Thinking Machines Lab: Connectionism*. <https://thinkingmachines.ai/blog/defeating-nondeterminism-in-llm-inference/>.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. `spacy`: Industrial-strength natural language processing in python. In *Proceedings of the 12th Language Resources and Evaluation Conference (LREC 2020)*, pages 4095–4100, Marseille, France. European Language Resources Association (ELRA).
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. [TextFooler: Generating adversarial texts with human readers in mind](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages

- 1080–1090. Association for Computational Linguistics.
- Won Ik Kim, Sungho Park, Hyojun Kang, et al. 2024. [Prometheus: Towards principled evaluations of large language models as evaluators](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*. Association for Computational Linguistics.
- Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, et al. 2018. [The NarrativeQA reading comprehension challenge](#). In *Transactions of the Association for Computational Linguistics (ACL)*, volume 6, pages 317–328. Association for Computational Linguistics.
- Kalpesh Krishna, Aurko Roy, and Mohit Iyyer. 2021. [Hurdles to Progress in Long-form Question Answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 4940–4957, Online. Association for Computational Linguistics.
- Haitao Li, Qian Dong, Junjie Chen, Huixue Su, Yujia Zhou, Qingyao Ai, Ziyi Ye, and Yiqun Liu. 2024. Llm-as-judges: a comprehensive survey on llm-based evaluation methods. *arXiv preprint arXiv:2412.05579*.
- S. Yifei Li, Allen Chang, Chaitanya Malaviya, and Mark Yatskar. 2025. [ResearchQA: Evaluating Scholarly Question Answering at Scale Across 75 Fields with Survey-Mined Questions and Rubrics](#). *arXiv preprint*. ArXiv:2509.00496.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Jianzhe Lin, Maurice Diesendruck, Liang Du, and Robin Abraham. 2024. [Batchprompt: Accomplish more with less](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. [G-eval: Nlg evaluation using gpt-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2511–2522, Singapore.
- Peter McCullagh. 1980. Regression models for ordinal data. *Journal of the Royal Statistical Society: Series B (Methodological)*, 42(2):109–142.
- John Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. 2020. [TextAttack: A framework for adversarial attacks, data augmentation, and adversarial training in NLP](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 119–126, Online. Association for Computational Linguistics.
- Daniel Naber. 2003. A rule-based style and grammar checker. In *Proceedings of the 2nd International Conference on Computational Linguistics and Intelligent Text Processing (CICLing 2003)*, pages 170–174, Mexico City, Mexico.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Tristan Hume, Pamela Mishkin Kosinski, Alethea Power Chowdhery, Jared Kaplan, et al. 2022. [Webgpt: Browser-assisted question-answering with human feedback](#). *arXiv preprint arXiv:2112.09332*.
- OpenRouter. 2024. Openrouter: Unified api for accessing large language models. <https://openrouter.ai/>. Accessed: 2025-10-13.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [BLEU: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Huzaifa Pardawala, Siddhant Sukhani, Agam Shah, Veer Kejriwal, Abhishek Pillai, Rohan Bhasin, Andrew DiBiasio, Tarun Mandapati, Dhruv Adha, and Sudheer Chava. 2024. [Subjective-qa: Measuring subjectivity in earnings call transcripts’ QA through six-dimensional feature analysis](#). In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Corby Rosset, Guoqing Zheng, Victor Dibia, Ahmed Hassan Awadallah, and Paul N. Bennett. 2021. [Axiomatic preference modeling for long-form question answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 3974–3987, Online. Association for Computational Linguistics.
- Tom Roth, Yansong Gao, Alsharif Abuadbbba, Surya Nepal, and Wei Liu. 2024. Token-modification adversarial attacks for natural language processing: A survey. *AI Communications*, 37(4):655–676.

- Rachneet Singh Sachdeva, Yixiao Song, Mohit Iyer, and Iryna Gurevych. 2025. [Localizing and mitigating errors in long-form question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 20437–20469, Vienna, Austria. Association for Computational Linguistics.
- Marzieh Saeidi, Max Bartolo, Patrick Lew, Pontus Stenetorp, and Sebastian Riedel. 2018. [Interpretation of natural language rules in conversational machine reading](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2087–2097, Brussels, Belgium. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur P. Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 7881–7892, Online. Association for Computational Linguistics.
- Yixiao Song, Yekyung Kim, and Mohit Iyer. 2023. [VERISCORE: Evaluating the Factuality of Verifiable Claims in Long-Form Text Generation](#). In *Proceedings of the 2023 Conference on Language Resources and Evaluation (LREC)*, pages XXXX–XXXX, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Ivan Stelmakh, Yi Luan, Bhuwan Dhingra, and Ming-Wei Chang. 2022. [ASQA: Factoid questions meet long-form answers](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8273–8288, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- S. S. Stevens. 1946. On the theory of scales of measurement. *Science*, 103(2684):677–680.
- Dan Su, Xiaoguang Li, Jindi Zhang, Lifeng Shang, Xin Jiang, Qun Liu, and Pascale Fung. 2022. [Read before generate! faithful long form question answering with machine reading](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 744–756, Dublin, Ireland. Association for Computational Linguistics.
- Martynas Subonis. 2024. [Zero-temperature randomness in llms](#). Substack blog post, accessed February 22, 2026.
- Samson Tan, Pengfei Liu, and Graham Neubig. 2024. [JudgeBench: A benchmark for evaluating llm-as-a-judge systems](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, et al. 2024. Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference. *arXiv preprint arXiv:2412.13663*.
- Ning Xian, Yixing Fan, Ruqing Zhang, Maarten de Rijke, and Jiafeng Guo. 2025. [An Empirical Study of Evaluating Long-form Question Answering](#). *arXiv preprint*. ArXiv:2504.18413.
- Fangyuan Xu, Yixiao Song, Mohit Iyer, and Eunsol Choi. 2023. [A critical evaluation of evaluations for long-form question answering](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3225–3245, Toronto, Canada. Association for Computational Linguistics.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. BARTScore: Evaluating generated text as text generation. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [BERTScore: Evaluating text generation with BERT](#). In *International Conference on Learning Representations (ICLR)*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#).
- 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 (EMNLP)*, pages 2023–2038, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

A. Appendix

A.1. LFQA Initial Definition

LFQA Initial Definition

A long-form question answer requires a multi-sentence answer that involves explanation, reasoning, exploration, and conceptual development. It cannot be answered sufficiently with a single fact, formula application, calculation, only code, or by generating creative content (poem, joke, story, music, song lyrics, audio, picture).

A.2. Performance of LLM Evaluators at $T = 1.0$ without Majority Voting

In the main experiment, we evaluated LLM-based evaluators in a deterministic setting (temperature $T = 0$) with majority voting across three independent attempts. In the table 9, we instead report results under a higher stochastic setting ($T = 1.0$) with a single independent generation. Although minor fluctuations are observed, deterministic decoding is not enforced; the evaluator’s performance remains stable, suggesting limited sensitivity to temperature-induced randomness.

Model	Acc.	Prec.	Rec.	F1
Logistic Regression	0.6826	0.6821	0.6816	0.6817
<i>LLM-as-a-judge</i>				
GPT-4o	0.6818	0.6841	0.6834	0.6817
Gemini-2.5-flash	0.6926	0.6962	0.6947	0.6923
Llama-4-Scout	0.6891	0.6941	0.6917	0.6887
<i>Rubrics</i>				
Gemini-2.5-flash (R)	0.6848	0.6876	0.6866	0.6846

Table 9: Performance analysis of LLM evaluators at $T = 1.0$ without majority voting

A.3. LFQA Few-Shot Prompt with Initial Definition

LLMs prompt with initial definition to classify to classify LFQA

Instruction: You must answer only “yes” or “no”. Do not add any explanation. Judge whether the following question qualifies as a long-form question (LFQA) according to the definition and annotation flow below. Important: Strict response format: yes or no Use your own intuition instead of simple keyword matching Definition: A long-form question answer requires a multi-sentence answer that involves explanation, reasoning, exploration, and conceptual development. It cannot be answered sufficiently with a single fact, formula application, calculation, only code, or by generating creative content (poem, joke, story, music, song lyrics, audio, picture).

Annotation flow: Step 1: Simple question check Is the question asking for a single fact, yes/no response, only code generation, direct calculation using formulas, or creative content (poem, joke, story, music, song lyrics, audio, picture)? If yes, classify it as not LFQA and reply no If not, move to step 2

Step 2: Require multi-step answers with explanation, reasoning, exploration check Does answering the question require a multi-sentence answer involving explanation, reasoning, exploration or conceptual development? If yes, it is LFQA and reply yes If not, it is not LFQA and reply no

Examples:

Question: How do you feel? Answer: no

Question: How is pure mathematics researched? Answer: yes

Question: Who is the president of the USA? Answer: no

Question: How a bill gets introduced to the U.S. Congress Just curious about this.? Answer: Yes

Question: I remember learning about an ex prisoner who was brought to America to help train the soldiers. But the details escape me. Can anyone provide details to who he was? Answer: Yes

Input: Question:

Output: Answer:

A.4. LFQA Few-Shot Prompt with Updated Definition

LLMs prompt to classify LFQA

Instruction: Decide whether the user's question qualifies as a Long-Form Question Answering (LFQA) according to the definition.

Output Format: Strict response format: You must answer only "yes" or "no". Please do not add any explanation or rationale behind it.

Definition: A LFQA must express a single, well-defined information need that requires a detailed, multi-sentence answer involving exposition, explanation, reasoning, exploration, or description of the process. These questions are usually complex or open-ended, often asking "why" or "how" about a process, reason, or concept in an objective manner, which means answers should be grounded in facts, reasoning, or conceptual explanation.

It must not: - Combine multiple distinct or loosely related sub-questions - Request personal advice, express individual preferences, seek personal opinions, and recommendations (e.g., product, place, city, food, restaurant, flight, school, course) - Be answerable with a single fact, a yes/no response, only code generation, direct formula calculation, or creative content (e.g., poem, joke, story, music, song lyrics, image, or audio)

Examples:

Question: How do you feel? Answer: no Rationale: It expects a short, personal response, not an explanation.

Question: How is pure mathematics researched? Answer: yes Rationale: It requires a detailed explanation of abstract methods, proof techniques, and theoretical exploration—not a short answer.

Question: I have \$100 come up with a specific plan that can help me generate or increase my revenue from \$100 to \$1000 in one month Answer: no Rationale: It is a personal advice or planning request, not a factual or conceptual question requiring an explanatory long-form answer.

Question: Best documentary about evolution for someone who only learned about creationism in school? I grew up in a fundamentalist cult learning exclusively about creationism. Now that I'm free of that, I want to learn about evolution. Can someone recommend a good documentary that is suitable for someone with NO background knowledge of the subject, that's not necessarily directed at kids? Thanks in advance! Answer: no Rationale: Request for recommendations, not a question requiring a conceptual or explanatory long-form answer.

Question: Is it true that chameleons change color primarily to blend in with their surroundings? Answer: yes Rationale: It requires more than a yes/no—explaining the biological reasons for color change demands a multi-sentence explanation.

Input: Question:

Output: Answer:

A.5. LFQA Few-Shot Batch Prompt

LLMs batch prompt to classify to classify LFQA

Instruction: Decide whether the user's question qualifies as a Long-Form Question Answering (LFQA) according to the definition.

Output Format: Strict response format: You must answer only "yes" or "no". Please do not add any explanation or rationale behind it.

Definition: A LFQA must express a single, well-defined information need that requires a detailed, multi-sentence answer involving exposition, explanation, reasoning, exploration, or description of the process. These questions are usually complex or open-ended, often asking "why" or "how" about a process, reason, or concept in an objective manner, which means answers should be grounded in facts, reasoning, or conceptual explanation.

It must not: - Combine multiple distinct or loosely related sub-questions - Request personal advice, express individual preferences, seek personal opinions, and recommendations (e.g., product, place, city, food, restaurant, flight, school, course) - Be answerable with a single fact, a yes/no response, only code generation, direct formula calculation, or creative content (e.g., poem, joke, story, music, song lyrics, image, or audio)

Examples:

Question: How do you feel? Answer: no Rationale: It expects a short, personal response, not an explanation.
Question: How is pure mathematics researched? Answer: yes Rationale: It requires a detailed explanation of abstract methods, proof techniques, and theoretical exploration—not a short answer.

Question: I have \$100 come up with a specific plan that can help me generate or increase my revenue from \$100 to \$1000 in one month Answer: no Rationale: It is a personal advice or planning request, not a factual or conceptual question requiring an explanatory long-form answer.

Question: Best documentary about evolution for someone who only learned about creationism in school? I grew up in a fundamentalist cult learning exclusively about creationism. Now that I'm free of that, I want to learn about evolution. Can someone recommend a good documentary that is suitable for someone with NO background knowledge of the subject, that's not necessarily directed at kids? Thanks in advance! Answer: no Rationale: Request for recommendations, not a question requiring a conceptual or explanatory long-form answer.

Question: Is it true that chameleons change color primarily to blend in with their surroundings? Answer: yes Rationale: It requires more than a yes/no—explaining the biological reasons for color change demands a multi-sentence explanation.

—
Input Questions 1: **According to the LFQA definition & instruction, judge this question and answer yes or no.**

Input Questions 2: **According to the LFQA definition & instruction, judge this question and answer yes or no.**

Input Questions 3: **According to the LFQA definition & instruction, judge this question and answer yes or no.**

Input Questions 4: **According to the LFQA definition & instruction, judge this question and answer yes or no.**

Input Questions 5: **According to the LFQA definition & instruction, judge this question and answer yes or no.**

Input Questions 6: **According to the LFQA definition & instruction, judge this question and answer yes or no.**

Input Questions 7: **According to the LFQA definition & instruction, judge this question and answer yes or no.**

Input Questions 8: **According to the LFQA definition & instruction, judge this question and answer yes or no.**

Input Questions 9: **According to the LFQA definition & instruction, judge this question and answer yes or no.**

Input Questions 10: **According to the LFQA definition & instruction, judge this question and answer yes or no.**

—
Output:

A.6. G-Eval Prompt for Coherence Assessment

LLMs prompt to measure coherence

You will be given one answer written for a question.

Your task is to rate the answer on one metric.

Please make sure you read and understand these instructions carefully. Please keep this document open while reviewing, and refer to it as needed.

Evaluation Criteria:

Coherence (1-5) - the collective quality of all sentences. We align this dimension with the DUC quality question of structure and coherence whereby "the answer should be well-structured and well-organized. The answer should not just be a heap of related information, but should build from sentence to a coherent body of information about a topic."

Evaluation Steps:

1. Read the question carefully and identify the main topic and key points.
2. Read the answer and compare it to the question. Check if the answer covers the main topic and key points of the question, and if it presents them in a clear and logical order.
3. Assign a score for coherence on a scale of 1 to 5, where 1 is the lowest and 5 is the highest based on the Evaluation Criteria.

Example:

Question:

Answer:

Evaluation Form (scores ONLY):
- Coherence:

A.7. G-Eval Prompt for Fluency Assessment

LLMs prompt to measure fluency

You will be given one answer written for a question.
Your task is to rate the answer on one metric.
Please make sure you read and understand these instructions carefully. Please keep this document open while reviewing, and refer to it as needed.
Evaluation Criteria:
Fluency (1-3): the quality of the answer in terms of grammar, spelling, punctuation, word choice, and sentence structure.
- 1: Poor. The answer has many errors that make it hard to understand or sound unnatural.
- 2: Fair. The answer has some errors that affect the clarity or smoothness of the text, but the main points are still comprehensible.
- 3: Good. The answer has few or no errors and is easy to read and follow.
Example:
Answer:
Evaluation Form (scores ONLY):
- Fluency (1-3):

A.8. G-Eval Prompt for Relevance Assessment

LLMs prompt to measure relevance

You will be given one answer written for a question.
Your task is to rate the answer on one metric.
Please make sure you read and understand these instructions carefully. Please keep this document open while reviewing, and refer to it as needed.
Evaluation Criteria:
Relevance (1-5) - selection of important content from the question. The answer should include only relevant important information that is related to the question. Annotators were instructed to penalize answers which contained redundancies and excess information.
Evaluation Steps:
1. Read the answer and the question carefully.
2. Compare the answer to the question and identify the main points of the question.
3. Assess how well the answer covers the main points of the question, and how much irrelevant or redundant information it contains.
4. Assign a relevance score from 1 to 5.
Example:
Question:
Answer:
Evaluation Form (scores ONLY):
- Relevance:

A.9. G-Eval Prompt Adaptation for Evaluating Completeness

LLMs prompt to measure completeness

You will be given a question. You will then be given one answer written for this question.
Your task is to rate the answer on one metric.
Please make sure you read and understand these instructions carefully. Please keep this document open while reviewing, and refer to it as needed.
Evaluation Criteria:
Completeness (1-5) - the degree to which the answer is a comprehensive encapsulation of the relevant information required by the question. A complete answer addresses all parts of the question thoroughly, accurately captures the main points, and does not omit crucial elements.

Evaluation Steps:

1. Read the question carefully and identify the main aspects that must be covered for a complete response.
2. Read the answer and compare it to the question. Check if the answer addresses all essential aspects, main ideas, and pertinent details required by the question.
3. Assign a score for completeness based on the Evaluation Criteria.

Example:

Question:

Answer:

Completeness Evaluation Form (return numeric scores ONLY):