

000 001 LFQA-E: CAREFULLY BENCHMARKING LONG-FORM 002 QA EVALUATION 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

011 Long-Form Question Answering (LFQA) involves generating comprehensive,
012 paragraph-level responses to open-ended questions, which poses a significant
013 challenge for evaluation due to the richness of information and flexible response
014 format. Existing LFQA-evaluation benchmarks often lack reference answers and
015 are limited in size and topic coverage, reducing their reliability. To address
016 this gap, we introduce LFQA-E, a well-constructed, multilingual, and reference-
017 based benchmark designed to rigorously evaluate automatic metrics for LFQA.
018 LFQA-E comprises 1,625 questions and 7,649 pairwise comparisons across 15
019 topics, drawn from diverse sources such as online queries and examination ques-
020 tions, thereby enabling a comprehensive assessment of evaluation metrics. We
021 examine five categories of metrics, encompassing 17 specific methods, using
022 LFQA-E. The results demonstrate that none of the existing automatic metrics
023 perform comparably to human judgments, highlighting their inability to capture
024 the dense information in long-form responses. Furthermore, we present a detailed
025 analysis of the failure cases and the generalization capacity of these metrics, of-
026 fering insights to guide the future development of LFQA evaluation methods.
027

028 1 INTRODUCTION

030 Long-form Question Answering (LFQA) (Fan et al., 2019) targets at generating in-depth, paragraph-
031 level responses to open-ended questions. It requires models to have comprehensive domain-specific
032 knowledge or use evidence from retrieved documents (Nakano et al., 2022; Akash et al., 2023) to
033 provide relevant and accuracy answers. Despite efforts to enhance the quality of long-form answers,
034 developing automatic and reliable evaluation metrics for LFQA is still underexplored.

035 Evaluating long-form answers presents significant challenges, as evaluators must possess compre-
036 hensive domain knowledge. Previous manual evaluations typically employed crowd-sourced work-
037 ers for annotation. However, their limited domain expertise inevitably compromises reliability. In
038 contrast, expert annotation would ensure higher quality, while the cost of employing experts to an-
039 notate large-scale datasets is prohibitive. Consequently, automatic evaluation metrics are essential.
040 In automatic evaluation of LFQA, ROUGE (Lin, 2004) has been widely adopted. However, Krishna
041 et al. (2021) argue that ROUGE provides limited informativeness in long-form contexts, weakening
042 its reliability. With the advancement of LLMs (OpenAI, 2023; 2024) and Large Reasoning Mod-
043 els (LRMs) (DeepSeek-AI et al., 2025a), numerous studies have leveraged these models to develop
044 evaluation metrics through various approaches (Chang et al., 2023), including prompting (Wei et al.,
045 2023), fine-tuning (Li et al., 2023; Jiang et al., 2024), and training LLMs as Reward Models (RMs)
046 (Liu et al., 2024a; Chen et al., 2025), either generative or scalar-based. Despite the advances of eval-
047 uation metrics, determining which metrics are most effective and best aligned with human judgment
048 for LFQA evaluation requires systematic verification and benchmarking.

049 Previous benchmark for LFQA evaluation (Xu et al., 2023) samples records from reddit/ELI5, hir-
050 ing experts to annotate the better one between two responses without references, and test alignment
051 between automatic evaluation metrics and expert labels. However, their benchmark has several
052 limitations: **1) Lack of authorized references** A reference answer provides a baseline for assess-
053 ing whether a response covers key details and maintains factual accuracy. Without ground-truth
054 references, the comparison between metrics may be unfair, and evaluations without clear criteria
055 or rubrics are inherently unreliable. **2) Limited diversity** The benchmark consists of only 260

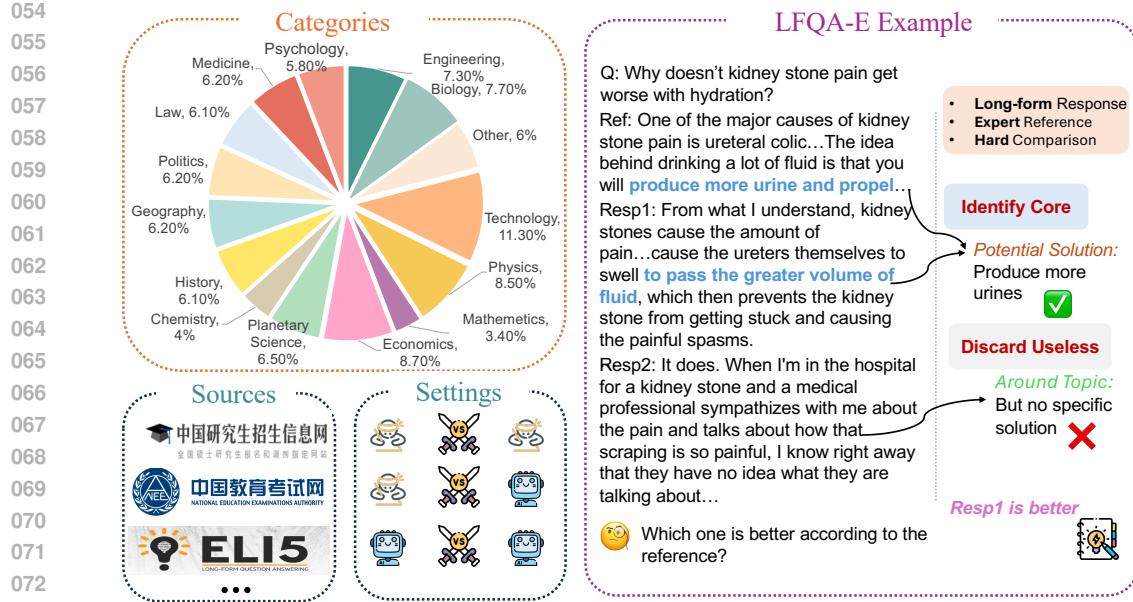


Figure 1: The figure shows the overview of LFQA-E. The left side displays the categories, sources, and three settings, showcasing its diversity. The right side illustrates an example of LFQA-E examples, all in English, constraining its linguistic and topical diversity. Moreover, it treats the comparison as an A/B task, but in real scenarios, a “tie” option always exists.

To fill the gap, we introduce LFQA-E, towards evaluating the ability of different metrics. 1) To evaluate whether current automatic evaluation metrics can select a better one from two nuanced responses, we gather references that are examined by the experts, and judge based on them. To ensure the difficulty, we choose human responses based on their upvotes or their scores, and model responses based on two models with comparable capabilities. 2) To analyze the systematic differences in validity among evaluation metrics, we rigorously assess the performance from several aspects. First we evaluate them based on three settings, i.e. human vs human (h v. h), human vs model (h v. m), and model vs model (m v. m). Moreover, we collect multilingual responses, i.e. English and Chinese, and multiple domain-specific responses, e.g., Engineering, Law, Medicine, to ensure a thorough analysis. 3) To prevent data contamination, we collect data from offline examination, i.e., College Entrance Examination Simulation Questions (CEESQ) and Postgraduate Entrance Examination Questions (PEEQ) and online platform questions (reddit/ELI5) from the recent half-year. The overview of LFQA-E Benchmark is shown in Figure 1.

Using LFQA-E, we critically assess the efficacy of 17 evaluation metrics. The experimental results show that current leading evaluation metrics fail to capture core information as human beings from verbose responses when differentiating the better one between two responses with similar quality. Furthermore, we provide analysis on why automatic evaluation metrics fail in LFQA evaluation and find the misalignment between evaluation metrics. Lastly, we try TTRL (Zuo et al., 2025) to improve the evaluation performance of model-based metrics and provide some actionable insights.

2 RELATED WORK

Development of LFQA LFQA (Fan et al., 2019) requires models to generate paragraph-level responses to open-ended questions which is more complex compared to datasets like SQuAD (Rajpurkar et al., 2016), TriviaQA (Joshi et al., 2017), and NarrativeQA (Kočiský et al., 2017), where answers are primarily words or phrases extracted directly from documents. In LFQA, models must generate comprehensive yet correct responses based on their knowledge or existing evidence documents. Several studies have analyzed the discourse structure of long-form answers (Xu et al., 2022) and have sought to enhance the performance on LFQA (Chen et al., 2023; Akash et al., 2023).

Evaluation of LFQA The automatic evaluation of LFQA remains challenging and underexplored. For human annotation, HURDLES (Krishna et al., 2021) and WEBGPT (Nakano et al., 2022) employ

108 A / B testing, where crowd-sourced annotators are instructed to choose the best of two candidate an-
 109 swers. Since annotation of LFQA requires high expertise, the results of crowd-sourced workers may
 110 be unreliable. To address the gap, Xu et al. (2023) employs experts for annotation, and tests several
 111 evaluation metrics, such as ROUGE (Lin, 2004), BERTScore (Zhang et al., 2020), and BARTScore
 112 (Yuan et al., 2021), on an expert-annotated dataset. Their findings validated that no existing metrics
 113 fully align with human judgment. However, the dataset they used lacks expert-written references,
 114 which are sourced from Reddit/ELI5, and is limited in scale, comprising only about 260 samples.
 115 More recently, since the development of LLMs and LRMs, many work uses them for evaluation of
 116 free-form answers, using prompt (Wei et al., 2023), fine-tuning using specific data (Liu et al., 2023),
 117 and reinforcement-learning (Chen et al., 2025).

3 METHODOLOGY

3.1 OVERVIEW

123 To reasonably test the evaluation ability of different metrics for LFQA when having a reference, we
 124 introduce LFQA-E, a multilingual and comprehensive benchmark composed of different topics and
 125 questions. LFQA-E consists of the Chinese version LFQA-E-ZH and the English version LFQA-
 126 E-EN. Table 1 shows its overview. It includes 1625 questions and 7649 comparisons, consisting of
 127 1493 comparisons in Chinese and 6139 comparisons in English. It spans 15 topics, ranging from
 128 history to engineering, ensuring its diversity. LFQA-E comprises expert-annotated references for
 129 fair comparison and nuanced responses. Therefore, it is naturally a hard yet reasonable benchmark
 130 for LFQA evaluation.

131 **Reference-Based Evaluation** For LFQA-E,
 132 references are sourced from academic exami-
 133 nations or widely discussed questions in Red-
 134 dit/ELI5. After being reviewed by experts
 135 with relevant academic backgrounds, these ref-
 136 erences are ensured to cover all the key points
 137 needed to answer the question. This provides a
 138 baseline for evaluation metrics to look up and
 139 provide a more precise comparison.

140 **Difficult Comparisons** All the questions
 141 contained in LFQA-E have been carefully ex-
 142 amined by domain experts to ensure it is an-
 143 swerable and clear to understand. We ensure that models have not seen the data by collecting data
 144 from recent examinations and forum questions. The responses are collected from human-written re-
 145 sponses, with close scores or upvotes, and model responses generated by comparable LLMs. There-
 146 fore, it is hard to distinguish the better one at a glance.

147 **Diverse Benchmark** We collect 1493 questions and 7649 comparisons in 15 distinct domains,
 148 from natural science to social science, to guarantee a diverse and representative benchmark. Also,
 149 LFQA-E is multilingual, consisting of examples in both Chinese and English. Moreover, LFQA-E
 150 includes three kinds of comparisons, guaranteeing the comprehensibility of the benchmark.

3.2 DATA PROCESSING

154 **Data Collection** For LFQA-E-ZH, we source our data from CEESQ and PEEQ, where ques-
 155 tions, references, and scoring schemas are developed by domain experts, including teachers and
 156 professors from high schools and colleges. The records are sampled from 2024 and are based on
 157 local examinations restored from PDF files that have not been submitted to online platforms. These
 158 questions cover diverse subjects, including politics, history, medicine, psychology, law, and geog-
 159 graphy. We provide two examples of CEESQ and PEEQ in Appendix D.1. For LFQA-E-EN, data
 160 is sourced from Reddit/ELI5, where each question is explained without specialized terminology or
 161 complex concepts, and we use the top-ranked answer as our candidate reference. For LFQA-E-ZH,
 to prevent overlap with potential training data, we avoid using data from actual College Entrance

Table 1: Detailed statistics of LFQA-E. **Avg Que. Lens**, **Avg Ref. Lens**, **Avg Res. Lens** corresponds to question lengths, reference lengths, and response lengths, respectively.

	LFQA-E-EN	LFQA-E-ZH
# Topics	9	6
# Questions	1026	599
# Comparisons	6156	1493
# Avg Que. Lens	13.4	24.6
# Avg Ref. Lens	299.1	187.2
# Avg Res. Lens	245.0	308.3
Annotate	Expert	Expert

162 Examinations. All problems are sourced from PDF files and have not been uploaded online. Additionally, the questions captured from ELI5 are all from the past six months. We provide a benchmark
 163 contamination study in Sec 5.5. To ensure that all questions are clear and answerable, we instruct
 164 GPT-4o with temperature=0.7 to filter out questions with unclear descriptions. The instruction used
 165 is listed in Appendix E.1. We conducted an experiment on the effectiveness of leveraging GPT-4o as
 166 a filter, and the results reveal that GPT-4o demonstrates superior performance for this task, achieving
 167 97% accuracy. Subsequently, to ensure our references contain all information needed to answer
 168 the questions, we submit the remaining data for expert annotation. The annotation guidelines are
 169 presented in Table 25. Each reference was annotated by two annotators, and references were dis-
 170 carded if either annotator labeled them as invalid. The Cohen’s kappa coefficient is 0.78, indicating
 171 substantial inter-annotator agreement. Through this process, we obtained 1,625 questions.
 172

173 **Human Response Collection** For LFQA-E-ZH, we gather examination papers primarily in im-
 174 age format and employ Optical Character Recognition (OCR) systems to extract student responses.
 175 Specifically, we choose student answers with close scores to ensure the comparison difficulty. The
 176 OCR is conducted using the Volcano Engine API. For LFQA-E-EN, we collect responses from the
 177 forum section of the corresponding question. Also, we select answers within the many-voted yet
 178 close up-votes to make them hard to differentiate. However, the responses we collect for LFQA-
 179 E-ZH are mainly written during examination, it is concise and well structured, and the responses
 180 we collect for LFQA-E-EN include some special characters like URLs, which deteriorate our data
 181 quality. To handle it, we use GPT-4o to paraphrase and clean our human responses. The instruction
 182 we used is shown in Appendix E.1. We randomly sample 100 records and annotate the paraphrased
 183 responses to validate the performance of GPT-4o, with the results in Appendix B.1, which indicates
 184 that leveraging LLMs doesn’t introduce any error. We also provide a case study in Appendix D.2.

185 **Model Response Generation** When generating model responses, we focus on evaluating whether
 186 LLMs can understand the semantic meaning of texts well and properly select the better response.
 187 Therefore, we do not impose extremely strict requirements on answer quality. Instead, we ensure
 188 the difficulty of LFQA-E BENCH by selecting models with similar ranking in the LMSYS Arena
 189 (Chiang et al., 2024; Zheng et al., 2023; 2024). On account of responses generated by stronger mod-
 190 els like GPT-4o or Claude-4-sonnet will pose a great challenge for our annotators to differentiate,
 191 increasing the cost of annotation, we leverage Llama-3-8B-Instruct (Dubey et al., 2024) and GPT-
 192 3.5-turbo (OpenAI, 2023) for response generation. For model-generated answers, we use *“Generate
 193 reasonable answers to the following questions. Use references or examples if needed”* to prompt
 194 LLMs. The generation temperature is set to 1.0 to encourage diverse and creative responses.

196 3.3 HUMAN ANNOTATION

197 **Annotator Decision** We hire 10 annotators from relevant aspects or who have taken relevant
 198 courses. Then we provide them with clear annotation recipes for better quality control. The an-
 199 notation recipe is in Appendix F. Each annotator receives 2\$ for annotating a question, including
 200 4-6 comparisons. To ensure the effectiveness of the annotation, we pre-annotate a subset consisting
 201 of 35 records. The annotator will start their work until they reach a 90% consistency on the subset.
 202

203 **Annotation Setting** Guided by Xu et al. (2023), our evaluation criteria mainly focus on factuality
 204 and completeness according to the reference, since almost all responses we collect are already very
 205 fluent. Unlike typical A/B testing, our method employs a triple-choice format, giving a tie option,
 206 to better capture the subtle differences between answers, as they often show comparable levels of
 207 information overlap with the reference while the other information is useless or verbose according
 208 to the central topic which can be dropped without hindering the comprehension of the response.
 209

210 **Annotation Process** The annotators assess two responses against a given reference and select the
 211 more informative and complete answer or declare a “tie” if both are of similar quality. During the
 212 process, we treat a piece of information as the basic unit as FActScore (Min et al., 2023). Initially,
 213 annotators extract the key information needed to answer the question from the provided reference
 214 and check whether the responses under evaluation contain similar statements. Then, they will select
 215 a better one based on the overlapped information. To minimize bias and subjectivity, each record
 is annotated by two independent reviewers. Each comparison takes around 7 minutes to annotate.

216 Table 2: Performance of evaluation metrics on LFQA-E. The largest value is denoted in **bold**.
217

218 219 220 221 222 223 224 225 226 227 228 229 230 231 232 233 234 235 236 237 238 239 240 241 242 243 244 245 246 247 248 249 250 251 252 253 254 255 256 257 258 259 260 261 262 263 264 265 266 267 268 269	Model	LFQA-E-EN		LFQA-E-ZH		Avg _{F1}	Avg _{Acc}
		F1	Accuracy	F1	Accuracy		
Static Evaluation Metric							
Human Baseline		77.7	83.3	68.9	76.5	73.3	79.9
Length		26.0	42.8	33.5	52.6	30.8	47.7
ROUGE		37.5	55.5	34.0	49.7	35.8	52.6
BERTScore		35.9	54.1	36.6	52.4	36.3	53.3
LLMs-based Evaluation Metric							
Qwen2.5-32B-Instruct		45.8	63.5	41.8	56.7	43.8	60.1
Qwen2.5-72B-Instruct		43.1	61.2	39.0	53.0	41.1	57.1
Llama3.1-70B-Instruct		42.5	59.6	29.4	30.7	36.0	45.2
GPT-4o		46.4	61.7	42.6	53.2	44.5	57.5
DeepSeek-V3		39.3	57.9	41.1	53.8	40.2	55.9
RM-based Evaluation Metric							
Skywork-Reward-Llama		37.3	54.4	38.2	53.6	37.8	54.0
Skywork-Reward-Gemma		37.5	56.0	33.0	48.3	35.3	52.2
RM-R1-Qwen2.5-Instruct-14B		36.4	64.9	35.1	51.9	35.8	58.4
RM-R1-DeepSeek-Distilled-Qwen-14B		43.9	65.7	36.7	53.2	40.3	59.5
LRM-based Evaluation Metric							
o1-mini		45.9	62.9	45.2	58.9	45.6	60.9
Deepseek-R1		42.9	59.6	42.4	57.8	42.7	58.7
Trained Evaluation Metric							
Auto-J-6B-bilingual		46.0	66.8	35.4	51.9	40.7	59.4
Prometheus-7B-v2.0		41.8	64.2	34.1	50.1	38.0	57.2
M-Prometheus-14B		41.6	60.8	33.9	49.4	37.8	55.1

244 For some hard-to-differentiate comparisons, detailed justification is saved to help understand. After
245 annotation, we find the Cohen’s kappa correlation of inter-annotator agreement is approximately
246 0.65, indicating a substantial agreement. We show a screenshot in Figure 5.

248 4 EXPERIMENTS

250 4.1 MODELS

252 We evaluate various metrics on LFQA-E-EN and LFQA-E-ZH respectively, including: **Static Metrics**: We use Length-orientation, ROUGE-1 (Lin, 2004) and BERTScore (F1) Zhang et al. (2020)
253 since they are widely used as the evaluation metric for LFQA. **LLMs**: We select Qwen2.5-32B-
254 Instruct (Qwen et al., 2025), Qwen2.5-72B-Instruct, Llama-3.1-70B-Instruct (Dubey et al., 2024),
255 Deepseek-V3 (DeepSeek-AI et al., 2025b), and GPT-4o (OpenAI et al., 2024). **LRMs**: Considering
256 the high time complexity and cost, we use o1-mini and Deepseek-R1 (DeepSeek-AI et al., 2025a).
257 **RM**s: We test on Skywork-Reward-Gemma-2-27B-v0.2 (Liu et al., 2024a), Skywork-Reward-
258 Llama-3.1-8B-v0.2 considering their leading position on Reward Bench (Lambert et al., 2024).
259 We also test RM-R1-Qwen2.5-Instruct-14B (Liu et al., 2024b) and RM-R1-Deepseek-Distilled-
260 Qwen-14B since they represent another paradigm of reward models. We refer to Skywork-Reward-
261 Gemma-2-27B-v0.2 and Skywork-Reward-Llama-3.1-8B-v0.2 as Skywork-Reward-Gemma and
262 Skywork-Reward-Llama for simplification. **Evaluation-Specific Models**: There are some mod-
263 els trained to be evaluation models. Among these models, we select Auto-J-Bilingual (Li et al.,
264 2023), Prometheus-7B-v2.0, and M-Prometheus-14B (Kim et al., 2024).

266 4.2 IMPLEMENTATION DETAILS

268 We evaluate all the metrics in both LFQA-E-EN and LFQA-E-ZH. We use Jieba cut for ROUGE-
269 zh. For BertScore, we use roberta-large for LFQA-E-EN evaluation and bert-base-chinese for
LFQA-E-ZH. We set the temperature at 1.0 for all LLM-based evaluation metrics to encourage

270 diverse responses. We show the results of temperature = 0 in Sec 5.1. The prompts we used are
 271 shown in Appendix E.2. For models with specific training templates, we adopt them. We include
 272 references for models to look up in all our settings. We use accuracy and macro-F1 as our indicators:
 273

$$274 \text{Acc} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\text{pred}_i = \text{label}_i) \quad (1)$$

$$275 \text{F1}_{\text{macro}} = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \left(2 \cdot \frac{P_c R_c}{P_c + R_c} \right) \quad (2)$$

276 where $\mathcal{C} = \{\text{A}, \text{B}, \text{tie}\}$. For LLM-based methods, we include a "tie" option in the instruction, while
 277 for other methods that return a scalar, we round the scalar to 3 decimal places. For the human
 278 baseline, we hire another 3 annotators with doctor's degrees to ensure the quality. The final scores
 279 are obtained by averaging annotators' results. We use the annotation recipe as previously claimed.
 280

281 4.3 MAIN RESULTS

282 Table 2 lists our results. The overall low accuracies and F1-scores of all evaluation metrics
 283 indicate the challenge LFQA-E poses to current models and methods. We also provide a
 284 cost analysis in Appendix B.4.

285 **Comparison Between Metrics** Though
 286 none of the evaluation metrics achieves a
 287 high performance on LFQA-E, we observe
 288 that scaling model size doesn't definitely
 289 yield a better result. For example, Qwen2.5-
 290 32B-Instruct beats Qwen2.5-72B-Instruct
 291 by 3%. What's more, LRM-s show a great
 292 performance compared with LLMs, thanks to their long CoT and extended thinking. RM-based
 293 evaluation metrics don't show promising results when generalizing to LFQA evaluation, perhaps
 294 because they are trained to give a better one between two responses, renouncing the "tie" option.
 295 We will analyze further and give a fairer comparison in Section B.3.

296 **Comparison Between Indicators** All evaluation metrics struggle to give a tie as good as human
 297 beings. Table 3 indicates that among the evaluation metrics we test, the best result is just 9.2%
 298 for LFQA-E-EN and 14.6% for LFQA-E-ZH. Observing the responses, we find that they are too
 299 conservative to claim two responses are of equal quality. This explains why accuracy is always
 300 larger than Macro-F1. The low accuracy on tie comparison reflects the difficulty of LFQA-E again.
 301

302 **Specialized Tuned models Help Boost Performance.** Observing the table above, we find that
 303 tuning a base model to be a robust generative reward model helps in LFQA evaluation. Moreover,
 304 when changing the base model to a strong reasoning model, the performance gains continually. In
 305 addition, SFT models can also achieve performance comparable with models of larger sizes. These
 306 phenomenon indicates that tuning is essential for LFQA evaluation.
 307

308 5 ANALYSIS

309 5.1 TEMPERATURE MATTERS

310 To ablate the effect of temperature on our validation, we further experiment using a temperature
 311 equals 0.0 for deterministic results. The result in Table 4 shows a noticeable shift in metrics. For
 312 instance, the performance of LLM-based Evaluation Metrics such as Qwen2.5-32B-Instruct and
 313 GPT-4o are higher compared to when the temperature is set to 1.0. This suggests that deterministic
 314 behavior likely reduces the model's variability in responses, leading to more consistent evaluation
 315 results. However, the LRM-based evaluation metrics show a significant drop, indicating that models
 316 relying on more exploration perform worse or even collapse under deterministic conditions.

324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
Table 4: Performance of metrics on LFQA-E when temperature is set to 0. The largest value is denoted in **bold**.

Model	LFQA-E-EN		LFQA-E-ZH		Avg_{F1}	Avg_{Acc}
	F1	Accuracy	F1	Accuracy		
LLMs-based Evaluation Metric						
Qwen2.5-32B-Instruct	46.0	64.5	40.0	54.0	43.0	59.3
Qwen2.5-72B-Instruct	41.8	58.6	38.5	52.7	40.2	55.7
Llama3.1-70B-Instruct	36.8	53.5	40.5	50.4	38.6	52.0
GPT-4o	49.5	63.1	43.9	52.7	46.7	57.9
DeepSeek-V3	40.1	57.8	40.9	55.1	40.5	56.5
LRM-based Evaluation Metric						
o1-mini	43.4	56.2	4.3	5.8	23.9	31.0
Deepseek-R1	2.0	2.7	32.1	45.3	17.1	24.0

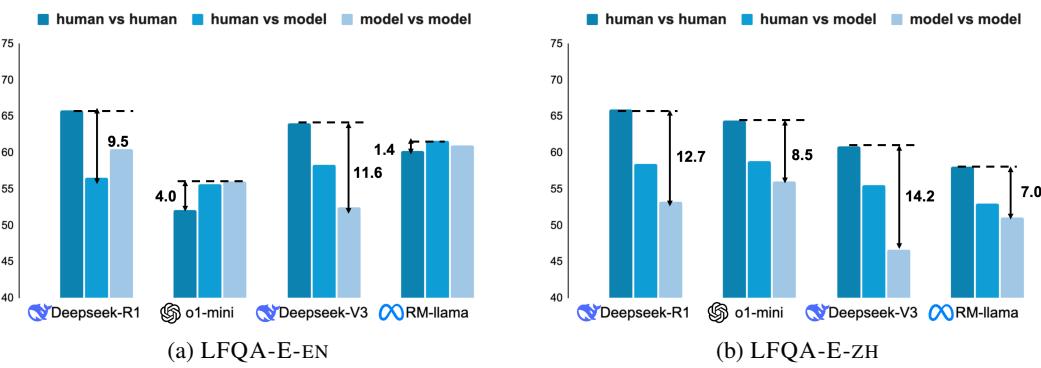


Figure 2: Performance of different models on our three settings on LFQA-E.

5.2 EVALUATION METRICS CAN'T EXCEL AT ALL SETTINGS.

To have a thorough understanding of whether the model evaluates human response or model response differently, we experiment on a different group of LFQA-E. We break it into three groups, i.e., h v. h , h v. m , and m v. m , where h indicates human response and m represents model response, and see the accuracy changes. The results are listed in Figure 2a for LFQA-E-EN and Figure 2b for LFQA-E-ZH. We can observe that for many evaluation metrics, there exists a huge difference between different comparison settings. In LFQA-E-EN, the RMs show steady ability while others exhibit degradation when model responses are introduced. In LFQA-E-ZH, all the metrics show a drastic accuracy decline under m v. m , with a maximum drop of 14.2% from Deepseek-V3. This further validates our assumption that current evaluation metrics can't differentiate between two nuanced responses. we also show the performance across subjects in Appendix B.2.

5.3 REASONS EVALUATION METRICS FAIL WHEN EVALUATING LONG-FORM RESPONSES.

For LM-based Evaluation Metric We observe the outputs of several LLMs and find that almost all errors arise from the following aspects.

- *Keypoints Identification Error*: The model fails to correctly identify and separate bullet keypoints or enumerated lists in responses, leading to poorly structured answers.
- *Irrelevant/Incorrect Information Error*: The model does not penalize or filter out irrelevant or factually incorrect details in its responses, reducing accuracy.
- *Contradiction Error*: During reasoning, the model generates inconsistent or contradictory statements due to factual hallucinations.
- *Formatting Error*: The model produces responses with an improper format.

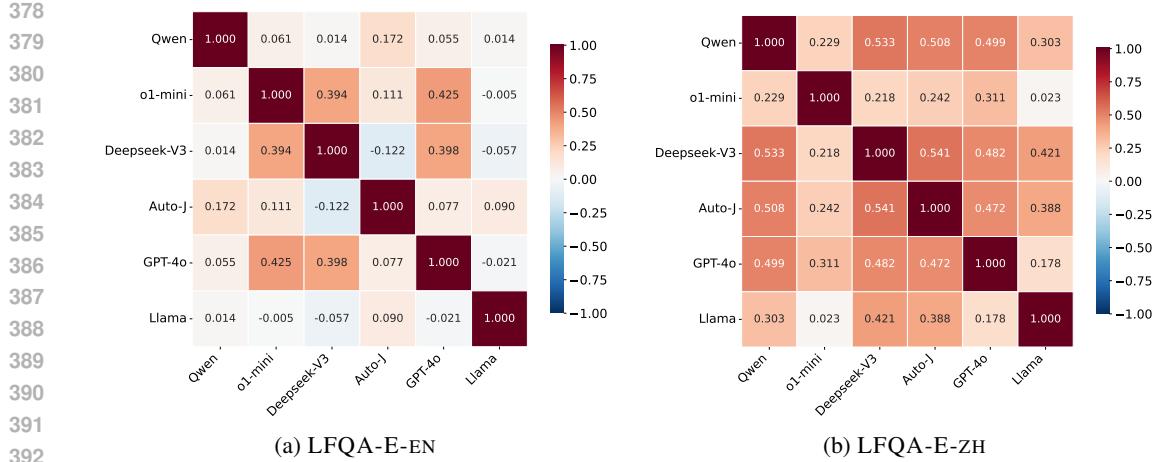


Figure 4: The Cohen’s Kappa Correlation Matrix in LFQA-E.

We show the probability of each error occurring in Figure 3. We choose Deepseek-V3 and o1-mini for representation. *Point Identification Error* and *Irrelevant/Incorrect Information Error* happen most time, indicating the relatively low inherent ability for LMs when evaluating long-form answers.

Static Evaluation Metrics These methods simply leverage word-level or embedding-level similarities, which scratch on surface when evaluating. As described in Fan et al. (2024), considering evaluating two long responses around a topic, there may be many words overlapping. Also, overly long responses dilute semantics, making originally important key information trivial, so metrics fail to consider informativeness, but only focus on similarity.

5.4 DIFFERENT EVALUATION METRICS DON’T AGREE WITH EACH OTHER.

To find whether there is a correlation between different evaluation metrics, we observe detailed evaluation results. Specifically, we select ROUGE, Qwen2.5-32B-Instruct (simplified as Qwen), GPT-4o, Skywork-Reward-Llama (simplified as Llama), o1-mini, and Auto-J-6B-bilingual (simplified as Auto-J), considering their relatively better performance on LFQA-E. Figure 4a and 4b show the results. We observe that neither of the two metrics achieves a high correlation, indicating two metrics may contradict each other to a large degree. There are even some negative correlations between the two metrics under LFQA-E-EN. This phenomenon further illustrates that there is no stable evaluation result across different metrics.

5.5 BENCHMARK CONTAMINATION ANALYSIS

To examine whether our evaluation benchmark has potential contamination from pretraining corpora, we conducted perplexity (PPL) and n -gram overlap the same as Xu et al. (2024) analysis using two widely adopted open-source instruction-tuned models: Qwen2.5-7B-Instruct and Llama-3.1-8B-Instruct. The underlying intuition is that if evaluation data is memorized during pretraining, the models would exhibit abnormally low perplexity and high n -gram overlap on the evaluation benchmark.

Perplexity Analysis. Table 5 reports the PPL values on both English and Chinese subsets. For Qwen2.5-7B-Instruct, the PPL is 11.60 (en) and 11.73 (zh). For Llama-3.1-8B-Instruct, the PPL is 11.70 (en) and lower at 7.21 (zh), which all fall within tolerable limits and low data contamination potential.

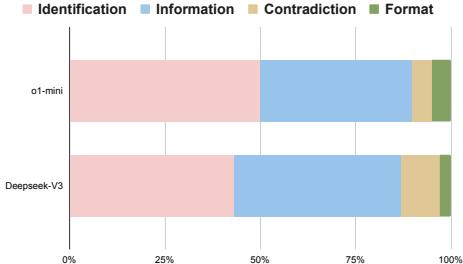


Figure 3: Percentage of error types for LMs on the LFQA-E dataset.

432 Table 5: Perplexity (PPL) on the benchmark.
433

434 Model	435 en	436 zh
436 Qwen2.5-7B-Instruct	437 11.60	11.73
437 Llama-3.1-8B-Instruct	11.70	7.21

438 **n-gram Overlap.** We further computed n-gram exact match accuracy between the benchmark
439 and model generations. As shown in Table 6, the overlap scores remain low across both models,
440 e.g., 0.025–0.030 in English and 0.047–0.093 in Chinese. These values are significantly below
441 contamination thresholds observed in prior work, reinforcing the view that large-scale memorization
442 is unlikely to happen in our benchmark.

443 Overall, both analyses suggest that the benchmark is not heavily contaminated. The English
444 subset appears relatively safe across both models. For the Chinese subset, while Llama-3.1-
445 8B-Instruct shows somewhat lower perplexity and higher n-gram overlap, though n-gram
446 higher than expected, remain within acceptable bounds and don't indicate severe memorization.
447

448 5.6 TTRL TO BOOST PERFORMANCE

449 Here, we try to improve the performance on LFQA-E through prompting and reinforcement
450 learning (RL). Firstly, we use a structured prompt to instruct models to embrace their
451 response within a `<answer>...</answer>` tag. This leads to a non-trivial performance in-
452 crease, considering the relatively lower instruction following ability of small language models.
453 Building on this, since RL is widely used to improve the reasoning ability of LLMs (Cui et al., 2025;
454 Fan et al., 2025), we use RL to leverage the performance of LLMs on LFQA evaluation. Consid-
455 ering the lack of high-quality data, we implement TTRL (Zuo et al., 2025) on our LFQA-E-EN
456 as an example. We configure our model with a batch size of 8, a rollout temperature of 1.0, and
457 generate 32 rollouts per prompt. The learning rate is $5e-7$. During validation, the temperature is 0.0
458 for consistent results. The reward signal is based on an outcome-based rule, similar to the approach
459 used in Deepseek-R1 (DeepSeek-AI et al., 2025a). As shown in Table 7, TTRL yield a substantial
460 performance boost, demonstrating the effectiveness of using RL for this task. Also, the response
461 length grows steadily during the training until the training rewards converge. However, we observe
462 a rapid convergence where all rollouts produce identical preferences, which limits further improve-
463 ments. We attribute this to the underlying three-category classification which is easily overfitted
464 when the model is over-confident. For a remedy, we implement clip-higher mechanism used in
465 DAPO (Yu et al., 2025). Also, we observe a performance boost, indicating that when we increase
466 the diversity during rollouts, RL can take effect more stably. However, more sophisticated methods
467 and high-quality data are needed for a better evaluation metric that mimics human preferences.
468

469 6 CONCLUSION

470 We introduce LFQA-E, a multilingual benchmark for LFQA evaluation. It consists of 1625 ques-
471 tions and 7649 comparisons, spanning 15 topics, from natural science to social science, consisting
472 of 3 settings, i.e., h v. h , h v. m , and m v. m . Each records include a clear question, an authorized ref-
473 erence, and two hard-to-differentiate responses, ensuring its difficulty. We conduct experiments on
474 15 automatic evaluation metrics. The results show that none of the metrics can evaluate long-form
475 responses as well as human beings. We further analyze the generalization of different metrics across
476 languages and settings. The results further indicate that all models struggle to generalize well to all
477 comparisons. We find that LRM and specifically trained evaluation models lead on LFQA-E. The
478 test-time-scaled evaluation model may be used to enhance the performance of LFQA evaluation.
479

480 Table 6: n -gram accuracy on the benchmark.
481

482 Model	483 en	484 zh
484 Qwen2.5-7B-Instruct	485 0.030	0.047
485 Llama-3.1-8B-Instruct	0.025	0.093

486 Table 7: Performance of TTRL.
487

488 Model	489 LFQA-E-EN
490 Qwen2.5-3B-Instruct	
491 CoT	49.6
492 Structured Prompt	59.7
493 TTRL	63.9
494 TTRL + Clip Higher	66.5
495 Qwen2.5-7B-Instruct	
496 CoT	53.3
497 Structured Prompt	60.6
498 TTRL	68.2
499 TTRL + Clip Higher	68.6

500

501

502

503

504

505

506

507

508

509

510

511

512

513

514

515

516

517

518

519

520

521

522

523

524

525

526

527

528

529

530

531

532

533

534

535

536

537

538

539

540

541

542

543

544

545

546

547

548

549

550

551

552

553

554

555

556

557

558

559

560

561

562

563

564

565

566

567

568

569

570

571

572

573

574

575

576

577

578

579

580

581

582

583

584

585

586

587

588

589

590

591

592

593

594

595

596

597

598

599

600

601

602

603

604

605

606

607

608

609

610

611

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648

649

650

651

652

653

654

655

656

657

658

659

660

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702

703

704

705

706

707

708

709

710

711

712

713

714

715

716

717

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756

757

758

486 ETHICS STATEMENT
487488 This paper proposes a benchmark consisting of questions and references from examination papers
489 and online forums. We perform data filtering to ensure that there is no offense in the benchmark
490 data.
491492 REPRODUCIBILITY STATEMENT
493494 We have provide our settings comprehensively for data construction and evaluation so that the result
495 can be reproduced using the same setting and prompt.
496497 REFERENCES
498499 Pritom Saha Akash, Kashob Kumar Roy, Lucian Popa, and Kevin Chen-Chuan Chang. Long-form
500 question answering: An iterative planning-retrieval-generation approach, 2023. URL <https://arxiv.org/abs/2311.09383>.
501502 Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan
503 Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and
504 Xing Xie. A survey on evaluation of large language models, 2023. URL <https://arxiv.org/abs/2307.03109>.
505506 Hung-Ting Chen, Fangyuan Xu, Shane Arora, and Eunsol Choi. Understanding retrieval augmen-
507 tation for long-form question answering, 2023. URL <https://arxiv.org/abs/2310.12150>.
508509 Xiusi Chen, Gaotang Li, Ziqi Wang, Bowen Jin, Cheng Qian, Yu Wang, Hongru Wang, Yu Zhang,
510 Denghui Zhang, Tong Zhang, Hanghang Tong, and Heng Ji. Rm-r1: Reward modeling as reason-
511 ing, 2025. URL <https://arxiv.org/abs/2505.02387>.
512513 Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolaos Angelopoulos, Tianle Li,
514 Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica.
515 Chatbot arena: An open platform for evaluating llms by human preference, 2024.
516517 Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu
518 Yu, Qixin Xu, Weize Chen, et al. Process reinforcement through implicit rewards. *arXiv preprint*
519 *arXiv:2502.01456*, 2025.
520521 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,
522 Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu,
523 Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao
524 Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan,
525 Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao,
526 Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding,
527 Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang
528 Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai
529 Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang,
530 Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang,
531 Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang,
532 Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang,
533 R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng
534 Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing
535 Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen
536 Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong
537 Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu,
538 Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaoqin Shen, Xi-
539 aoshia Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia
Wang, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng
Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong
Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong,

540 Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou,
 541 Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying
 542 Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda
 543 Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu,
 544 Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu
 545 Zhang, and Zhen Zhang. Deepseek-r1: Incentivizing reasoning capability in llms via reinforce-
 546 ment learning, 2025a. URL <https://arxiv.org/abs/2501.12948>.

547 DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Cheng-
 548 gang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang,
 549 Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting
 550 Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui
 551 Ding, Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi
 552 Ni, Jiashi Li, Jiawei Wang, Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li,
 553 Junxiao Song, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang,
 554 Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Litong Wang, Liyue Zhang, Meng Li, Miaojun
 555 Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan
 556 Huang, Peiyi Wang, Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J.
 557 Chen, R. L. Jin, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin Xu, Ruoyu Zhang,
 558 Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhua Chen, Shaoqing Wu, Shengfeng
 559 Ye, Shengfeng Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Shut-
 560 ing Pan, T. Wang, Tao Yun, Tian Pei, Tianyu Sun, W. L. Xiao, Wangding Zeng, Wanjia Zhao,
 561 Wei An, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, X. Q. Li, Xiangyue
 562 Jin, Xianzu Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaojin Shen, Xiaokang Chen, Xi-
 563 aokang Zhang, Xiaosha Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Cheng, Xin
 564 Liu, Xin Xie, Xingchao Liu, Xingkai Yu, Xinnan Song, Xinxia Shan, Xinyi Zhou, Xinyu Yang,
 565 Xinyuan Li, Xuecheng Su, Xuheng Lin, Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Yang
 566 Zhang, Yanhong Xu, Yanhong Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui
 567 Li, Yaohui Wang, Yi Yu, Yi Zheng, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Ying
 568 Tang, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yu Wu,
 569 Yuan Ou, Yuchen Zhu, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yukun Zha, Yunfan
 570 Xiong, Yunxian Ma, Yuting Yan, Yuxiang Luo, Yuxuan Liu, Yuyang Zhou, Z. F.
 571 Wu, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhen Huang, Zhen Zhang, Zhenda
 572 Xie, Zhengyan Zhang, Zhewen Hao, Zhibin Gou, Zhicheng Ma, Zhigang Yan, Zhihong Shao,
 573 Zhipeng Xu, Zhiyu Wu, Zhongyu Zhang, Zhuoshu Li, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li,
 Ziwei Xie, Ziyang Song, Ziyi Gao, and Zizheng Pan. Deepseek-v3 technical report, 2025b. URL
<https://arxiv.org/abs/2412.19437>.

574 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
 575 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony
 576 Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark,
 577 Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere,
 578 Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris
 579 Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong,
 580 Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny
 581 Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino,
 582 Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael
 583 Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Ander-
 584 son, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah
 585 Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan
 586 Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vrane, Jason Park, Jay Ma-
 587 hadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy
 588 Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak,
 589 Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Al-
 590 wala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini,
 591 Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Lauren Rantala-Yeary, Laurens van der
 592 Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo,
 593 Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Man-
 natt Singh, Manohar Paluri, Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova,
 Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal,

594 Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur
 595 Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhar-
 596 gava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong,
 597 Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic,
 598 Roberta Raileanu, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sum-
 599 baly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa,
 600 Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang,
 601 Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende,
 602 Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney
 603 Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom,
 604 Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta,
 605 Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Voeti, Vladan Petro-
 606 vic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang,
 607 Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur,
 608 Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre
 609 Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aaron Grattafiori, Abha
 610 Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay
 611 Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baevski, Allie Feinstein, Amanda
 612 Kallet, Amit Sangani, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew
 613 Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita
 614 Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh
 615 Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De
 616 Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon
 617 Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina
 618 Mejia, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris
 619 Cai, Chris Tindal, Christoph Feichtenhofer, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel
 620 Li, Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Di-
 621 ana Liskovich, Didem Foss, Dingkang Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa
 622 Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Es-
 623 teban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat Ozgenel,
 624 Francesco Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Flo-
 625 rez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Govind Thattai, Grant Her-
 626 man, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou,
 627 Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren,
 628 Hunter Goldman, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman,
 629 James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer
 630 Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe
 631 Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie
 632 Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khandelwal, Katayoun
 633 Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal
 634 Chawla, Kushal Lakhotia, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva,
 635 Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madiam
 636 Khabsa, Manav Avalani, Manish Bhatt, Maria Tsimpoukelli, Martynas Mankus, Matan Hasson,
 637 Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Ke-
 638 neally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel
 639 Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mo-
 640 hammad Rastegari, Munish Bansal, Nandini Santhanam, Natascha Parks, Natasha White, Navy-
 641 ata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning Dong,
 642 Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli,
 643 Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux,
 644 Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvaraj, Qian Liang, Rachad Alao,
 645 Rachel Rodriguez, Rafi Ayub, Raghatham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li,
 646 Rebekkah Hogan, Robin Battey, Rocky Wang, Rohan Maheswari, Russ Howes, Ruty Rinott,
 647 Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Sa-
 tadrup Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lind-
 648 say, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang
 649 Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen
 650 Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho,
 651 Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser,

648 Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Tim-
 649 othy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan,
 650 Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu,
 651 Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xi-
 652 aocheng Tang, Xiaofang Wang, Xiaojian Wu, Xiaolan Wang, Xide Xia, Xilun Wu, Xinbo Gao,
 653 Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin
 654 Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick,
 655 Zhao du Wen, Zhenyu Yang, and Zhiwei Zhao. The llama 3 herd of models, 2024. URL
 656 <https://arxiv.org/abs/2407.21783>.

657 Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. Eli5:
 658 Long form question answering, 2019. URL <https://arxiv.org/abs/1907.09190>.

659 Yuchen Fan, Xin Zhong, Yazhe Wan, Chengsi Wang, Haonan Cheng, Gaoche Wu, Ning Ding, and
 660 Bowen Zhou. Eva-score: Evaluating abstractive long-form summarization on informativeness
 661 through extraction and validation, 2024. URL <https://arxiv.org/abs/2407.04969>.

662 Yuchen Fan, Kaiyan Zhang, Heng Zhou, Yuxin Zuo, Yanxu Chen, Yu Fu, Xinwei Long, Xuekai
 663 Zhu, Che Jiang, Yuchen Zhang, et al. Ssrl: Self-search reinforcement learning. *arXiv preprint*
 664 *arXiv:2508.10874*, 2025.

665 Dongfu Jiang, Yishan Li, Ge Zhang, Wenhao Huang, Bill Yuchen Lin, and Wenhui Chen. Tiger-
 666 score: Towards building explainable metric for all text generation tasks, 2024. URL <https://arxiv.org/abs/2310.00752>.

667 Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly
 668 supervised challenge dataset for reading comprehension, 2017. URL <https://arxiv.org/abs/1705.03551>.

669 Seungone Kim, Juyoung Suk, Shayne Longpre, Bill Yuchen Lin, Jamin Shin, Sean Welleck, Graham
 670 Neubig, Moontae Lee, Kyungjae Lee, and Minjoon Seo. Prometheus 2: An open source language
 671 model specialized in evaluating other language models, 2024.

672 Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis,
 673 and Edward Grefenstette. The narrativeqa reading comprehension challenge, 2017. URL <https://arxiv.org/abs/1712.07040>.

674 Kalpesh Krishna, Aurko Roy, and Mohit Iyyer. Hurdles to progress in long-form question answering,
 675 2021. URL <https://arxiv.org/abs/2103.06332>.

676 Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu,
 677 Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, Noah A. Smith, and Hannaneh Hajishirzi.
 678 Rewardbench: Evaluating reward models for language modeling. <https://huggingface.co/spaces/allenai/reward-bench>, 2024.

679 Junlong Li, Shichao Sun, Weizhe Yuan, Run-Ze Fan, Hai Zhao, and Pengfei Liu. Generative judge
 680 for evaluating alignment, 2023. URL <https://arxiv.org/abs/2310.05470>.

681 Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization
 682 Branches Out*, pp. 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics.
 683 URL <https://aclanthology.org/W04-1013>.

684 Chris Yuhao Liu, Liang Zeng, Jiacai Liu, Rui Yan, Jujie He, Chaojie Wang, Shuicheng Yan, Yang
 685 Liu, and Yahui Zhou. Skywork-reward: Bag of tricks for reward modeling in llms. *arXiv preprint*
 686 *arXiv:2410.18451*, 2024a.

687 Xiusi Liu, Gaotang Li, Ziqi Wang, Bowen Jin, Cheng Qian, Yu Wang, Hongru Wang, Yu Zhang,
 688 Denghui Zhang, Tong Zhang, Hanghang Tong, and Heng Ji. Rm-r1: Reward modeling as reason-
 689 ing, 2024b. URL <https://arxiv.org/abs/2505.02387>.

690 Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. G-eval: Nlg
 691 evaluation using gpt-4 with better human alignment, 2023. URL <https://arxiv.org/abs/2303.16634>.

702 Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen tau Yih, Pang Wei Koh, Mohit Iyyer,
 703 Luke Zettlemoyer, and Hannaneh Hajishirzi. Factscore: Fine-grained atomic evaluation of fac-
 704 tual precision in long form text generation, 2023. URL <https://arxiv.org/abs/2305.14251>.

705

706 Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher
 707 Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou,
 708 Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. Webgpt:
 709 Browser-assisted question-answering with human feedback, 2022. URL <https://arxiv.org/abs/2112.09332>.

710

711

712 OpenAI. Chatgpt: Chat generative pre-trained transformer. <https://chat.openai.com/>,
 713 2023. Accessed: 2024-08-05.

714

715 OpenAI. Hello gpt-4o. <https://openai.com/index/hello-gpt-4o/>, 2024. Accessed:
 716 2024-08-05.

717 OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan
 718 Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-
 719 Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol,
 720 Alex Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Con-
 721 neau, Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian,
 722 Amin Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein,
 723 Andrew Cann, Andrew Codispoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrew
 724 Mishchenko, Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia,
 725 Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben
 726 Leimberger, Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake
 727 Samic, Bob McGrew, Bobby Spero, Bogo Giertler, Bowen Cheng, Brad Lightcap, Brandon
 728 Walkin, Brendan Quinn, Brian Guaraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo
 729 Lugaressi, Carroll Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li,
 730 Chan Jun Shern, Channing Conger, Charlotte Barette, Chelsea Voss, Chen Ding, Cheng Lu,
 731 Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim,
 732 Christine Choi, Christine McLeavey, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley
 733 Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, Dane Sherburn, Daniel Kappler,
 734 Daniel Levin, Daniel Levy, David Carr, David Farhi, David Mely, David Robinson, David Sasaki,
 735 Denny Jin, Dev Valladares, Dimitris Tsipras, Doug Li, Duc Phong Nguyen, Duncan Findlay,
 736 Edede Oiwoh, Edmund Wong, Ehsan Asdar, Elizabeth Proehl, Elizabeth Yang, Eric Antonow,
 737 Eric Kramer, Eric Peterson, Eric Sigler, Eric Wallace, Eugene Brevdo, Evan Mays, Farzad Kho-
 738 rasani, Felipe Petroski Such, Filippo Raso, Francis Zhang, Fred von Lohmann, Freddie Sulit,
 739 Gabriel Goh, Gene Oden, Geoff Salmon, Giulio Starace, Greg Brockman, Hadi Salman, Haiming
 740 Bao, Haitang Hu, Hannah Wong, Haoyu Wang, Heather Schmidt, Heather Whitney, Heewoo Jun,
 741 Hendrik Kirchner, Henrique Ponde de Oliveira Pinto, Hongyu Ren, Huiwen Chang, Hyung Won
 742 Chung, Ian Kivlichan, Ian O'Connell, Ian O'Connell, Ian Osband, Ian Silber, Ian Sohl, Ibrahim
 743 Okuyucu, Ikai Lan, Ilya Kostrikov, Ilya Sutskever, Ingmar Kanitscheider, Ishaan Gulrajani, Ja-
 744 cob Coxon, Jacob Menick, Jakub Pachocki, James Aung, James Betker, James Crooks, James
 745 Lennon, Jamie Kiro, Jan Leike, Jane Park, Jason Kwon, Jason Phang, Jason Teplitz, Jason Wei,
 746 Jason Wolfe, Jay Chen, Jeff Harris, Jenia Varavva, Jessica Gan Lee, Jessica Shieh, Ji Lin, Jiahui
 747 Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joanne Jang, Joaquin Quinonero Candela, Joe Beutler, Joe
 748 Landers, Joel Parish, Johannes Heidecke, John Schulman, Jonathan Lachman, Jonathan McKay,
 749 Jonathan Uesato, Jonathan Ward, Jong Wook Kim, Joost Huizinga, Jordan Sitkin, Jos Kraaijeveld,
 750 Josh Gross, Josh Kaplan, Josh Snyder, Joshua Achiam, Joy Jiao, Joyce Lee, Juntang Zhuang,
 751 Justyn Harriman, Kai Fricke, Kai Hayashi, Karan Singh, Katy Shi, Kavin Karthik, Kayla Wood,
 752 Kendra Rimbach, Kenny Hsu, Kenny Nguyen, Keren Gu-Lemberg, Kevin Button, Kevin Liu, Kiel
 753 Howe, Krithika Muthukumar, Kyle Luther, Lama Ahmad, Larry Kai, Lauren Itow, Lauren Work-
 754 man, Leher Pathak, Leo Chen, Li Jing, Lia Guy, Liam Fedus, Liang Zhou, Lien Mamitsuka,
 755 Lilian Weng, Lindsay McCallum, Lindsey Held, Long Ouyang, Louis Feuvrier, Lu Zhang, Lukas
 Kondraciuk, Lukasz Kaiser, Luke Hewitt, Luke Metz, Lyric Doshi, Mada Afshak, Maddie Simens,
 Madelaine Boyd, Madeleine Thompson, Marat Dukhan, Mark Chen, Mark Gray, Mark Hudnall,
 Marvin Zhang, Marwan Aljubeh, Mateusz Litwin, Matthew Zeng, Max Johnson, Maya Shetty,
 Mayank Gupta, Meghan Shah, Mehmet Yatbaz, Meng Jia Yang, Mengchao Zhong, Mia Glaese,

756 Mianna Chen, Michael Janner, Michael Lampe, Michael Petrov, Michael Wu, Michele Wang,
 757 Michelle Fradin, Michelle Pokrass, Miguel Castro, Miguel Oom Temudo de Castro, Mikhail
 758 Pavlov, Miles Brundage, Miles Wang, Minal Khan, Mira Murati, Mo Bavarian, Molly Lin, Murat
 759 Yesildal, Nacho Soto, Natalia Gimelshein, Natalie Cone, Natalie Staudacher, Natalie Summers,
 760 Natan LaFontaine, Neil Chowdhury, Nick Ryder, Nick Stathas, Nick Turley, Nik Tezak, Niko Felix,
 761 Nithanth Kudige, Nitish Keskar, Noah Deutsch, Noel Bundick, Nora Puckett, Ofir Nachum,
 762 Ola Okelola, Oleg Boiko, Oleg Murk, Oliver Jaffe, Olivia Watkins, Olivier Godement, Owen
 763 Campbell-Moore, Patrick Chao, Paul McMillan, Pavel Belov, Peng Su, Peter Bak, Peter Bakkum,
 Peter Deng, Peter Dolan, Peter Hoeschele, Peter Welinder, Phil Tillet, Philip Pronin, Philippe
 764 Tillet, Prafulla Dhariwal, Qiming Yuan, Rachel Dias, Rachel Lim, Rahul Arora, Rajan Troll, Ran-
 765 dall Lin, Rapha Gontijo Lopes, Raul Puri, Reah Miyara, Reimar Leike, Renaud Gaubert, Reza
 766 Zamani, Ricky Wang, Rob Donnelly, Rob Honsby, Rocky Smith, Rohan Sahai, Rohit Ramchan-
 767 dani, Romain Huet, Rory Carmichael, Rowan Zellers, Roy Chen, Ruby Chen, Ruslan Nigmat-
 768 ullin, Ryan Cheu, Saachi Jain, Sam Altman, Sam Schoenholz, Sam Toizer, Samuel Miserendino,
 769 Sandhini Agarwal, Sara Culver, Scott Ethersmith, Scott Gray, Sean Grove, Sean Metzger, Shamez
 770 Hermani, Shantanu Jain, Shengjia Zhao, Sherwin Wu, Shino Jomoto, Shirong Wu, Shuaiqi Xia,
 771 Sonia Phene, Spencer Papay, Srinivas Narayanan, Steve Coffey, Steve Lee, Stewart Hall, Suchir
 772 Balaji, Tal Broda, Tal Stramer, Tao Xu, Tarun Gogineni, Taya Christianson, Ted Sanders, Tejal
 773 Patwardhan, Thomas Cunningham, Thomas Degry, Thomas Dimson, Thomas Raoux, Thomas
 774 Shadwell, Tianhao Zheng, Todd Underwood, Todor Markov, Toki Sherbakov, Tom Rubin, Tom
 775 Stasi, Tomer Kaftan, Tristan Heywood, Troy Peterson, Tyce Walters, Tyna Eloundou, Valerie Qi,
 776 Veit Moeller, Vinnie Monaco, Vishal Kuo, Vlad Fomenko, Wayne Chang, Weiyi Zheng, Wenda
 777 Zhou, Wesam Manassra, Will Sheu, Wojciech Zaremba, Yash Patil, Yilei Qian, Yongjik Kim,
 778 Youlong Cheng, Yu Zhang, Yuchen He, Yuchen Zhang, Yujia Jin, Yunxing Dai, and Yury Malkov.
 779 Gpt-4o system card, 2024. URL <https://arxiv.org/abs/2410.21276>.

780 Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan
 781 Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang,
 782 Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin
 783 Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li,
 784 Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang,
 785 Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025.
 786 URL <https://arxiv.org/abs/2412.15115>.

787 Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions
 788 for machine comprehension of text, 2016. URL <https://arxiv.org/abs/1606.05250>.

789 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc
 790 Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models,
 791 2023. URL <https://arxiv.org/abs/2201.11903>.

792 Fangyuan Xu, Junyi Jessy Li, and Eunsol Choi. How do we answer complex questions: Discourse
 793 structure of long-form answers, 2022. URL <https://arxiv.org/abs/2203.11048>.

794 Fangyuan Xu, Yixiao Song, Mohit Iyyer, and Eunsol Choi. A critical evaluation of evaluations for
 795 long-form question answering, 2023. URL <https://arxiv.org/abs/2305.18201>.

796 Ruijie Xu, Zengzhi Wang, Run-Ze Fan, and Pengfei Liu. Benchmarking benchmark leakage in
 797 large language models. *arXiv preprint arXiv:2404.18824*, 2024. URL <https://arxiv.org/abs/2404.18824>.

798 Qiyi Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai,
 799 Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guang-
 800 ming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu,
 801 Jiaze Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Yuxuan Song, Xiangpeng Wei, Hao
 802 Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingx-
 803 uan Wang. Dapo: An open-source llm reinforcement learning system at scale, 2025. URL
 804 <https://arxiv.org/abs/2503.14476>.

805 Weizhe Yuan, Graham Neubig, and Pengfei Liu. Bartscore: Evaluating generated text as text gener-
 806 ation, 2021. URL <https://arxiv.org/abs/2106.11520>.

810 Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. Bertscore: Evalu-
811 ating text generation with bert, 2020. URL <https://arxiv.org/abs/1904.09675>.
812

813 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,
814 Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Sto-
815 ica. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Thirty-seventh Confer-
816 ence on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL
817 <https://openreview.net/forum?id=uccHPGDlao>.

818 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Tianle Li, Siyuan Zhuang, Zhanghao Wu, Yonghao
819 Zhuang, Zhuohan Li, Zi Lin, Eric Xing, Joseph E. Gonzalez, Ion Stoica, and Hao Zhang. Lmsys-
820 chat-1m: A large-scale real-world llm conversation dataset. In *The Twelfth International Confer-
821 ence on Learning Representations*, 2024. URL <https://openreview.net/forum?id=BOFDKxfwt0>.
822

823 Yuxin Zuo, Kaiyan Zhang, Li Sheng, Shang Qu, Ganqu Cui, Xuekai Zhu, Haozhan Li, Yuchen
824 Zhang, Xinwei Long, Ermo Hua, et al. Ttrl: Test-time reinforcement learning. *arXiv preprint
825 arXiv:2504.16084*, 2025.
826

827

828

829

830

831

832

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864

A USE OF LLMs

865
866 We use LLMs to refine our writing using Gemini2.5-Pro, GPT-5, and Claude-4.1. We check the
867 refined phrases after generation.
868869

B ADDITIONAL RESULTS

870

B.1 ANNOTATION OF LLM PARAPHRASING

871
872 We randomly sample 100 records to compare the original responses and the paraphrased responses.
873 We define three types of errors: 1) The paraphrased response contains factual errors; 2) The para-
874 phrased response adds new core points to answer the problem, which may reverse the order; 3) The para-
875 phrased response drops original points or introduces inaccuracy or inconsistency. After annota-
876 tion, we find *error 2* in only 1 out of 100 paraphrased response, indicating the effectiveness of using
877 GPT-4o to paraphrase responses.
878879

B.2 THE ABILITY OF EVALUATION METRICS VARIES ACROSS DOMAINS.

880
881 The performance of evaluation metrics in LFQA-E varies significantly across different domains. As
882 shown in Tables 8 and 9, models and metrics exhibit distinct strengths and weaknesses depending
883 on the subject area. For instance, on the LFQA-E-ZH dataset, models like Qwen2.5-32B-Instruct
884 and GPT-4o consistently excel in subjects such as Geography, Law, and Medicine but perform
885 less effectively in more complex domains like Psychology and History. Conversely, models like
886 DeepSeek-V3 and RM-based Evaluation Metrics show particular strengths in fields like Politics and
887 Law, where the emphasis is on factual accuracy and legal context. Similarly, in the LFQA-E-EN
888 dataset, LLM-based models such as Qwen2.5-32B-Instruct perform exceptionally well in Engineering
889 and Technology but show a noticeable decline in subjects like Psychology and Mathematics,
890 possibly due to the more abstract nature of these domains. In contrast, RM-based Evaluation Met-
891 rics, such as RM-R1-DeepSeek-Distilled-Qwen-14B, demonstrate impressive performance in fields
892 like Planetary Science and Chemistry, where the data is often more structured and less ambiguous.
893894

Table 8: Results of different topics from LFQA-E-ZH. The largest value is denoted using **bold**.

METRICS	GEOGRAPHY	HISTORY	POLITICS	PSYCHOLOGY	MEDICINE	LAW
LLM-based Evaluation Metric						
Qwen2.5-32B-Instruct	64.3	58.2	50.4	44.8	65.1	49.0
Qwen2.5-72B-Instruct	65.5	54.3	51.1	42.7	58.3	54.9
Llama-3.1-70B-Instruct	41.5	34.4	37.0	46.3	46.6	62.6
GPT-4o	62.0	51.8	49.1	44.8	58.3	46.1
Deepseek-V3	60.0	50.5	54.4	40.6	57.3	49.0
RM-based Evaluation Metric						
Skywork-Reward-Llama	58.3	54.1	51.4	34.4	59.2	54.9
Skywork-Reward-Gemma	54.3	50.0	41.0	36.5	48.5	56.9
RM-R1-Qwen2.5-Instruct-14B	58.6	51.0	45.5	44.8	58.4	54.9
RM-R1-DeepSeek-Distilled-Qwen-14B	61.4	51.8	62.5	37.5	54.5	57.8
LRM-based Evaluation Metric						
o1-mini	53.5	46.9	42.5	49.0	66.0	56.9
Deepseek-R1	64.0	58.4	53.2	42.7	67.0	54.9
Trained Evaluation Metric						
Auto-J-6B-bilingual	58.5	55.4	41.7	38.5	58.4	59.8
Prometheus-7B-v2.0	54.0	52.6	43.2	40.6	51.5	59.8
M-Prometheus-14B	56.0	51.3	43.7	41.7	44.7	50.0

912
913

B.3 LLMs ARE BETTER AT FINDING SOMETHING BETTER.

914
915 Considering that giving a tie option is difficult for both humans and models, we drop out the records
916 that are labeled as a tie and conduct the experiments again. We show the results in Table 10. After
917 discarding the tied comparison, all the evaluation metrics show nontrivial performance boosts. GPT-
918 4o even gets a 6.4% bonus. This increase matches what we find when comparing indicators. Similar

918 Table 9: Results of different topics from LFQA-E-EN. The largest value is denoted using **bold**.
919

920 METRICS	921 ENGINEERING	922 BIOLOGY	923 TECHNOLOGY	924 PHYSICS	925 MATHEMATICS	926 ECONOMICS	927 PLANETARY SCIENCE	928 CHEMISTRY	929 OTHER
LLM-based Evaluation Metric									
Qwen2.5-32B-Instruct	59.0	64.7	59.0	52.3	55.6	63.0	64.7	64.7	65.7
Qwen2.5-72B-Instruct	59.8	65.7	60.2	52.2	57.1	61.9	61.6	67.5	63.4
Llama-3.1-70B-Instruct	59.7	61.4	53.2	51.0	55.0	61.2	60.5	62.6	58.0
GPT-4o	60.5	65.1	60.4	55.9	58.1	63.0	64.4	64.4	65.5
Deepseek-V3	52.6	64.9	52.8	51.9	54.4	62.7	64.4	64.7	57.1
RM-based Evaluation Metric									
Skywork-Reward-Llama	55.9	57.1	54.1	55.9	56.2	57.1	58.0	63.1	53.5
Skywork-Reward-Gemma	56.9	59.0	59.2	56.3	56.8	57.0	58.5	56.6	60.0
RM-R1-Qwen2.5-Instruct-14B	56.9	59.0	59.2	56.3	56.8	57.0	58.5	56.6	60.0
RM-R1-DeepSeek-Distilled-Qwen-14B	65.7	72.1	64.5	63.9	61.3	67.7	73.0	71.4	69.0
LRM-based Evaluation Metric									
o1-mini	56.2	64.7	53.9	54.0	58.7	64.3	67.3	62.6	57.0
Deepseek-r1	54.0	58.6	54.3	48.3	47.1	58.8	56.3	62.6	56.1
Trained Evaluation Metric									
Auto-J-6B-bilingual	72.3	68.3	73.1	65.2	64.3	67.1	67.4	69.6	72.9
Prometheus-7B-v2.0	66.8	66.4	69.0	65.1	65.1	63.9	64.0	72.0	57.0
M-Prometheus-14B	64.1	63.5	62.0	59.5	55.2	62.1	64.9	70.6	65.5

931 Table 10: Performance of different evaluation metrics on LFQA-E. The examples whose labels are
932 tie are discarded for fairer comparison. The largest value is denoted using **bold**.
933

934 MODEL	935 LFQA-E-EN	936 LFQA-E-ZH	937 Avg
Static Evaluation Metric			
Length	42.7	56.9	49.8 ↑ 2.1%
ROUGE	57.5	53.8	55.7 ↑ 3.1%
BERTScore	56.0	56.6	56.3 ↑ 3.0%
LLM-based Evaluation Metric			
Qwen2.5-32B-Instruct	66.8	62.8	64.8 ↑ 4.7%
Qwen2.5-72B-Instruct	63.4	57.6	60.5 ↑ 3.4%
Llama-3.1-70B-Instruct	66.1	34.0	50.1 ↑ 4.9%
GPT-4o	66.3	61.4	63.9 ↑ 6.4%
Deepseek-V3	60.0	60.0	60.0 ↑ 4.1%
RM-based Evaluation Metric			
Skywork-Reward-Llama	56.7	58.4	57.6 ↑ 3.6%
Skywork-Reward-Gemma	58.2	52.5	55.4 ↑ 3.2%
RM-R1-Qwen2.5-Instruct-14B	67.5	56.2	61.9 ↑ 3.5%
RM-R1-DeepSeek-Distilled-Qwen-14B	68.0	57.5	62.8 ↑ 3.3%
LRM-based Evaluation Metric			
o1-mini	67.3	64.2	65.8 ↑ 4.9%
Deepseek-R1	61.6	63.1	62.4 ↑ 3.7%
Trained Evaluation Metric			
Auto-J-6B-bilingual	70.0	57.3	63.7 ↑ 4.3%
Prometheus-7B-v2.0	66.5	54.1	60.3 ↑ 3.1%
M-Prometheus-14B	63.7	53.3	58.5 ↑ 3.4%

962 to what we observe above, LRM_s remain leading on the fairer comparison, and RM_s still struggle to
963 generalize to long-form response evaluation. Static evaluation metrics, however, show the least im-
964 provement. The experimental results demonstrate the potential of test-time scaling, while reflects the
965 generalization problem of RM_s. What's more, specific evaluation models show their great potential
966 once again, ranking first on LFQA-E-EN, displaying its future for LFQA evaluation.

967

B.4 HOW TO BALANCE THE COST AND TIME EFFICIENCY BETWEEN MODELS?

970 To further discuss the tradeoff between efficiency and effectiveness among language models, we cal-
971 culate the average inference time per question of several LLM_s, LRM_s and RM_s. Table 13 lists the
972 results. Specifically, for the closed-source models that need to invoke api for inference, we calculate

972 their costs on both LFQA-E-EN and LFQA-E-ZH in Table 14. We observe that **o1-mini** incurs
 973 the highest cost, which is aligned with its relatively higher accuracy and F1 scores. **DeepSeek-
 974 V3**, on the other hand, offers a more cost-effective alternative, though its performance is somewhat
 975 lower. For a more practical and cost-efficient evaluation, models like Skywork-Reward-Llama and
 976 Skywork-Reward-Gemma have much smaller parameter scales and can be deployed on local re-
 977 sources, showing strong promise for reducing costs while still providing a reasonable evaluation.
 978

979 C DISCUSSION

982 We suggest training open-domain RMs or evaluation models, which may help for the evaluation of
 983 LFQA, considering their relatively low cost and GPU requirements, but with a decent score. Also,
 984 we recommend future work to focus on evaluation workflows that combine the strengths of both
 985 LLM-based models and more efficient reward models. Particularly, for open-sourced models, we
 986 don’t observe steady performance gains as model size scales. Therefore, smaller models with more
 987 training data may help more than larger models with some well-designed prompts.
 988

989 D CASE STUDY

991 D.1 CASE FROM CEESQ AND PEEQ

993 We provide two cases from CEESQ and PEEQ in Table 11. We translate them into English for easier
 994 comprehension.

996 Table 11: Case Study from CEESQ and PEEQ.

998 QUESTION FROM PEEQ:

999 What are the classifications and percentages of white blood cells?

1000 REFERENCE:

1001 White blood cells are divided into granular cells and non-granular cells. Granular cells include:
 1002 neutrophils (50%-70%); basophils (0-1%); eosinophils (0.5%-5%). Non-granular cells include:
 1003 monocytes (3%-8%); lymphocytes (20%-40%)

1004 QUESTION FROM CEESQ:

1005 Analyze the natural reasons for the numerous sandbars in the Yangtze River estuary area.

1006 REFERENCE

1008 The river has a large discharge volume and carries a large amount of sediment; located at the
 1009 river estuary, the terrain is low and flat, the flow velocity is slow, with deposition as the main pro-
 1010 cess, leading to massive sediment accumulation; situated at the river-sea interface, tidal backing
 1011 enhances the deposition process, forming numerous sandbars in the estuary area.

1015 D.2 CASE FOR PARAPHRASING

1016 We provide a case study from LFQA-E-ZH for paraphrasing in Table 12.

1019 D.3 FAILED CASE

1020 We show error cases in Table 15, Table 16, and Table 17. Table 15 shows the *Incorrect Information*
 1021 *Error*. Table 16 shows the *Point Identification Error*. Table 17 shows the *Format Error*.

1024 D.4 SUCCESS CASE

1025 We show successful cases in Table 18 and Table 19.

1026 Table 12: Case Study for parahrasing.
10271028 **QUESTION FROM LFQA-E-ZH:**
10291030 Reasons for the extremely fragile ecological environment in Guizhou Province.
10311032 **ORIGINAL RESPONSE:**
10331034 1. Guizhou features karst topography with undulating terrain
1035 2. Serious soil erosion
1036 3. Guizhou Province has few plains and is susceptible to natural disasters
1037 4. Rugged terrain with inconvenient transportation
10381039 **PARAPHRASED RESPONSE:**
10401041 The ecological environment of Guizhou Province is extremely fragile for four main reasons:
1042 First, Guizhou features karst topography with undulating terrain and complex landforms; Sec-
1043 ond, serious soil erosion has caused severe depletion of land resources; Third, Guizhou Province
1044 has few plains, with most areas being mountainous, making it susceptible to natural disasters;
1045 Fourth, the rugged terrain with inconvenient transportation creates certain difficulties for ecolog-
1046 ical environment protection and restoration.
10471048 Table 13: The average inference time of different language models per question, arranged in de-
1049 scending order.
1050

1051 Model	1052 TIME (SECONDS)
Skywork-Reward-Llama	0.21
Qwen2.5-32B-Instruct	0.27
Skywork-Reward-Gemma	0.64
Qwen2.5-72B-Instruct	1.03
o1-mini	1.80
GPT-4o	3.56
DeepSeek-V3	12.00
DeepSeek-R1	14.13

1053 Table 14: Overall cost on closed-source models using API. The largest value is denoted in **bold**.
1054

1055 Model	1056 LFQA-E-EN	1057 LFQA-E-ZH
DeepSeek-V3	3.6\$	0.7\$
DeepSeek-R1	18.0\$	6.9\$
o1-mini	67.0\$	14.2\$
GPT-4o	35.6\$	5.7\$

1058 **E INSTRUCTIONS**
10591060 The following are instructions we used.
10611062 **E.1 DATA COLLECTION INSTRUCTIONS**
10631064 The instruction for the question filter is in Table 20. The instruction for the paraphrase is in Table
1065 21.
10661067 **E.2 ENGLISH LLM EVALUATION INSTRUCTION**
10681069 The instruction for all LLMs and LRMs is in Table 22. The instruction for Prometheus series is in
1070 Table 23. The instruction for Auto-J is in Table 24.
1071

1080

1081

1082 Table 15: Case Study from LFQA-E from GPT-4o. This results from an incorrect information error.

1083

1084

1085

QUESTION:

How does law enforcement track down a criminal during a manhunt?

1086

1087

1088

1089

1090

1091

1092

1093

1094

1095

1096

REFERENCE:

- 1) Do we know where he/she is likely to “go to ground?” We try to find people he/she relies on from phone taps, social media, prior crimes etc. You’d be amazed how many guys are found at **their aunt’s/ex’s/former cellmate’s place**.
- 2) Put the word out to **local TV news**. In my experience they plaster the manhunt news and mugshot on their website landing page right away.
- 3) License plate scanner hits on **patrol cars and traffic cameras**, if available. This is surprisingly effective.
- 4) **Social media phishing**. Is he/she into certain drugs? Certain types of prostitutes? Is he trying to find his wife or kids who are working with victim services and in hiding (very common in bad domestic abuse cases)?

RESPONSE 1:

There are a lot of ways that law enforcement can track someone during a manhunt. First and foremost, we live in the age of technology, and technology is a really powerful tool for tracking someone. A cell phone is pretty much always transmitting your approximate location, and law enforcement can gain access to the satellite that keeps track of that location.

There are also a few tried and true methods: the public. Law enforcement will almost always provide a physical description of a suspect, complete with a photograph or artist rendition, and urge people to report anyone who fits the bill close enough to be the person. This is a powerful tool when to get anything modern, you need to interact with other humans or come into close contact with human society.

There are also recordings, photographs, and footage captured by **CCTV cameras everywhere**. Police may also deploy tracking hounds if they know who you are and have anything that has your smell on it and escaping on foot from persistent, trained hunting dogs is unbelievably difficult.”

RESPONSE 2:

Ohhhhh so many ways. Really depends on what resources are on hand.

1. **Suspect probably has friends and family.** Check their houses, interview them, see what they know about the suspect’s hobbies and motives. Do they have a goal they’re trying to achieve? Is that goal local or would it send them running elsewhere?
2. Flowing from above, roadblocks. If you generally know where your suspect is, where they might try to go, and what means of transportation are available to them, you can monitor those points of access until they show up.
3. Technology. Where did the suspect use their credit card. Oh, they pulled out cash to avoid being traced? We can see where they made the withdrawal and go from there. Is their phone still online? Can we call and trace them? **Were they recently active on social media? How about their friends?** A lot of those posts have geotags in the metadata

MODEL JUDGE:

- Response 1 aligns well with the standard answer by emphasizing technology (cell phone tracking), public involvement (providing descriptions and photos), and surveillance methods (CCTV cameras and tracking dogs). It touches upon finding the suspect’s connections, similar to the standard answer’s first point about identifying likely places to “go to ground.”
- Response 2 also covers relevant points such as investigating friends and family, setting up roadblocks, and using technology like phone tracing and social media analysis. However, it delves into areas like credit card tracking, which the standard answer specifically mentions as ineffective.

Overall, **Response 1** more closely mirrors the strategies outlined in the standard answer. (**The Judgement is incorrect by incorrectly estimating the weights of irrelevant or unimportant key-points.**)

1132

1133

1134 Table 16: Case Study from LFQA-E from GPT-4o. This results from an identification error.
11351136 **QUESTION:**1137 Why do swear words exist? What is it about swear words that make us feel better? Why don't
1138 other normal words make us feel the same way swear words do?
11391140 **REFERENCE:**1141 Wear words fall broadly into two categories: [profanity](#) and [obscenity](#). Profanity originates with
1142 the Ten Commandments, specifically "Do not take the Lord's name in vain" in Abrahamic reli-
1143 gions. Generally the idea though is "don't disparage or devalue the sacred". This is where the
1144 term "curse words" comes from, as you would curse someone by saying "God damn you". It's
1145 also where the term "swear words" comes from, i.e. "I swear to God". Obscenity, by contrast, is
1146 primarily words that [relate to sex](#) or [using the toilet](#). These are activities during which you would
1147 be vulnerable to being attacked, and also that produce excrement that could be harmful for oth-
1148 ers. There is sometimes a need to discuss these things in polite company, [so terms were divided](#)
1149 [into "polite" and "rude"](#). In English, this distinction largely comes from the language of origin.
1150 English has linguistic roots in [both Romance languages](#) (ones based on Latin like French) and
1151 [Germanic languages](#). At the time when the language was developing, [higher society had more](#)
1152 [background in Romance languages while the lower classes were more Germanic](#). So the former
1153 words were considered refined and the latter crude and crass. Most English obscenities have
1154 close analogs in German. In modern usage, there's actually official definitions of profanity and
1155 obscenity by the FCC. Profanity is using these words in an [exclamatory sense](#), while obscenity
1156 is using these words to [describe the actions](#). The former is more allowable than the latter. As for
1157 why it feels good to swear, [it's mostly because breaking a taboo gives a feeling of exhilaration](#),
1158 and if nobody actually suffers as a result, you also get to feel like you're "getting away with it".
1159 Basically, if we didn't put so much emphasis on NOT saying the words, it wouldn't feel so good
1160 when we do say them.
11611162 **RESPONSE 1:**1163 Language is a tool of communication. Expletives communicate something that you basically
1164 cannot express using polite language. There's certainly an element to them where the social
1165 transgression is part of what makes them function as they do. What is or isn't considered a swear
1166 word largely depends on cultural values. In the Middle Ages it's pretty apparent that words
1167 [related to bodily functions or sexuality were not considered particularly offensive](#), while blas-
1168 phemies were generally considered a lot more offensive than they are in most modern societies.
1169 Medieval folks appear to have been simultaneously less prudish but more fearful of the almighty.
11701171 **RESPONSE 2:**1172 Language has been used to separate classes for a very long time. Do you know why a pig is a pig
1173 in the pen but pork on your plate? Or why a cow becomes beef when you serve it up? Because
1174 one word was used by the lower class who raised the animals and the other word was used by
1175 the upper class who could afford to consume it (in this case, Normans). [Words used by the lower](#)
1176 [classes frequently, but not by upper classes](#), will be found to be taboo in the upper classes, leading
1177 them to be essentially banned in polite society. They have power because they are taboo, and for
1178 no other reason. The word 'merde' has no power to you unless you speak French, and the word
1179 'crap' has no power to you unless you're 6. [You've been conditioned to believe these words are](#)
1180 [strong, so you feel strong when saying them](#). It really is that simple. It's nothing more than the
1181 placebo effect.
11821183 **MODEL JUDGE:**1184 Upon evaluating both responses against the reference answer:
11851186 - Response 1 aligns closely with the standard answer by discussing the role of swear words
1187 in communication, the significance of social transgression, and cultural variations in what is
1188 considered offensive. It also touches on historical perspectives, similar to how the reference
1189 answer differentiates profanity and obscenity based on origins and societal values.
1190 - Response 2 introduces the concept of language separating classes and the power of taboo stem-
1191 ming from class distinctions. While it addresses the conditioning and perceived strength of swear
1192 words, it deviates more from the primary focus of the standard answer, which centers on cultural
1193 and linguistic origins rather than class-based language differentiation.
1194 Therefore, **Response 1** is better. (*The Judgement is incorrect by observing the number of infor-
1195 mation coexisting.*)
1196

1188
1189

Table 17: Case Study For LFQA-E from GPT-4o. This results from a format error.

1190
1191
1192
1193**QUESTION:**

Can someone explain Benford's Law. I get that certain numbers show up more often in large data sets, but why?

1194
1195
1196
1197
1198
1199
1200
1201
1202**REFERENCE:**

Benford's Law is typically stated about data that span multiple orders of magnitude; so we have values that are many times larger than others instead of them being close together. Next is that we have at least a bit of what is called scale-invariance: unlike many simpler probability problems, we don't assume that all numbers are equally likely, but instead that larger ones are rarer. More precisely, we want roughly the same chance to land between x and $2x$ as for landing between y and $2y$. If those two are satisfied, Benford's Law states that the leading digit 1 is more likely than an other, and the larger digits are the less common ones. There can be no perfect examples of scale-invariance in probability. But between those lie dozens of orders of magnitude, easily more than enough within reason...

1203
1204
1205
1206
1207
1208
1209
1210**RESPONSE 1:**

Benford's law says if you gather up a bunch of random numbers that appear "in the wild" and you tally up the first digits of each of these numbers then you will tend to find lots of 1's and not many 9's. This tends to happen when you have a range of numbers which is spread out over multiple orders of magnitude. The reason is when numbers are spread out over many orders of magnitude, it tends to be roughly uniform when measured on a "log scale". This means that the number of data points between x and $2x$ will be about the same for different values of x . In other words, there should be as many values between 100 and 200 as there are between 200 and 400.

1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222**RESPONSE 2:**

Benford's law says the first digit of a number picked from a large range of numbers tends to be a 1. Think of it like this between 1 and 20, 11 of the 20 numbers start with a 1. Between 1 and 99, 11 of the 99 numbers start with a 1. This repeats for 0-200 and 0-999 etc always holding the max of "just over half" and the min of "about 11 percent" so if you average that for all ranges, you get that about 30 percent of numbers in an unknown range start with 1. Between 1 and 89, 1 number starts with 9, basically 1/89 percent. Going up to 99 brings us back to 11/99, but now 11 is the max and "almost 0" is the min, so again average it over all ranges, and you get more like 4 percent of lead numbers. You can then apply this to some fraud cases. If the numbers span multiple orders of magnitude AND should be roughly random, AND there are a lot of them, you should expect them to match Benford's law pretty well. If they don't, one of the 3 requirements is probably false. If you know the first and last are true, you can say "these probably aren't actually random".

1223

MODEL JUDGE:**- Strengths of Response 1:**

Correctly identifies that Benford's Law applies to numbers spanning large ranges.

Explains the higher probability of 1 as a leading digit by comparing intervals which aligns with the scale-invariance idea.

Mentions the decreasing probability for higher digits.

Correctly notes the conditions for applying Benford's Law and its use in fraud detection.

- Weaknesses of Response 1:

Doesn't explicitly mention scale-invariance, though the intuition is somewhat captured.

The explanation for why 1 is more likely is a bit simplistic.

- Strengths of Response 2: Introduces Benford's Law in a simple, relatable context. Correctly notes that 1 and 2 are more common as leading digits.

- Weaknesses of Response 2:

Misrepresents the core reasons for Benford's Law. The explanations given are incorrect or misleading. Benford's Law is not about human behavior but about inherent properties of scale-invariant data.

Doesn't mention orders of magnitude or scale-invariance at all, which are central to the reference answer.

The references at the end are irrelevant to the actual content of the response.

No preference provided!

Table 18: Case Study For LFQA-E from Deepseek-V3.

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

QUESTION:

how do metals put together with electricity cause technology and chips and computers and data etc?

REFERENCE:

Silicon along with a little bit of some other elements combined form what is called a semiconductor, of which there are 2 types, the n type and the p type. When you then sandwich these semiconductors in a specific way, it has the property that applying a current to these sandwiches causes it to become conductive, or if you invert the sandwich it has the property that it's conductive when you don't apply a current and then not conductive when you do. So essentially, what you have created is an electrical switch, which is controlled by electricity, and we call these transistors.

If you're looking at how these transistors then form a computer, you can arrange transistors to perform logical functions. A simple one would be the NOT gate, which just uses the second switch I mentioned that turns off when there is an input. An AND gate would be if you took two switches, placed them one after the other, and each switch is then controlled by an input. Both inputs need to be on for there to be an output.

RESPONSE 1:

For that, you need to look at boolean values and at turing machines. A boolean value is just a yes no answer. Yes vs no, on vs off, one vs zero. Technically you use boolean values all the time in speech. "Is Tommy in the bathroom?" You could answer "Yes", but you could also say "one", and anyone who knows Boolean will understand.

A Turing machine is just a name for any machine that uses Boolean values. It has rules, for example "if you see seven zeros in a row followed by a one, then do". It seems pretty abstract. But math people have actually proved that you can do any calculation, think any thought, perform any task, if you have the right rules designed for your turing machine.

So once we figured that out, it just became a matter of inventing the right rules for the machine, and inventing the best form for a machine to store boolean values. A transistor is always either a One or a Zero, and they can be made very small, so it makes sense to use them for your Turing Machine.

RESPONSE 2:

When metals are put together with electricity in the context of technology, it typically refers to the process of creating integrated circuits or microchips. These chips are made up of layers of different metals, such as silicon, aluminum, and copper, that are carefully arranged on a silicon wafer.

Electricity is used to power the various components on the chip, allowing them to perform their specific functions. The flow of electricity through these metal layers enables the chip to process and store information, ultimately allowing computers and other electronic devices to function.

Additionally, metals such as gold and copper are used in the construction of circuit boards and connectors within electronic devices. These metals have excellent conductivity properties, which allow for the efficient flow of electricity throughout the device.

MODEL JUDGE:

Response 1 delves into the concept of Boolean values and Turing machines, which extends the topic to how computers "think" and process data. It links the use of transistors to Turing machines and emphasizes the binary nature of data processing, which aligns well with the reference answer's explanation of logical functions and binary systems accessible through transistors.

Response 2 focuses on the physical aspects of creating chips, such as the use of various metals and their conductivity. While it provides a general overview of chip construction and functionality, it lacks depth on the conceptual workings, like the function of transistors and binary systems, which is a significant aspect of the reference answer.

The reference answer provides an explanation centered around the role of transistors as switches and their organization to perform logical operations. It links this to the representation of numbers and data processing in computers through binary systems, which Response 1 addresses more comprehensively compared to Response 2.

Therefore, **Response 1** is better. (*The evaluation is correct!*)

1296 Table 19: Case Study For LFQA-E from Deepseek-V3.
12971298 **QUESTION:**1299 How did human calculate time?
13001301 **REFERENCE:**1302 For a long time we didn't and no one cared. However, once you start farming, you need to know
1303 what season it is, what hour of the day it is more or less, so you know what you can and should do
1304 on the fields for a bountiful yield. The easiest way to do is to **watch the celestial bodies, namely**
1305 **the two major ones, the Sun and the Moon** (the Stars also helped, but eventually it was more
1306 of a tool for navigation). Over time people figured out that the Sun shifts, days are not evenly
1307 distributed over the year, but luckily it goes in cycles, so it is predictable. Once you make a chart
1308 of these shifts throughout the year, you got yourself more or less a **calendar**. **Civilizations with**
1309 **accurate-enough calendars survived and thrived**. Calendars got more and more accurate over time
1310 as our math and science got better and here we are, keeping time down to **split seconds**. By the
1311 way, it does not fit perfectly. We accumulate quite the error over time, but the Gregorian Calendar
1312 is designed as a way to correct for these accumulative errors. The Persian Solar Calendar is the
1313 most accurate one, because the new year starts after a celestial event. Gregorian calendar has a
1314 bigger focus on religious events rather than any celestial event, so it needs corrections. We have
1315 a very accurate time measurement, but that is NOT what is used on the daily, because we don't
1316 need it.
13171318 **RESPONSE 1:**1319 Absolutely nothing about the way humans decided to calculate time "fits neatly" into Earth's
1320 behaviours. Years are actually 365 days plus 6 hours and 9 minutes, days are actually 23 hours
1321 and 56 minutes long. Other time systems have been attempted, with the Romans dividing day and
1322 night into 12 hours, which **varied in actual length as the days and nights shifted** through the year.
1323 They also had the "Civil Day" system: formal names given to time according to certain markers
1324 (such as separate named times for "rooster is crowing" and "rooster stops crowing", or "time we
1325 light candles" and "bed time"). Humans have been struggling with measuring it forever.
13261327 **RESPONSE 2:**1328 Time has been calculated and measured in various ways throughout history. One of the earliest
1329 methods used by ancient civilizations was **observing the movement of celestial bodies such as**
1330 **the sun, moon, and stars**. By tracking the positions of these objects in the sky, early astronomers
1331 were able to develop **calendars and timekeeping systems**. The modern concept of time mea-
1332 surement is based on the International System of Units (SI). In terms of dividing the day into
1333 **hours, minutes, and seconds**, this system was developed by the ancient Babylonians and Egyp-
1334 tians. The Babylonians used a base-60 system for counting, which led to the division of the hour
1335 into 60 minutes and the minute into 60 seconds. The Egyptians also used a similar system of
1336 timekeeping. Overall, **the calculation of time has evolved over centuries through a combination**
1337 **of astronomical observations, scientific advancements, and cultural practices**. The precise mea-
1338 surement of time today is achieved through the use of atomic clocks and coordinated universal
1339 time (UTC) to ensure accurate timekeeping worldwide.
13401341 **MODEL JUDGE:**1342 The reference answer emphasizes the historical development of time measurement, focusing on
1343 the transition from no precise timekeeping to the necessity of using celestial bodies for farming.
1344 Response 1 discusses the inconsistency and challenges of time measurements historically, noting
1345 the Roman system and difficulties faced by humans in tracking time over time. While it touches
1346 on early attempts to measure time, it doesn't delve specifically into celestial observations or the
1347 evolution of calendars as much.
13481349 Response 2 covers the historical aspect of using celestial bodies and tracks the progression to
1350 modern timekeeping with atomic clocks. It also discusses the Babylonian and Egyptian contribu-
1351 tions to dividing time. This response aligns more closely with the standard answer by tracing
1352 the transition from ancient methods to modern precision, even mentioning cultural influences on
1353 timekeeping.
13541355 Overall, while both responses provide historical context, Response 2 is closer to the reference
1356 answer as it better follows the progression of timekeeping from ancient observations leading to
1357 the precise systems we have today.
13581359 Therefore, **Response 2** is better. (*The evaluation is correct!*)
1360

1350
1351
1352
1353
1354
1355
1356
1357
1358

Table 20: Prompt used for LLM-based question filtering.

PROMPT FOR LLM FILTER	
Question Filtering Instructions	
Objective	Filter out questions that are either unclear in their description or too broad to provide a meaningful reference.
Filtering Criteria	
1. Unclear Questions	Reject questions that exhibit:
	<ul style="list-style-type: none"> * Ambiguous wording or phrasing * Multiple possible interpretations * Missing critical context or parameters * Vague or undefined terms * Grammatical issues that obscure meaning * Incomplete or fragmented thoughts
2. Overly Broad Questions	Reject questions that:
	<ul style="list-style-type: none"> * Request information on topics with no reasonable boundaries * Would require encyclopedic or book-length answers * Ask for opinions on vast, multi-faceted subjects * Lack specific focus or scope constraints * Would yield references too general to be useful * Cover multiple unrelated topics simultaneously
Process	
	<ol style="list-style-type: none"> 1. Read the question carefully and completely 2. Evaluate against both clarity and breadth criteria 3. Make a filtering decision:
	PASS: Question is clear and appropriately scoped
	REJECT - UNCLEAR: Question lacks clarity (provide specific reason)
	REJECT - TOO BROAD: Question is overly broad (provide specific reason)
Examples of Questions to Reject	
	<ul style="list-style-type: none"> * "What about technology?" (unclear) * "Explain everything about human history" (too broad) * "How does stuff work in general?" (both unclear and too broad) * "What are all the factors affecting everything in the world?" (too broad)
Examples of Questions to Pass	
	<ul style="list-style-type: none"> * "What is the boiling point of water at sea level?" * "How does photosynthesis work in green plants?" * "What were the main causes of World War I?"

1397
1398
1399
1400
1401
1402
1403

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

Table 21: Prompt for LLM Paraphrase.

PROMPT FOR LLM PARAPHRASE**Objective:**

Transform the provided response into a more verbose version while strictly preserving the original meaning and information.

Requirements:

- Expand the original text by adding descriptive language, elaborations, and explanatory phrases
- Maintain complete fidelity to the original information—do not introduce any new facts, claims, or insights
- Preserve the tone and intent of the original message
- Use stylistic techniques such as:
 - * Adding clarifying phrases and parenthetical explanations
 - * Employing more elaborate sentence structures
 - * Incorporating synonyms and varied vocabulary
 - * Adding transitional phrases between ideas
 - * Expanding brief points into full explanations
- Ensure the final text feels natural and not artificially inflated

Process:

1. Thoroughly analyze the original response to understand its complete meaning
2. Identify core points and supporting details
3. Expand each point methodically while maintaining the original structure
4. Review to confirm no new information has been introduced
5. Polish the text for readability and flow

Table 22: Prompt for English LLM Evaluation.

PROMPT FOR ENGLISH LLM EVALUATION**Objective:**

Transform the provided response into a more verbose version while strictly preserving the original meaning and information.

Requirements:

- Expand the original text by adding descriptive language, elaborations, and explanatory phrases
- Maintain complete fidelity to the original information—do not introduce any new facts, claims, or insights
- Preserve the tone and intent of the original message
- Use stylistic techniques such as:
 - * Adding clarifying phrases and parenthetical explanations
 - * Employing more elaborate sentence structures
 - * Incorporating synonyms and varied vocabulary
 - * Adding transitional phrases between ideas
 - * Expanding brief points into full explanations
- Ensure the final text feels natural and not artificially inflated

Process:

1. Thoroughly analyze the original response to understand its complete meaning
2. Identify core points and supporting details
3. Expand each point methodically while maintaining the original structure
4. Review to confirm no new information has been introduced
5. Polish the text for readability and flow

1458

1459

1460

1461

1462

1463

Table 23: Prompt for Prometheus Evaluation.

1464

PROMPT FOR PROMETHEUS EVALUATION

1466

Task Description:

1467

1468

An instruction (might include an Input inside it), a response to evaluate, and a score rubric representing a evaluation criteria are given.

1469

1470

1. Write a detailed feedback that assess the quality of two responses strictly based on the given score rubric, not evaluating in general.
2. After writing a feedback, choose a better response between Response A and Response B. You should refer to the score rubric.
3. The output format should look as follows: "Feedback: (write a feedback for criteria) [RESULT] (A or B)"
4. Please do not generate any other opening, closing, and explanations.

1475

1476

Instruction: {orig_instruction}
Response A: {response_A}
Response B: {response_B}
Score Rubric: {score_rubric}

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

Table 24: Prompt for Auto-J Evaluation.

1492

1493

1494

PROMPT FOR AUTO-J EVALUATION

1495

1496

1497

1498

1499

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511

You are a helpful and precise assistant for checking the quality of the feedback.

Two pieces of feedback have been provided for the same response to a particular query. Which one is better with regard to their correctness, comprehensiveness, and specificity to the query?

[BEGIN DATA]**[Query]:** {prompt}**[Response]:** {response}**[Feedback 1]:** {feedback1}**[Feedback 2]:** {feedback2}**[END DATA]**

Please choose from the following options, and give out your reason in the next line.

A: Feedback 1 is significantly better.

B: Feedback 2 is significantly better.

C: Neither is significantly better.

1512 **F ANNOTATION**
15131514 We show the annotation recipe in Table 25 and Table 26. We show a screenshot of annotation
1515 pipeline in Figure 5.
15161517 Table 25: Annotation recipe for deciding valid references of LFQA-E.
15181519 **ANNOTATION RECIPE FOR DROPPING INVALID REFERENCES**
15201521 **Goal**
1522 Keep only real, helpful explanation. Discard those that are uninformative, incorrect, or not actual expla-
1523 nations.
15241525 **Keep if the answer:**
1526

- Directly answers the question.
- Gives a simple but meaningful explanation (even if simplified).
- Is factually reasonable — not misleading or false.
- Stands alone (no “see link” or “I don’t know”).

1527 **Good example:** “We yawn to help cool the brain and stay alert. It brings in oxygen and improves blood
1528 flow.”
15291530 **Discard if the answer is:**
1531

- **Not an explanation** – e.g., “Google it,” “It’s magic,” jokes, memes, one-liners (“Because science”).
- **Irrelevant** – Doesn’t address the question or misunderstands it.
- **Factually wrong** – Clear misinformation (e.g., “Rain comes from clouds crying”).
- **Too vague** – No real content: “It’s complicated,” “There are many reasons.”
- **Avoids answering** – “Great question!”, “Not sure, but here’s a thought...” with nothing useful.
- **Circular** – Repeats the question: “We sleep because we’re tired.”
- **Inappropriate** – Offensive, harmful, or unprofessional content.

1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565

1566

1567

1568

1569

1570

1571

1572

1573

1574

1575

1576

1577

1578

1579

1580

1581

1582

1583

1584

1585

1586

1587

1588

1589

1590

1591

1592

1593

1594

1595

1596

1597

1598

1599

1600

1601

1602

1603

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619

Question:

What is an instigator in Chinese criminal law?

1. Identify Key Information in reference

Reference:

Overall, an instigator in Chinese criminal law refers to an individual who incites others to commit a crime. Specifically, such individuals intentionally lead or provoke others to decide to carry out criminal activities.
 Firstly, the fundamental characteristic of an instigator is that they incite others to commit a crime without directly participating in the criminal behavior themselves. They play a role in inducing others to form criminal intentions. For instance, a person may encourage, persuade, or offer some benefit to successfully prompt another to decide to commit acts such as theft or fraud while not partaking in the actual criminal process.
 Secondly, certain conditions must be met to establish instigation. Subjectively, the instigator must have the intention to induce others to form the criminal intent and decide to commit the crime, which is known as the intent to instigate others to commit a crime. This intent must be explicit; the instigator must clearly know whom they are instigating and what crime they are instigating. Without explicit intent, instigation cannot be established; behavior that inadvertently causes others to form criminal intentions cannot constitute instigation. Objectively, the instigator must have carried out acts of instigation, commonly manifested as encouragement, deception, persuasion, requests, bribery, or coercion, prompting specific individuals to commit specific crimes. Whether the instigatee actualizes the criminal intent or determination and whether the instigatee actually carries out the criminal act do not affect the establishment of the instigator.
 Regarding the criminal liability of instigators, several key points are noteworthy. First, instigators are punished according to their role in the joint crime. If the instigator plays a major role in the crime process, they will be punished as the main perpetrator; if they only play a minor or auxiliary role, they will be punished according to the standard for accomplices. Although instigators typically play a major role and are generally punished as main perpetrators, it is possible for an instigator with a minor role to be punished as an accomplice. Second, if the instigatee does not commit the instigated crime, the instigator may still be deemed guilty but subject to lighter or reduced punishment. This scenario, usually referred to as "attempted instigation," indicates that attempted instigation is also punishable, underscoring the independent and punishable nature of instigative behavior under Chinese criminal law. Third, instigating minors to commit crimes should result in harsher punishment, especially in cases where those under 18 are incited to commit criminal acts.
 Though instigators possess independent criminality or punishability, they are not an independent crime category. The instigator should be charged according to the specific crime instigated. For example, if the instigation pertains to robbery, the instigator should be prosecuted under robbery charges. In summary, through these three points, we can gain a deeper understanding of the provisions concerning instigators in Chinese criminal law and clarify their legal responsibilities.
 Of course, this summary does not cover all possible scenarios, and specific legal determinations in practice require detailed analysis based on individual cases. However, these basic principles provide a relatively comprehensive understanding of instigators. These principles stem from Chinese criminal law and reflect our serious attitude towards and strict control over instigative behavior.

answer a:

An instigator in Chinese criminal law refers to a person who directly or indirectly instructs or incites others to commit a crime. According to the relevant provisions of the Criminal Law of the People's Republic of China, an instigator not only bears the corresponding criminal responsibility for the criminal act but may also be punished more severely based on their role and influence in the criminal process. For example, if a person induces and instructs others through an online platform, leading to the commission of a fraud crime by the incited person, the instructing behavior, although not directly involved in the fraud, may be recognized as an instigator, bearing corresponding criminal responsibility. In legal practice, the determination of an instigator generally requires a comprehensive judgment based on the specific circumstances of the case, including the nature, content, and result of the instruction, as well as the identity of the instructing person.

answer b:

2. Check for Key Information in Answers

According to Chinese criminal law, an instigator is a person who deliberately entices or incites others to commit criminal behavior through oral, written, or other means. The instigator usually does not directly commit the crime but their actions lead directly to the commission of the crime by others. In criminal law, an instigator, like a principal or accomplice, must bear corresponding legal responsibility. For example, if someone incites another person online to rob a store, and the incited person commits the robbery, the instigator constitutes an instigator.

3. Handling Responses

The handling responses criteria classify answers as totally correct, partially correct, or totally wrong, affecting the overlapping information score accordingly. Use purple to display common points.

4. Compare Overlapping Information

Which will be considered the better response?

Answer A

Answer B

Figure 5: The annotation pipeline.

1620
1621
1622

Table 26: Annotation recipe of LFQA-E.

1623

Overview

1624

This guide helps annotators evaluate and compare long-form responses against a reference to determine which response is more informative and complete. The process uses a triple-choice format (Response A Better, Response B Better, or Tie).

1625

Key Principles

1626

- Focus on **factuality** and **completeness** according to the reference
- Fluency is not a primary evaluation criterion (all responses are expected to be fluent)
- Use information units as the basic evaluation unit
- Minimize bias through systematic comparison

1627

Prerequisites

1628

- Domain knowledge relevant to the question topic
- Understanding of the subject matter through academic coursework or professional experience
- Ability to maintain focus during paragraph-level analysis

1629

Evaluation Process

1630

Step 1: Extract Key Information from Reference

1631

1. Read the question carefully to understand what information is being requested
2. Read the reference thoroughly
3. Identify and list all key information units that:
 - Directly answer the question
 - Provide necessary context or background
 - Support the main answer with evidence or examples
4. Organize key information into logical categories or themes

1632

Step 2: Check for Key Information in Responses

1633

For each response (A and B):

1634

1. Read the response completely
2. Map each key information unit from the reference to the response
3. Mark which key information units are:
 - Present and accurate
 - Present but inaccurate
 - Missing entirely
4. Note any additional information not in the reference

1635

Step 3: Handle Response Content

1636

1. Evaluate additional information:
 - Is it relevant to the central topic?
 - Does it enhance understanding or is it verbose/unnecessary?
2. Identify intertwined information:
 - For sentences containing both correct and incorrect information, separate the components
 - Assess the impact of any inaccuracies on the overall response quality

1637

Step 4: Compare Overlapping Information

1638

1. Compare how well each response covers the key information units
2. Consider:
 - Completeness: Which response includes more key information?
 - Accuracy: Which response presents information more correctly?
 - Relevance: Which response stays more focused on the question?
3. Compare the quality of overlapping information presentation

1639

Step 5: Make Final Decision

1640

Select one of three options:

1641

- **Response A is Better:** A contains more key information and/or presents it more accurately
- **Response B is Better:** B contains more key information and/or presents it more accurately
- **Tie:** Both responses are comparable in information coverage and accuracy

1642

Common Pitfalls to Avoid

1643

1. Losing focus due to long paragraphs - use the systematic approach
2. Allowing domain bias to influence decisions - stick to the reference
3. Confusing eloquence with accuracy
4. Missing subtle differences between comparable responses

1644