

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 NOT SEARCH, BUT SCAN: BENCHMARKING MLLMs ON SCAN-ORIENTED ACADEMIC PAPER REASONING

005 **Anonymous authors**

006 Paper under double-blind review

## 009 ABSTRACT

011 With the rapid progress of multimodal large language models (MLLMs), AI al-  
 012 ready performs well at literature retrieval and certain reasoning tasks, serving as  
 013 a capable assistant to human researchers, yet it remains far from autonomous re-  
 014 search. The fundamental reason is that current work on scholarly paper reason-  
 015 ing is largely confined to a search-oriented paradigm centered on pre-specified  
 016 targets, with reasoning grounded in relevance retrieval, which struggles to sup-  
 017 port researcher-style full-document understanding, reasoning, and verification. To  
 018 bridge this gap, we propose ScholScan, a new benchmark for scholarly paper rea-  
 019 soning. ScholScan introduces a scan-oriented task setting that asks models to read  
 020 and cross-check entire papers like human researchers, scanning the document to  
 021 identify consistency issues. The benchmark comprises 1,800 carefully annotated  
 022 questions drawn from 9 error families across 13 natural-science domains and 715  
 023 papers, and provides detailed annotations for evidence localization and reasoning  
 024 traces, together with a unified evaluation protocol. We assessed 15 models across  
 025 24 input configurations and conduct a fine-grained analysis of MLLM capabilities  
 026 across error families. Across the board, retrieval-augmented generation (RAG)  
 027 methods yield no significant improvements, revealing systematic deficiencies of  
 028 current MLLMs on scan-oriented tasks and underscoring the challenge posed by  
 029 ScholScan. We expect ScholScan to be the leading and representative work of the  
 030 scan-oriented task paradigm.

## 032 1 INTRODUCTION

034 Scientific papers are crystallizations of human intelligence. Enabling multimodal large language  
 035 models (MLLMs) (OpenAI, 2025; Anthropic, 2025; ByteDance Seed Team, 2025; Meta, 2025;  
 036 xAI, 2025) to conduct comprehensive understanding and generation based on academic literature  
 037 is the ultimate goal of Deep Research, and a critical milestone on the path toward artificial gen-  
 038 eral intelligence (AGI) (Ge et al., 2023; Morris et al., 2024; et al., 2025c). With rapid advances,  
 039 MLLMs are increasingly capable of supporting academic workflows through retrieval, reading, and  
 040 writing. For example, PaSa (He et al., 2025) can invoke a series of tools to answer complex aca-  
 041 demic queries with high-quality results, while Google Deep Research (et al., 2025b) is capable of  
 042 producing human-level research reports based on specific queries.

043 However, most of the existing work still follows *a search-oriented paradigm*, where models re-  
 044 trieve a few relevant passages and reason over local evidence based on prespecified targets (Gao  
 045 et al., 2023; Lou et al., 2025). Such methods are effective for tasks with clearly predefined tar-  
 046 gets, but struggle with researcher-style full-document reasoning and verification (Zhou et al., 2024).  
 047 **To function as researchers, models must move beyond reactive question answering and toward**  
 048 **proactive discovery of implicit problems.**

049 To fill this gap, as shown in Figure 1, we introduce *a scan-oriented paradigm*, where models address  
 050 queries with targets absent and are required to actively **construct a document-level evidence view**,  
 051 **perform exhaustive scanning over the full paper, and conduct evidence-based reasoning**. In  
 052 contrast to search-oriented tasks that assess a model’s ability to identify and reason over *relevant*  
 053 fragments, scan-oriented tasks emphasize *consistency*. **Instead of relying on prespecified targets or**  
**hints, models must derive all necessary concepts and inferences solely from given documents.**

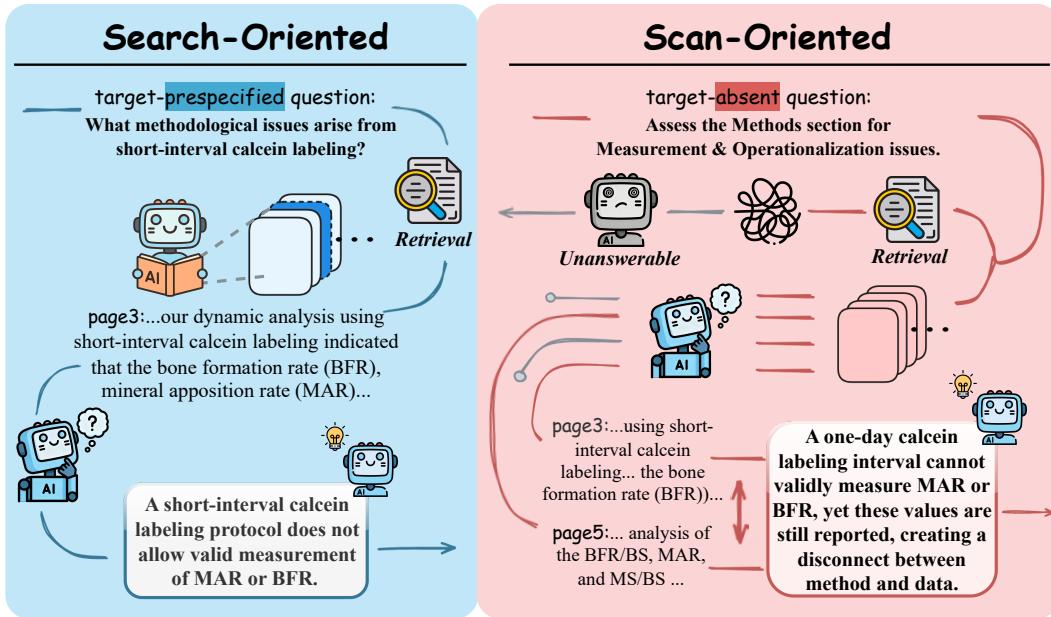


Figure 1: A comparison between search-oriented and scan-oriented task paradigms. Unlike the former, the scan-oriented paradigm provides no prespecified targets, requiring the model to actively scan the entire paper, construct a document-level evidence view.

We instantiate this setting via scientific error detection, as it naturally demands discovering non-obvious flaws without target cues, and present ScholScan, a new multimodal benchmark for scholarly reasoning. ScholScan features the following key highlights:

- **Scan-Oriented Task Paradigm.** ScholScan receive one or more complete academic papers together with target-absent queries, presenting a rigorous challenge to their evidence-based reasoning capabilities. The benchmark comprises 715 papers spanning 13 natural science disciplines.
- **Comprehensive Error Types.** ScholScan covers 9 categories of scientific errors across the entire research workflow. It also includes citation and referencing errors, providing a rigorous test of a model’s cross-source reasoning ability.
- **Process-Aware Evaluation Framework.** ScholScan provides fine-grained annotations for both evidence location and reasoning steps, enabling a comprehensive evaluation framework that assesses model performance in terms of both process and outcome.

We evaluate 15 models across 24 input configurations and 8 retrieval-augmented generation (RAG) frameworks. All models exhibit limited performance, and none of the RAG methods deliver significant improvements. These results highlight the inadequacy of search-oriented frameworks when applied to scan-oriented tasks, and underscore both the challenges and the potential of enabling MLLMs to perform reliable, document-level reasoning over full academic papers.

## 2 RELATED WORK

### 2.1 MULTIMODAL LARGE LANGUAGE MODELS

With the rapid progress of MLLMs, models have evolved beyond perception tasks (e.g., image recognition and explanation) (Liu et al., 2024) toward deep understanding of structured, multimodal long documents. Their strengths lie in the ability to integrate cross-modal information and perform multi-hop reasoning over extended contexts. These capabilities are not only valuable for specific question answering or instruction-following tasks (Yue et al., 2024) but are particularly well suited for simulating human thought processes and generating explainable reasoning trajectories (Zheng et al., 2023). Consequently, achieving comprehensive understanding of entire documents has emerged as a core challenge that MLLMs are inherently equipped to address.

108  
109

## 2.2 DOCUMENT UNDERSTANDING BENCHMARK

110  
111  
112  
113  
114  
115  
116  
117  
118  
119

Document understanding tasks challenge models to identify relevant context and perform accurate reasoning grounded in that information. Progress in document understanding benchmarks has followed two main axes. Along the input dimension, it has evolved from short to long contents, from everyday to specialized domains, and from plain text to multimodal format (Chen et al., 2021; Yang et al., 2018; Tito et al., 2021; Deng et al., 2025). Along the scenario dimension, it has shifted from limited-output formats to more open-ended responses (Pramanick et al., 2024). DocMath-Eval (Zhao et al., 2024) evaluates numerical reasoning on long, specialized documents, revealing large performance gaps even for strong models in expert domains, while MMLongBench-Doc (Ma et al., 2024) builds a multimodal benchmark with layout-rich documents. However, a comprehensive benchmark that integrates all challenges above has yet to be introduced.

120  
121

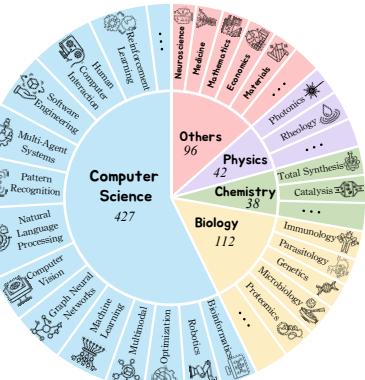
## 2.3 ACADEMIC PAPER UNDERSTANDING BENCHMARK

122  
123  
124  
125  
126  
127  
128  
129  
130  
131  
132

Compared with general documents, academic papers are distinguished by their rich domain knowledge and logical rigor. Reasoning over papers has emerged as a major challenge in recent research. Some studies ask for local elements like charts or snippets, leveraging their internal complexity, but neglect the need for cross-source integration and domain-specific interpretation within the full document (Wang et al., 2024; Li et al., 2024). Recent studies extend inputs to the document level and adopt image-based formats to better simulate real-world reading scenarios. (Auer et al., 2023; Yan et al., 2025) However, benchmarks based on the QA paradigm face inherent limitations, as they typically presuppose answer existence and embed explicit cues in the question itself, reducing the need for comprehensive understanding and information organization. Moreover, mainstream evaluation protocols focus on the final outcome, with limited assessment of whether intermediate reasoning is evidentially grounded and logically valid. More examples and analysis are shown in Appendix C.

133  
134

## 3 THE SCHOLEVAL BENCHMARK

135  
136  
137  
138  
139  
140  
141  
142  
143  
144  
145  
146  
147  
148

Benchmark	Mod.	Para.	Eval.	# Dom.
<i>Document Understanding</i>				
DocMath-Eval	T+TD	Search	A	N/A
CompLong	T+MD	Search	A	N/A
MMLongbench-Doc	T+MD	Search	A	N/A
LongDocURL	T+MD	Search	A	N/A
SlideVQA	T+MD	Search	A	N/A
<i>Academic Paper Understanding</i>				
CharXiv	I	Search	A	8
ArXivQA	I	Search	A	10
MMCR	T+MD	Search	A	CS
AAAR-1.0	T+MD	Search	A	CS
<b>ScholScan (ours)</b>	T+MD	Scan	A+P	13

Figure 2: Left: Overview of ScholScan. Right: Comparison to related benchmarks. **Mod.**: Modalities; **Para.**: Task Paradigm; **Eval.**: Evaluation; **T**: Text; **I**: Image; **TD**: Text-Form Document; **MD**: Multimodal Document; **A**: Answer; **P**: Process; **Dom**: Number of academic domains in the dataset.

152  
153

## 3.1 OVERVIEW OF SCHOLSCAN

154  
155  
156  
157  
158  
159  
160  
161

We introduce ScholScan, a benchmark designed to comprehensively evaluate MLLMs’ ability to detect scientific flaws in academic papers under scan-oriented task settings. As illustrated in Figure 2, ScholScan spans 13 disciplines across the natural sciences, including physics, chemistry, and computer science, and spans over 100 subfields such as immunology, total synthesis, and machine learning. The benchmark comprises 1,800 questions derived from 715 real academic papers, and covers 9 major error categories (Figure 3) that commonly observed in real-world research scenarios. These include issues in numerical and formulaic computation, experimental design, inference and conclusion, and citation misuse, among others. Figure 2 also provides a comparison ScholScan with existing benchmarks for multimodal paper understanding and long-document reasoning.

162	Research Question & Definitions	163	Design & Identifiability	164	Sampling & Generalizability
165		166		167	
168	<b>Explanation:</b> The definition of "actionable variants" shifts across sections (LOE 1–5 in Abstract, LOE 1–3 in Results), causing ambiguity.	169	<b>Explanation:</b> The design is described as probing both short- and long-range interactions, yet the paper still claims unique large- $q$ selectivity, creating a disconnect.	170	<b>Explanation:</b> The experiments use a narrow diabetic mouse substrain, yet the paper generalizes findings to all patients, creating an invalid sample-to-population inference.
171		172		173	
174	Measurement & Operationalization	175	Data Handling & Preprocessing	176	Computation & Formulae
177		178		179	
180	<b>Explanation:</b> First-harmonic demodulation is dominated by far-field background and cannot produce the reported high-quality near-field images.	181	<b>Explanation:</b> Feature selection for NSCLC and HCC models was done on the full dataset before splitting, causing data leakage, while the Discussion falsely claims unbiased validation.	182	<b>Explanation:</b> The Methods claim a 200-fold concentration, but the 200 $\mu\text{L}$ subsample is incorrectly said to represent $\sim 20 \text{ mL}$ instead of 40 $\text{mL}$ , creating a twofold calculation error.
183		184		185	
186	Inference & Conclusions	187	Referential & Citation Alignment	188	Language & Expression
189		190		191	
192	<b>Explanation:</b> The data show PGK1 promotes EGFR degradation, yet the Discussion claims inhibiting PGK1 as therapy, directly contradicting the results.	193	<b>Explanation:</b> Figure 1 report an LPS dose of 1.5 mg/kg, but Figure 5 reports 15 mg/kg, creating a tenfold discrepancy that makes the actual experimental dose unclear.	194	<b>Explanation:</b> The paper swaps <i>C. elegans</i> gene and protein nomenclature (e.g., 'unc-45' vs. 'UNC-45'), creating technically misleading references.
195		196		197	

Figure 3: Sampled ScholScan examples with 9 error types, covering the whole process of scientific research, each requiring the model to perform thorough cross-source evidence-based reasoning.

### 3.2 DATA COLLECTION & QUESTION GENERATION

We curated papers from ICLR 2024/2025 and Nature Communications, and collected public reviews for the former. Questions were constructed based on two dimensions, where the source is either generated or sampled, and the context is either within-paper or cross-paper.

**Generation.** On high-quality accepted papers, we prompt Gemini 2.5 Pro to perform coordinated sentence-level edits spanning multiple sections or pages. It then synthesizes composite errors and generates the corresponding question along with an explanation grounded in the edited context.

**Sampling.** From rejected ICLR submissions and their public reviews, we prompt Gemini 2.5 Pro to extract explicit, falsifiable scientific errors and convert them into questions with initial explanations. Subjective remarks about novelty or writing quality are excluded.

**Within-paper.** This setting focuses on verifiable facts and internal consistency within a single paper, and supports both Generation and Sampling.

**Cross-paper.** This setting examines citation consistency across papers. For each instance, Gemini 2.5 Pro receives an accepted paper and one of its cited sources, then edits the accepted paper to

introduce paraphrases or reasoning errors about the citation. As public reviews mainly address nonfalsifiable aspects such as appropriateness, all cross-paper instances are constructed exclusively using the generation method.

### 3.3 QUALITY CONTROL & ANNOTATION

Despite explicit instructions, initial outputs exhibited substantial hallucinations, logical inconsistencies, and low-quality questions. To ensure the quality, 10 domain experts conducted a rigorous annotation process. Each instance underwent independent dual review, and disagreements were resolved by a third expert. Among the 3,500 initially generated candidates, 1,700 were discarded, and 1,541 of the remaining were revised, including 535 question rewrites, 1,207 explanation edits, and 1,141 corrections to error categories or metadata. Further details are provided in Appendix D.

## 4 EXPERIMENTS

### 4.1 EXPERIMENTS SETTING

**Models.** We benchmark a total of 24 input configurations by feeding academic papers as either images or OCR text using the Tesseract (Smith, 2007) engine, covering 15 mainstream models (Yang et al., 2025; Bai et al., 2025; et al., 2025a; Guo et al., 2025; et al., 2025d).

**Evaluation Protocol.** Inspired by MMLongBench-Doc (Ma et al., 2024), we prompt models to generate necessary reasoning chains from evidence to detected anomalies without constraining the output format, which aims to assess the ability for evidence-grounded reasoning rather than mere instruction-following. For open-ended responses, we use GPT-4.1 (OpenAI, 2025) to extract cited evidence and reasoning steps, and quantify alignment with annotated explanations. Human evaluation confirms high agreement between our pipeline and expert annotations. Further implementation details are provided in Appendix F.

**Metrics.** We define a structured evaluation framework by parsing the model response  $a$  into a tuple:

$$\Psi(a) \Rightarrow (\mathbf{1}_{\text{exist}}, \mathbf{1}_{\text{contain}}, \hat{\mathcal{E}}, \hat{\mathcal{R}}, n). \quad (1)$$

Here,  $\mathbf{1}_{\text{exist}}$  and  $\mathbf{1}_{\text{contain}}$  are binary indicators for whether output contains any error and includes the annotated target error;  $\hat{\mathcal{E}}, \hat{\mathcal{R}}$  and  $\mathcal{E}^*, \mathcal{R}^*$  are the predicted and gold evidence sets and reasoning chains;  $\hat{g} = \text{prefix\_match}(\hat{\mathcal{R}}, \mathcal{R}^*)$  counts matched reasoning steps;  $n \in \mathbb{N}$  is the number of unrelated errors.  $\text{HasError}(a)$  is 1 if the output contains any predicted error, and 0 otherwise. Based on  $\Psi(a)$ , we define an end-to-end score  $S(m) \in [0, 1]$  that combines all aspects of prediction quality:

(i) *Existence.*  $S_{\text{exist}}(a) = 1$  if and only if the response includes the annotated target error.

$$S_{\text{exist}}(a) = \mathbf{1}\{\text{HasError}(a)\} \cdot \mathbf{1}\{\hat{\mathcal{E}} \cap \mathcal{E}^* \neq \emptyset\} \quad (2)$$

(ii) *Evidence location score.* Even when the target error is identified, the cited evidence may be incomplete or noisy. We compute a Dice score with a squared penalty for over-reporting:

$$S_{\text{location}} = \max\left\{0, \frac{2|\hat{\mathcal{E}} \cap \mathcal{E}^*| + \mathbf{1}\{|\hat{\mathcal{E}}| + |\mathcal{E}^*| = 0\}}{\max(|\hat{\mathcal{E}}| + |\mathcal{E}^*|, 1)} - 0.8 \left(\frac{|\hat{\mathcal{E}} \setminus \mathcal{E}^*|}{\max(|\hat{\mathcal{E}}|, 1)}\right)^2\right\}. \quad (3)$$

(iii) *Reasoning process score.* Even if the target error is detected, the reasoning may diverge from the gold chain. We use prefix match to assess reasoning completeness:

$$S_{\text{reasoning}} = \mathbf{1}\{g_r = 0\} + \mathbf{1}\{g_r > 0\} \left(\frac{\hat{g}}{g_r}\right)^2. \quad (4)$$

(iv) *Unrelated-error penalty.* Models may list unrelated items to inflate recall at the cost of precision. We penalize this with a rapidly increasing function of unrelated error count:

$$P_{\text{unrelated\_err}}(n) = 0.9^{\min(n, 2)} \exp\left(-0.6 [\max(n - 2, 0)]^{1.5}\right). \quad (5)$$

(v) *Overall outcome score.* The final score for  $a$  is defined as:

$$S(m) = S_{\text{exist}}(a) \sqrt{S_{\text{location}} \cdot S_{\text{reasoning}}} \cdot P_{\text{unrelated\_err}}(n). \quad (6)$$

270 Table 1: Model performance (scaled by 100) across input configurations. **RQD**: Research Question  
 271 & Definitions; **DI**: Design & Identifiability; **SG**: Sampling & Generalizability; **MO**: Measurement  
 272 & Operationalization; **DHP**: Data Handling & Preprocessing; **CF**: Computation & Formulae; **IC**:  
 273 Inference & Conclusions; **RCA**: Referential and Citation Alignment; **LE**: Language & Expression.

Models	Avg.	RQD	DI	SG	MO	DHP	CF	IC	RCA	LE
<b>MLLM (Image Input)</b>										
<i>Proprietary LLMs</i>										
Gemini 2.5 Pro	<u>15.6</u>	<b>11.9</b>	<b>12.6</b>	<b>35.7</b>	<u>12.3</u>	<b>27.0</b>	4.6	14.7	<u>15.2</u>	<b>7.4</b>
GPT-5	<b>19.2</b>	<u>10.1</u>	<u>9.7</u>	28.2	<u>14.6</u>	<u>26.6</u>	<b>13.8</b>	<b>25.3</b>	<b>25.3</b>	<u>6.9</u>
Grok 4	4.0	0.0	1.9	16.7	3.2	7.4	0.7	1.9	3.6	0.0
Douba-Seed-1.6-thinking	10.2	3.4	3.5	22.3	7.5	15.1	<u>10.2</u>	12.2	10.9	3.3
Douba-Seed-1.6	9.9	3.0	4.4	<u>29.2</u>	4.9	15.0	<u>6.3</u>	<u>17.9</u>	8.0	3.9
<i>Open-source LLMs</i>										
Llama 4 Maverick	7.0	7.0	7.3	9.4	4.5	4.0	6.5	6.7	8.8	3.0
Gemma 3 27B	1.7	0.5	2.7	2.3	1.7	1.0	1.0	1.3	2.6	0.0
Mistral Small 3.1	3.3	0.1	2.0	2.0	1.5	0.1	1.0	2.2	8.6	1.0
Qwen2.5 VL 72B	0.1	0.0	0.7	0.0	0.0	0.0	0.0	0.0	0.2	0.0
<b>OCR + LLM (Text Input)</b>										
<i>Proprietary LLMs</i>										
Gemini 2.5 Pro	<b>30.3</b>	<b>21.5</b>	<b>34.2</b>	<b>44.3</b>	<b>27.6</b>	<b>56.6</b>	<b>10.3</b>	<u>28.8</u>	<b>35.6</b>	<b>8.1</b>
GPT-5	<u>22.5</u>	<u>16.1</u>	<u>21.4</u>	26.0	<u>20.3</u>	<u>36.7</u>	4.7	<b>29.8</b>	30.0	2.6
Claude Sonnet 4	5.7	3.7	2.5	10.8	4.3	10.3	1.4	8.4	6.6	3.5
Grok 4	20.8	9.3	7.7	<u>37.4</u>	12.3	34.4	<u>9.0</u>	20.0	<u>31.2</u>	<u>7.2</u>
Douba-Seed-1.6-thinking	15.3	8.2	10.1	24.3	10.1	24.2	6.4	19.2	21.0	4.2
Douba-Seed-1.6	13.9	5.4	6.9	26.4	10.3	23.6	6.3	20.1	17.5	2.3
<i>Open-source LLMs</i>										
Qwen3 A22B (Thinking)	17.4	8.9	16.2	31.9	15.1	23.7	5.6	22.3	21.1	2.3
Qwen3 A22B	1.7	1.2	0.0	2.7	0.4	1.0	0.1	4.3	2.5	1.1
gpt-oss-120b	7.3	6.3	5.7	18.3	4.9	14.5	1.6	12.5	5.5	0.0
DeepSeek-R1	11.4	5.1	11.9	25.4	8.7	22.5	4.7	16.3	9.8	3.5
DeepSeek-V3.1	1.7	1.2	2.0	1.7	1.0	5.8	0.5	2.2	2.1	0.0
Llama 4 Maverick	2.3	1.5	2.0	4.8	3.0	3.6	0.0	5.8	1.6	0.2
Gemma 3 27B	2.0	2.1	1.6	3.0	2.7	0.2	0.7	7.7	1.0	0.0
Mistral Small 3.1	6.9	3.0	2.7	5.5	7.0	2.0	8.5	4.0	12.2	3.0
Qwen2.5 VL 72B	0.2	0.0	0.7	0.0	0.0	0.0	0.0	0.0	0.6	0.0

## 4.2 MAIN RESULT

Table 1 presents our evaluation results. Our main findings are summarized as follows:

**Overall performance remains unsatisfactory.** GPT-5 achieves the highest average score in the image input group (19.2), while Gemini 2.5 Pro, the best-performing model in the text input setting, still fails to surpass the 60-point threshold on any subtask. Even in the SG category, which yields the best performance overall, nearly half of the models receive single-digit scores. Most models perform poorly under the scan-oriented task formulation and fail to detect any issues in many papers. This challenge is particularly pronounced for open-source models.

**Reasoning-enhanced models demonstrate clear advantages.** Across both input configurations, reasoning-enhanced variants consistently achieve higher scores. Almost all top-performing models, measured by both subtask-specific and overall metrics, fall into this category. Notably, Qwen3-Thinking and Deepseek-R1 outperform their base versions by more than 10% in average scores, with substantial gains observed across all error types. These results indicate that reasoning-enhanced models are better able to simulate the iterative process of extraction followed by reasoning, which is essential for effectively handling scan-oriented tasks and producing higher-quality responses.

**MLLMs face significant bottlenecks in handling long multimodal inputs.** Across most evaluation metrics, text inputs outperform image inputs. Among the nine MLLMs tested, the average performance gap between text and image inputs reaches 4.81 points, highlighting visual processing as a key limitation in current MLLM capabilities.

324 **In most evaluation metrics, text inputs consistently outperform image inputs.** Among the nine  
 325 MLLMs evaluated, the average performance gap between text and image inputs is 4.81 points, un-  
 326 derscoring visual processing as a key limitation in current MLLM capabilities.  
 327

328 **Although overall performance is generally weaker, multimodal input remains indispensable.**  
 329 In certain categories such as CF, where OCR-based text extraction leads to substantial loss of formu-  
 330 laic or tabular content, image inputs outperform their text counterparts. This highlights the essential  
 331 role of multimodal reasoning and the irreplaceable value of visual information in addressing specific  
 332 types of errors.  
 333

### 334 4.3 FINE-GRAINED ANALYSIS

335 **Capability Dimensions.** We compute pairwise  
 336 Spearman correlations between error types across two  
 337 input configurations (text and image) for the eight  
 338 evaluated MLLMs excluding Qwen2.5-VL-72B, as  
 339 shown in Figure 4. We derive the following insights:  
 340

341 *(i) With image input, CF exhibits consistently low cor-  
 342 relations with other error categories, suggesting that  
 343 the skills required for mathematical reasoning are rel-  
 344 atively distinct.* In contrast, with text input, CF shows  
 345 moderate correlation with LE, indicating that OCR-  
 346 flattened formulas lose their structural specificity and  
 347 are interpreted by models in a manner more akin to  
 348 natural language. Combined with the overall poor  
 349 performance on CF tasks, this underscores the unique  
 350 challenges of this category and the need for targeted  
 351 improvements.  
 352

353 *(ii) Although DI is also related to experimental set-  
 354 tings, it does not exhibit strong correlations with SG,  
 355 MO, or DHP.* This indicates that DI primarily emphasizes causal framing and variable identifiability,  
 356 rather than the procedural understanding of experimental operations.  
 357

358 *(iii) OCR severely degrades structured content such as figures and formulas, making questions that  
 359 depend on multimodal information unanswerable.* This diminishes the expression of multimodal  
 360 reasoning capabilities and artificially inflates inter-category correlations under text input.  
 361

362 Based on the above analysis, we consolidate the original 9 error categories, each defined by its objec-  
 363 tive target, into 5 core latent skill dimensions evaluated by ScholScan under the image input setting.  
 364 While each dimension highlights the primary competence emphasized by its corresponding error  
 365 types, they are not mutually exclusive, as many questions involve overlapping reasoning abilities.  
 366

367 RQD and DI correspond to research concept comprehension, which requires models to **identify the**  
 368 **scope and definition** of research objectives by integrating contextual cues and prior knowledge.  
 369 SG, MO, and DHP fall under **experimental process modeling**, which tests a model’s ability to  
 370 reconstruct procedural workflows such as sampling, measurement, and data handling. CF captures  
 371 **formal reasoning and symbolic computation**, focusing on syntactic parsing and numerical logic. IC  
 372 evaluates causal inference, where models must **synthesize dispersed causal evidence** to reach sound  
 373 conclusions. RCA and LE reflect referential alignment and linguistic consistency, which assess the  
 374 ability to **verify citations and maintain coherent expression** throughout the document.  
 375

376 **Hidden Complexity in Scan-Oriented Tasks.** We analyze the reasoning traces of GPT-5 and Gemini  
 377 2.5 Pro under both input configurations, focusing on the number of evidence pieces scanned and  
 378 the reasoning steps performed. As illustrated in Figure 5, even the most advanced models often scan  
 379 up to 8 times more evidence and execute 3.5 times more reasoning steps than the reference answers,  
 380 merely to approximate a correct response, yet they still frequently fail. This highlights the substan-  
 381 tial hidden complexity inherent in scan-oriented tasks, which significantly amplifies the challenge  
 382 of successful task completion.  
 383

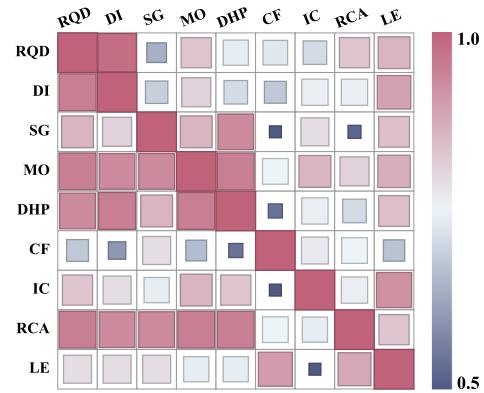


Figure 4: Spearman correlation matrix among the 9 error types.

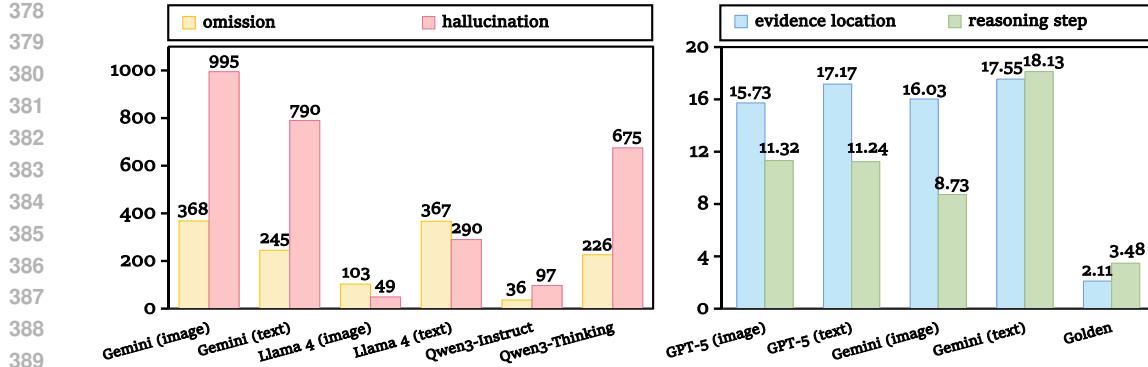


Figure 5: Left: Distribution of omission and hallucination errors. Right: Average reasoning steps and evidence locations involved in the answer generation, compared against the golden reference.

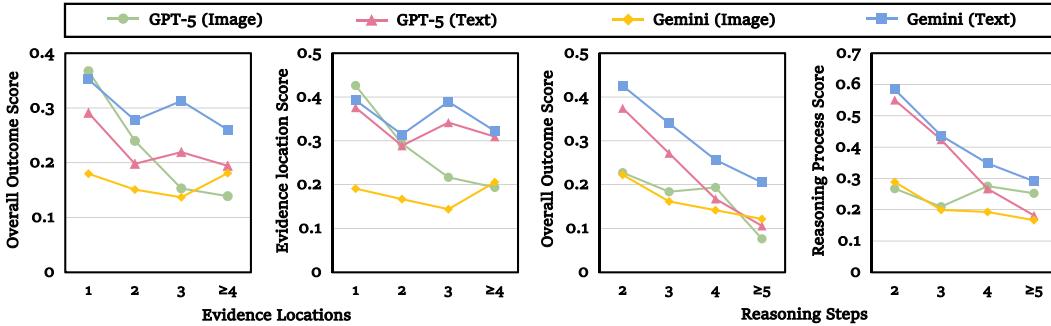


Figure 6: Performance trends across varying reasoning depths and evidence counts.

#### 4.4 ERROR ANALYSIS

**Omission and Hallucination.** Most zero-score cases fall into two categories: either the model fails to detect any errors in the paper, or it becomes overwhelmed by hallucinations and entirely overlooks the actual errors present in the reference answer. We analyze the number of zero-score questions and the proportion of these two failure modes across models, as shown in Figure 5. Stronger models tend to have fewer zero-score cases overall, but are more prone to overconfident hallucinations.

**Fragile Reasoning under Complex Evidence.** Figure 6 shows how top-performing models behave under different numbers of reasoning steps and evidence locations. As reasoning steps increase, both reasoning and overall scores steadily decline, revealing a clear bottleneck in MLLMs’ ability to construct long causal chains. In contrast, variation in evidence count has a weaker and less consistent impact. However, this does not imply that multi-evidence questions pose only marginal difficulty. Since the evaluation metric allows partial evidence omissions, more evidence items do not necessarily incur large score penalties. Still, heavier evidence loads often require longer reasoning chains, which substantially affect the coherence and completeness of inferred logic. These results highlight the persistent challenge for MLLMs in integrating evidence and maintaining logical structure as task complexity grows.

#### 4.5 RAG ANALYSIS

We evaluated 8 RAG methods under both input configurations (Robertson et al., 1994; Chen et al., 2024; Lee et al., 2025; Faysse et al., 2025; Yu et al., 2025; Wang et al., 2025; Izacard et al., 2022). Key findings are presented below, with detailed results shown in Tables 2 and 3.

**Oracle Condition Yields Significant Accuracy Gains.** Providing gold-standard images alleviates the scanning burden in long-context inputs, increasing the chances of generating correct answers. While overall performance improves, gains are limited for CF errors and minimal for LE errors. For

Table 2: Scores of RAG methods across the 9 error types (scaled by 100).

Models	Avg	RQD	DI	SG	MO	DHP	CF	IC	RCA	LE
<i>Text Input (Base Model: Qwen3 Thinking)</i>										
Baseline	17.4	8.9	16.2	31.9	15.1	23.7	5.6	22.3	21.1	2.3
Oracle	24.5	20.6	27.9	43.6	21.3	40.8	7.4	26.9	26.0	1.9
bm25	16.7	9.7	13.7	33.0	17.3	23.8	6.8	25.4	16.5	3.0
BGE-M3	11.3	8.6	7.5	24.8	9.1	15.4	5.3	15.6	11.4	1.0
Contriever-msmacro	16.6	9.7	18.2	33.7	10.7	20.8	6.4	18.5	19.8	1.8
nv-embed-v2	6.8	4.0	4.0	9.4	6.1	4.9	5.5	5.7	10.0	2.0
<i>Image Input (Base Model: Llama4 Maverick)</i>										
Baseline	7.0	7.0	7.3	9.4	4.5	4.0	6.5	6.7	8.8	3.0
Oracle	6.5	3.0	4.5	15.6	8.2	9.4	4.9	10.0	4.4	1.4
ColPali-v1.3	0.8	1.5	0.0	0.5	0.0	0.9	0.5	1.3	1.4	0.0
ColQwen2.5	1.2	2.1	0.7	0.5	0.0	1.2	0.2	2.7	2.0	0.0
VisRAG	1.0	2.0	0.0	1.0	0.0	1.0	1.6	1.3	1.2	0.0
VRAG-RL	10.9	9.8	11.6	17.8	8.2	11.0	6.8	13.1	10.8	8.1

CF, sparse formulaic content means gold images offer slight help. For LE, dense text distribution makes even direct access to target regions insufficient to reduce complexity for current models.

**In consistency-centric scan-oriented tasks, most retrieval-based enhancement methods show minimal effectiveness.** All embedding models exhibit poor retrieval accuracy. None achieves recall of 50% within the top-5 retrieved items. More critically, performance deteriorates after retrieval, especially for multimodal embedding models, where post-retrieval responses are almost entirely incorrect and scores approach 0.

**Complex embedding model architectures do not yield better performance.** Providing gold-standard images alleviates the scanning burden in long-context inputs, increasing the chances of retrieving correct answers. While overall performance improves, gains are limited for CF and minimal for LE errors. For CF, sparse formulaic content means gold images offer only slight localization help. For LE, dense error distribution makes even direct access to target regions insufficient to reduce task complexity for current models.

**Reinforcement learning frameworks with a visual-centric focus have distinguished themselves as leading approaches.** Despite being built on a compact 7B model, VRAG-RL consistently delivers improved performance and is the only method that achieves gains in the image-input setting following RL optimization. Its enhanced retrieval sharpens evidence selection, while strong reasoning provides effective guidance during document scanning. The retrieval and reasoning components are interleaved in design, with each stage informing the other in an iterative loop. This tightly coupled interaction contributes to the method’s superior performance potential.

## 5 CONCLUSION

In this paper, we introduce ScholScan, a benchmark designed to evaluate the performance of MLLMs on scan-oriented tasks that require detecting scientific errors across entire academic papers. We conduct a comprehensive evaluation and in-depth analysis of mainstream MLLMs and RAG methods. The results demonstrate that current MLLMs remain far from capable of reliably addressing such tasks, and that existing RAG approaches provide little to no improvement. This highlights the complexity, integrative demands, and originality of the ScholScan benchmark. Looking ahead, we aim to develop scan-oriented task paradigms suited to diverse academic scenarios and explore new techniques for enhancing model performance on target-suppressed inputs. These directions support the broader goal of advancing MLLMs from passive assistants to active participants in scientific research.

Table 3: Summary of retrieval performance for RAG methods.

Models	MRR@5	Recall@5
<i>Text Input (Base Model: Qwen3 Thinking)</i>		
bm25	0.41	0.48
BGE-M3	0.16	0.21
Contriever-msmacro	0.31	0.39
nv-embed-v2	0.30	0.38
<i>Image Input (Base Model: Llama4 Maverick)</i>		
ColPali-v1.3	0.26	0.31
ColQwen2.5	0.30	0.35
VisRAG	0.41	0.46

486 

## 6 ETHICS STATEMENT

487  
488 All data used in this paper were constructed by the authors and do not include any external public or  
489 proprietary datasets. The included academic papers and author names are publicly available through  
490 arXiv and OpenReview and can be freely accessed.491 A team of 10 domain experts was assembled to comprehensively review all task instances initially  
492 generated by Gemini 2.5 Pro. All annotators gave informed consent to participate. To ensure the ac-  
493 curacy and neutrality of both model-generated and human-verified content, we employed a rigorous  
494 multi-stage validation process involving cross-review and third-party adjudication.495 Evaluation across 15 mainstream models and 24 input configurations was conducted via legally  
496 authorized API access through the VolcEngine, Alibaba Cloud’s LLM services, and OpenRouter.497 ScholScan is fully open-sourced and freely available for academic and non-commercial research pur-  
498 poses. We provide the complete download link and documentation through an anonymous GitHub  
499 repository. All personally identifiable information has been removed from the dataset, and its col-  
500 lection and release comply with the ethical and legal requirements in place at the time of data acqui-  
501 sition.502 

## 7 REPRODUCIBILITY STATEMENT

503  
504 All results presented in this paper are fully reproducible. To facilitate verification and ex-  
505 tension, we provide an anonymous repository ([https://anonymous.4open.science/r/  
506 ScholScan-6657/](https://anonymous.4open.science/r/ScholScan-6657/)) that contains the complete dataset, source code, and detailed documentation.  
507 The repository also includes step-by-step instructions and the exact hyperparameter configurations  
508 used in our experiments, ensuring that other researchers can replicate our findings with minimal  
509 effort.510 The retrieval components in all retrieval-augmented generation (RAG) experiments were executed  
511 on a server equipped with 8 NVIDIA A40 GPUs.512 

## REFERENCES

513  
514 Anthropic. System card: Claude opus 4 & claude sonnet 4. <https://www.anthropic.com/claude-4-system-card>, May 2025. Updated Sep 2, 2025.515 S. Auer, Dante Augusto Couto Barone, Cassiano Bartz, E. Cortes, Mohamad Yaser Jaradeh,  
516 Oliver Karras, Manolis Koubarakis, Dmitry I. Mouromtsev, Dmitrii Pliukhin, Daniil Radyush,  
517 Ivan Shilin, Markus Stocker, and Eleni Tsalapati. The sciqa scientific question answering  
518 benchmark for scholarly knowledge. *Scientific Reports*, 13, 2023. URL <https://api.semanticscholar.org/CorpusID:258507546>.519  
520 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,  
521 Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan,  
522 Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng,  
523 Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-v1 technical report, 2025.  
524 URL <https://arxiv.org/abs/2502.13923>.525 ByteDance Seed Team. Introduction to techniques used  
526 in seed1.6. <https://seed/bytedance.com/en/blog/introduction-to-techniques-used-in-seed1-6>, June 2025. Official blog  
527 post describing Seed1.6 techniques.528 Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. Bge m3-embedding:  
529 Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge dis-  
530 tillation, 2024. URL <https://arxiv.org/abs/2402.03216>.531 Zhiyu Chen, Wenhui Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema  
532 Moussa, Matt Beane, Ting-Hao Huang, Bryan Routledge, and William Yang Wang. FinQA: A  
533 dataset of numerical reasoning over financial data. In Marie-Francine Moens, Xuanjing Huang,

540 Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical  
541 Methods in Natural Language Processing*, pp. 3697–3711, Online and Punta Cana, Dominican  
542 Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.  
543 emnlp-main.300. URL <https://aclanthology.org/2021.emnlp-main.300/>.  
544

545 Chao Deng, Jiale Yuan, Pi Bu, Peijie Wang, Zhong-Zhi Li, Jian Xu, Xiao-Hui Li, Yuan Gao, Jun  
546 Song, Bo Zheng, and Cheng-Lin Liu. LongDocURL: a comprehensive multimodal long docu-  
547 ment benchmark integrating understanding, reasoning, and locating. In Wanxiang Che, Joyce  
548 Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd An-  
549 nual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp.  
550 1135–1159, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-  
551 8-89176-251-0. doi: 10.18653/v1/2025.acl-long.57. URL <https://aclanthology.org/2025.acl-long.57/>.  
552

553 DeepSeek-AI et al. Deepseek-v3 technical report, 2025a. URL <https://arxiv.org/abs/2412.19437>.  
554

555 Gheorghe Comanici et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality,  
556 long context, and next generation agentic capabilities, 2025b. URL <https://arxiv.org/abs/2507.06261>.  
557

558 Long Phan et al. Humanity’s last exam, 2025c. URL <https://arxiv.org/abs/2501.14249>.  
559

560 OpenAI et al. gpt-oss-120b & gpt-oss-20b model card, 2025d. URL <https://arxiv.org/abs/2508.10925>.  
561

562 Manuel Faysse, Hugues Sibille, Tony Wu, Bilel Omrani, Gautier Viaud, CELINE HUDELOT,  
563 and Pierre Colombo. Colpali: Efficient document retrieval with vision language models. In  
564 *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=ogjBpZ8uSi>.  
565

566 Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Qianyu  
567 Guo, Meng Wang, and Haofen Wang. Retrieval-augmented generation for large language models:  
568 A survey. *ArXiv*, abs/2312.10997, 2023. URL <https://api.semanticscholar.org/CorpusID:266359151>.  
569

570 Yingqiang Ge, Wenyue Hua, Kai Mei, jianchao ji, Juntao Tan, Shuyuan Xu, Zelong Li, and  
571 Yongfeng Zhang. Openagi: When llm meets domain experts. In A. Oh, T. Naumann,  
572 A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Infor-  
573 mation Processing Systems*, volume 36, pp. 5539–5568. Curran Associates, Inc., 2023.  
574 URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/1190733f217404edc8a7f4e15a57f301-Paper-Datasets\\_and\\_Benchmarks.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/1190733f217404edc8a7f4e15a57f301-Paper-Datasets_and_Benchmarks.pdf).  
575

576 Daming Guo, Dongdong Yang, Hongyi Zhang, et al. Deepseek-r1 incentivizes reasoning in llms  
577 through reinforcement learning. *Nature*, 645:633–638, 2025. doi: 10.1038/s41586-025-09422-z.  
578

579 Yichen He, Guanhua Huang, Peiyuan Feng, Yuan Lin, Yuchen Zhang, Hang Li, and Weinan E. Pasa:  
580 An llm agent for comprehensive academic paper search, 2025. URL <https://arxiv.org/abs/2501.10120>.  
581

582 Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand  
583 Joulin, and Edouard Grave. Unsupervised dense information retrieval with contrastive learning,  
584 2022. URL <https://arxiv.org/abs/2112.09118>.  
585

586 Chankyu Lee, Rajarshi Roy, Mengyao Xu, Jonathan Raiman, Mohammad Shoeybi, Bryan Catan-  
587 zaro, and Wei Ping. NV-embed: Improved techniques for training LLMs as generalist embedding  
588 models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL  
589 <https://openreview.net/forum?id=lgsyLSSdRe>.  
590

594 Lei Li, Yuqi Wang, Runxin Xu, Peiyi Wang, Xiachong Feng, Lingpeng Kong, and Qi Liu. Multi-  
 595 modal ArXiv: A dataset for improving scientific comprehension of large vision-language models.  
 596 In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meet-  
 597 ing of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14369–14387,  
 598 Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/  
 599 2024.acl-long.775. URL <https://aclanthology.org/2024.acl-long.775/>.

600  
 601 Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee.  
 602 Llava-next: Improved reasoning, ocr, and world knowledge, January 2024. URL <https://llava-vl.github.io/blog/2024-01-30-llava-next/>.

603  
 604 Renze Lou, Hanzi Xu, Sijia Wang, Jiangshu Du, Ryo Kamoi, Xiaoxin Lu, Jian Xie, Yuxuan Sun,  
 605 Yusen Zhang, Jihyun Janice Ahn, Hongchao Fang, Zhuoyang Zou, Wenchao Ma, Xi Li, Kai  
 606 Zhang, Congying Xia, Lifu Huang, and Wenpeng Yin. AAAR-1.0: Assessing AI’s potential to  
 607 assist research. In *Forty-second International Conference on Machine Learning*, 2025. URL  
 608 <https://openreview.net/forum?id=RHAWcjIy12>.

609  
 610 Yubo Ma, Yuhang Zang, Liangyu Chen, Meiqi Chen, Yizhu Jiao, Xinze Li, Xinyuan Lu, Ziyu  
 611 Liu, Yan Ma, Xiaoyi Dong, Pan Zhang, Liangming Pan, Yu-Gang Jiang, Jiaqi Wang, Yixin Cao,  
 612 and Aixin Sun. MMLONGBENCH-DOC: Benchmarking long-context document understanding  
 613 with visualizations. In *The Thirty-eight Conference on Neural Information Processing Systems  
 614 Datasets and Benchmarks Track*, 2024. URL <https://openreview.net/forum?id=1oJM1acwzf>.

615  
 616 Meta. Llama 4 | model cards and prompt formats. [https://www.llama.com/docs/  
 617 model-cards-and-prompt-formats/llama4/](https://www.llama.com/docs/model-cards-and-prompt-formats/llama4/), 2025. Official model card and prompt  
 618 format documentation.

619  
 620 Meredith Ringel Morris, Jascha Sohl-Dickstein, Noah Fiedel, Tris Warkentin, Allan Dafoe, Alek-  
 621 sandra Faust, Clement Farabet, and Shane Legg. Position: Levels of AGI for operationalizing  
 622 progress on the path to AGI. In *Forty-first International Conference on Machine Learning*, 2024.  
 623 URL <https://openreview.net/forum?id=0ofzEysK2D>.

624  
 625 OpenAI. Gpt-5 system card. Technical report, OpenAI, August 2025. URL <https://cdn.openai.com/gpt-5-system-card.pdf>. Version: August 13, 2025.

626  
 627 OpenAI. gpt-4.1 — openai api documentation, 2025. URL <https://platform.openai.com/docs/models/gpt-4.1>. Accessed: 2025-09-25.

628  
 629 Shraman Pramanick, Rama Chellappa, and Subhashini Venugopalan. SPIQA: A dataset for mul-  
 630 timodal question answering on scientific papers. In *The Thirty-eight Conference on Neu-  
 631 ral Information Processing Systems Datasets and Benchmarks Track*, 2024. URL <https://openreview.net/forum?id=h3lddsY5nf>.

632  
 633 Stephen E. Robertson, Steve Walker, Susan Jones, Micheline Hancock-Beaulieu, and Mike  
 634 Gatford. Okapi at trec-3. In *Text Retrieval Conference*, 1994. URL <https://api.semanticscholar.org/CorpusID:41563977>.

635  
 636 R. Smith. An overview of the tesseract ocr engine. In *Ninth International Conference on Document  
 637 Analysis and Recognition (ICDAR 2007)*, volume 2, pp. 629–633, 2007. doi: 10.1109/ICDAR.  
 638 2007.4376991.

639  
 640 Rubèn Tito, Minesh Mathew, C. V. Jawahar, Ernest Valveny, and Dimosthenis Karatzas. ICDAR  
 641 2021 competition on document visualquestion answering. *CoRR*, abs/2111.05547, 2021. URL  
 642 <https://arxiv.org/abs/2111.05547>.

643  
 644 Quchen Wang, Ruixue Ding, Yu Zeng, Zehui Chen, Lin Chen, Shihang Wang, Pengjun Xie,  
 645 Fei Huang, and Feng Zhao. Vrag-rl: Empower vision-perception-based rag for visually rich  
 646 information understanding via iterative reasoning with reinforcement learning, 2025. URL  
 647 <https://arxiv.org/abs/2505.22019>.

648 Zirui Wang, Mengzhou Xia, Luxi He, Howard Chen, Yitao Liu, Richard Zhu, Kaiqu Liang,  
 649 Xindi Wu, Haotian Liu, Sadhika Malladi, Alexis Chevalier, Sanjeev Arora, and Danqi Chen.  
 650 Charxiv: Charting gaps in realistic chart understanding in multimodal llms, 2024. URL <https://arxiv.org/abs/2406.18521>.  
 651

652 xAI. Grok 4 fast model card. Technical report, xAI, September 2025. URL <https://data.xai/2025-09-19-grok-4-fast-model-card.pdf>. Last updated: September 19, 2025.  
 653

654

655 Dawei Yan, Yang Li, Qing-Guo Chen, Weihua Luo, Peng Wang, Haokui Zhang, and Chunhua Shen.  
 656 Mmcrl: Advancing visual language model in multimodal multi-turn contextual reasoning, 2025.  
 657 URL <https://arxiv.org/abs/2503.18533>.  
 658

659 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang  
 660 Gao, Chengan Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu,  
 661 Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin  
 662 Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang,  
 663 Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui  
 664 Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang  
 665 Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger  
 666 Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan  
 667 Qiu. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.  
 668

669 Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov,  
 670 and Christopher D. Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question  
 671 answering. *CoRR*, abs/1809.09600, 2018. URL <http://arxiv.org/abs/1809.09600>.  
 672

673 Shi Yu, Chaoyue Tang, Bokai Xu, Junbo Cui, Junhao Ran, Yukun Yan, Zhenghao Liu, Shuo Wang,  
 674 Xu Han, Zhiyuan Liu, and Maosong Sun. VisRAG: Vision-based retrieval-augmented generation  
 675 on multi-modality documents. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=zG459X3Xge>.  
 676

677 Xiang Yue, Yuansheng Ni, Tianyu Zheng, Kai Zhang, Ruoqi Liu, Ge Zhang, Samuel Stevens,  
 678 Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun,  
 679 Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and  
 680 Wenhui Chen. Mmmu: A massive multi-discipline multimodal understanding and reasoning  
 681 benchmark for expert agi. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9556–9567, 2024. doi: 10.1109/CVPR52733.2024.00913.  
 682

683 Yilun Zhao, Yitao Long, Hongjun Liu, Ryo Kamoi, Linyong Nan, Lyuhao Chen, Yixin Liu, Xiangru  
 684 Tang, Rui Zhang, and Arman Cohan. DocMath-eval: Evaluating math reasoning capabilities of  
 685 LLMs in understanding long and specialized documents. In Lun-Wei Ku, Andre Martins, and  
 686 Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Com-  
 687 putational Linguistics (Volume 1: Long Papers)*, pp. 16103–16120, Bangkok, Thailand, August  
 688 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.852. URL  
 689 <https://aclanthology.org/2024.acl-long.852/>.  
 690

691 Ge Zheng, Bin Yang, Jiajin Tang, Hong-Yu Zhou, and Sibei Yang. Ddcot: Duty-distinct  
 692 chain-of-thought prompting for multimodal reasoning in language models. In A. Oh,  
 693 T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neu-  
 694 ral Information Processing Systems*, volume 36, pp. 5168–5191. Curran Associates, Inc.,  
 695 2023. URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/108030643e640ac050e0ed5e6aace48f-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/108030643e640ac050e0ed5e6aace48f-Paper-Conference.pdf).  
 696

697 Ruiyang Zhou, Lu Chen, and Kai Yu. Is LLM a reliable reviewer? a comprehensive evaluation  
 698 of LLM on automatic paper reviewing tasks. In Nicoletta Calzolari, Min-Yen Kan, Veronique  
 699 Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (eds.), *Proceedings of the 2024  
 700 Joint International Conference on Computational Linguistics, Language Resources and Evalu-  
 701 ation (LREC-COLING 2024)*, pp. 9340–9351, Torino, Italia, May 2024. ELRA and ICCL. URL  
<https://aclanthology.org/2024.lrec-main.816/>.

# 702 703 704 705 706 707 Not Search, But Scan: Benchmarking MLLMs 708 on Scan-Oriented Academic Paper Reasoning 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 Supplementary Material

### 708 Table of Contents in Appendix 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 <b>A Use of LLMs</b>	<b>16</b>
<b>B Prompts</b>	<b>17</b>
B.1 Within-Generate Prompt . . . . .	17
B.2 Within-Sample Prompt . . . . .	20
B.3 Extractor Prompt . . . . .	22
B.4 System Prompt . . . . .	24
<b>C Examples from Existing Datasets</b>	<b>25</b>
C.1 Example from DocMath-Eval . . . . .	25
C.2 Example from SlideVQA . . . . .	26
C.3 Example from MMLongBench-Doc . . . . .	27
C.4 Example from LongDocURL . . . . .	28
C.5 Example from ArXivQA . . . . .	29
C.6 Example from Charxiv . . . . .	30
C.7 Example from AAAR . . . . .	31
C.8 Example from MMCR . . . . .	32
C.9 Example from DocVQA . . . . .	33
C.10 Example from SPIQA . . . . .	34
<b>D Dataset Annotation and Construction</b>	<b>35</b>
D.1 Human Annotator Guidelines . . . . .	35
D.2 Annotation Statistics . . . . .	35
D.3 Examples of Annotation . . . . .	36
D.3.1 Case 1: Discard Directly . . . . .	36
D.3.2 Case 2: Modify Question . . . . .	37
D.3.3 Case 3: Modify Explanation . . . . .	38
D.3.4 Case 4: Modify Category . . . . .	39
<b>E Common Failure Cases of MLLMs</b>	<b>40</b>
E.1 RQD (Research Question & Definitions) . . . . .	40
E.2 DI (Design & Identifiability) . . . . .	42
E.3 SG (Sampling & Generalizability) . . . . .	44
E.4 RCA (Referential and Citation Alignment) . . . . .	46

756	E.5 MO (Measurement & Operationalization) . . . . .	48
757	E.6 DHP (Data Handling & Preprocessing) . . . . .	50
758	E.7 CF (Computation & Formulae) . . . . .	52
759	E.8 IC (Inference & Conclusions) . . . . .	54
760	E.9 LE (Language & Expression) . . . . .	56
761		
762		
763		
764	<b>F Human-Machine Consistency Evaluation</b>	<b>58</b>
765		
766		
767		
768		
769		
770		
771		
772		
773		
774		
775		
776		
777		
778		
779		
780		
781		
782		
783		
784		
785		
786		
787		
788		
789		
790		
791		
792		
793		
794		
795		
796		
797		
798		
799		
800		
801		
802		
803		
804		
805		
806		
807		
808		
809		

810 **A USE OF LLMs**  
811812 Large language models (LLMs) were used solely to assist in language editing and stylistic refinement  
813 during manuscript preparation. All technical content, experiments, dataset construction, evaluation  
814 protocols, and analysis were conceived, implemented, and validated entirely by the authors. No  
815 LLMs were involved in the generation of benchmark data, research methodology design, or result  
816 interpretation. The use of LLMs did not influence the scientific conclusions of this paper.  
817  
818  
819  
820  
821  
822  
823  
824  
825  
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  
865**B PROMPTS**866  
867**B.1 WITHIN-GENERATE PROMPT**868  
869**Within-Generate Prompt**870  
871

You will receive a high-quality, already accepted scientific paper as a PDF. Working only with the PDF itself (and any appendix embedded in the same PDF), edit specific textual spans to inject one or more errors chosen only from the taxonomy below, such that the errors are hard yet clearly identifiable by a professional reviewer reading the PDF alone.

872  
873

Error Type (fixed):

874  
875

Research Question & Definitions

876  
877

Definition: The core construct/hypothesis/variable is insufficiently or inconsistently defined (conceptual vs operational), leaving the estimand ambiguous.

Design & Identifiability

878  
879

Definition: Given a clear estimand, the design violates structural identification conditions so the effect is not identifiable even with infinite data and perfect measurement.

880  
881

Sampling & Generalizability

882  
883

Definition: The sampling frame/process/composition or cluster/power setup does not support valid or stable sample→population claims.

884  
885

Measurement & Operationalization

886  
887

Definition: Measures/manipulations lack feasibility/reliability/validity/timing, so observed variables systematically diverge from the intended construct/treatment.

888  
889

Data Handling & Preprocessing

890  
891

Definition: Pipeline choices in missing handling, joins/keys, temporal splitting, feature construction, or partitioning introduce bias (incl. leakage or unit/scale conflicts).

892  
893

Computation & Formulae

894  
895

Definition: Arithmetic/algebra/notation errors (totals/ratios, unit conversion, CI vs point estimate, p-value vs label, symbol reuse, undefined variables, dimension mismatch).

896  
897

Inference & Conclusions

898  
899

Definition: Interpretations or causal statements exceed what methods/data support, or contradict the shown statistics/tables/captions.

900  
901

Referential and Citation Alignment

902  
903

Definition: Contradictions about the same quantity/term across text, tables, captions, or appendix within the paper.

904  
905

Language & Expression

906  
907

Definition: Terminology/capitalization/grammar ambiguities that affect meaning or domain-critical term consistency (not cosmetic typos).

908  
909910  
911912  
913914  
915916  
917

918

## Within-Generate Prompt (Continued)

919

920

921 Global constraints (must comply)

922 1. Each error must map to exactly one primary category in the  
923 taxonomy. Do not mix causes.924 2. Each error must involve more than 2 micro-edits (each edit  
925 ≤ 20 English words) spread across distinct pages or  
926 paragraphs.927 3. If an edit would create an immediate contradiction in the  
928 same sentence/paragraph/caption, you may add shadow patch  
929 (es) for the same error to keep the text natural (still  
930 counted as edit locations).

931 4. Independence across errors (per-copy generation)

932 Generate each error on a separate copy of the original PDF  
933 . Different errors must be logically and operationally  
934 independent:935 No progression or variant relations: an error must not be  
936 a stricter/looser version, superset/subset, or minor  
937 wording variant of another error.938 No anchor reuse: do not target the same sentence/caption/  
939 table cell or reuse the same old\_str (or a near-  
940 duplicate paraphrase) across different errors.941 Applying any single error in isolation to the original PDF  
942 must still yield a detectable, clearly categorizable  
943 error according to the taxonomy.

944 5. Every error must be supportable using text inside the PDF.

945 Do not rely on external supplementary files or prior  
946 knowledge.947 6. Design as difficult as possible but clean errors. Prefer  
948 edits that force cross-checking between two spots (e.g.,  
949 Methods vs Results). Avoid trivialities. Edits must  
950 remain locally plausible and not advertise themselves via  
951 obviously artificial phrases (e.g., avoid contrived  
952 tokens purely added to be detectable).953 7. ‘‘No cosmetic issues’’ applies except for I (Language &  
954 Expression). For I, edits must affect meaning or domain-  
955 critical terminology (e.g., ambiguous phrasing,  
956 inconsistent technical terms). Pure typos, punctuation  
957 tweaks, or layout nits are not allowed.958 8. Do not edit titles, author lists, bibliography entries,  
959 equation numbering, figure images, or add new figures/  
960 tables/references.961 9. Frame each question as a neutral imperative that asks for  
962 a decision about a specific condition, using (but not  
963 limited to) Decide/Determine/Judge/Evaluate/Assess  
964 whether.... Do not presuppose an outcome or use  
965 suggestive intensifiers (e.g., clearly/obviously/likely/  
966 suspicious as examples).

967

968

969

970

971

972  
973  
974  
975  
976  
977

## Within-Generate Prompt (Continued)

```

10. Output English-only and strictly follow the JSON schema
    below. Do not include any additional text outside the
    JSON:
    [
        {
            "id": "1-based integer as string",
            "modify": [
                {
                    "location": "Page number + short unique nearby quote ("
                        "≤15 tokens).",
                    "old_str": "Exact original text from the PDF (verbatim)"
                        ".",
                    "new_str": "Edited text after your change."
                }
            /* Add 1-2 more locations; each location ≤ 20 words
            changed.
            Shadow patches for local coherence count as locations.
            */
        ],
        "question": "One neutral audit-style task (1-25 words).",
        "explanation": "Explain in 2-4 sentences why a reviewer can
            detect this error from the edited PDF alone.",
        "Type": "Name the primary category (e.g., Inference &
            Conclusions).",
    }
    /* More Errors */
]

```

1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025

1026  
1027

## B.2 WITHIN-SAMPLE PROMPT

1028  
1029

## Within-Sample Prompt

1030  
1031

You will receive a paper PDF and the weaknesses mentioned in its peer-review comments. Your task is, based only on the content of that PDF, to sample from the review comments and verify possible errors related to the categories below, and for each confirmed or highly plausible error, generate one question and one explanation.

1032  
1033  
1034  
1035  
1036

Error Type (fixed):

1037  
1038

Research Question & Definitions

1039  
1040  
1041

Definition: The core construct/hypothesis/variable is insufficiently or inconsistently defined (conceptual vs operational), leaving the estimand ambiguous.

1042  
1043  
1044  
1045

Design & Identifiability

1046  
1047  
1048  
1049

Definition: Given a clear estimand, the design violates structural identification conditions so the effect is not identifiable even with infinite data and perfect measurement.

1050  
1051  
1052  
1053  
1054

Sampling & Generalizability

1055  
1056  
1057  
1058

Definition: The sampling frame/process/composition or cluster/power setup does not support valid or stable sample→population claims.

1059

Measurement & Operationalization

1060  
1061  
1062  
1063

Definition: Measures/manipulations lack feasibility/reliability/validity/timing, so observed variables systematically diverge from the intended construct/treatment.

1064  
1065  
1066  
1067

Data Handling & Preprocessing

1068  
1069  
1070

Definition: Pipeline choices in missing handling, joins/keys, temporal splitting, feature construction, or partitioning introduce bias (incl. leakage or unit/scale conflicts).

1071  
1072  
1073  
1074

Computation & Formulae

1075  
1076  
1077  
1078  
1079

Definition: Arithmetic/algebra/notation errors (totals/ratios, unit conversion, CI vs point estimate, p-value vs label, symbol reuse, undefined variables, dimension mismatch).

Inference & Conclusions

1075  
1076  
1077  
1078  
1079

Definition: Interpretations or causal statements exceed what methods/data support, or contradict the shown statistics/tables/captions.

Referential and Citation Alignment;

1075  
1076  
1077  
1078  
1079

Definition: Contradictions about the same quantity/term across text, tables, captions, or appendix within the paper.

Language & Expression

1075  
1076  
1077  
1078  
1079

Definition: Terminology/capitalization/grammar ambiguities that affect meaning or domain-critical term consistency (not cosmetic typos).

1080

1081

1082

1083

Within-Sample Prompt (Continued)

1083 Global constraints (must comply)  
 1084 Output only the specified categories; even if other error  
 1085 types appear in the reviews, do not output them.  
 1086 Sample first, then verify: extract candidates from the review  
 1087 comments, then confirm them in the PDF. If you cannot  
 1088 locate supporting anchors in the PDF (page number plus  
 1089 phrase/label), do not output that candidate.  
 1090 Questions must be neutral and non-leading: use an "audit task  
 1091 + decision" style, avoiding yes/no bias.  
 1092 Independence: each question must target a different figure or  
 1093 different textual anchor; no minor variants of the same  
 1094 issue.  
 1095 Evidence first: the explanation must cite locatable anchors  
 1096 in the PDF (page number + original phrase/caption). You  
 1097 may mention a key short phrase from the review as a clue,  
 1098 but write the question and explanation in your own words.  
 1099 Language & format: both question and explanation must be in  
 1100 English; output JSON only, with no extra text.  
 1101 Quantity: sort by evidence strength and output up to 5 items;  
 1102 if none qualify, output an empty array [].  
 1103 Example output  
 1104 [  
 1105 {  
 1106 "id": "1",  
 1107 "question": "Audit y-axis baselines and possible axis  
 1108 breaks in Figure 2; decide presence/absence and cite  
 1109 evidence.",  
 1110 "explanation": "The review flags possible exaggeration in  
 1111 Fig.2. In the PDF (p.6, caption 'Performance vs  
 1112 baseline'), the y-axis starts at 0.85 with a break,  
 1113 magnifying small differences; panels use different  
 1114 ranges."  
 1115 "Type": "Visualization & Presentation Bias"  
 1116 }  
 1117 ]

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

1134  
1135 B.3 EXTRACTOR PROMPT

1136 Extractor Prompt

1138 You will receive three inputs:  
1139 Q: the open-ended question;  
1140 E: the gold explanation (describes exactly one error; extra  
1141 details still belong to the same single error);  
1142 A: the model's answer to be evaluated.  
1143 Your job is to extract counts only and output a single JSON  
1144 object with the exact schema below. Do not compute any  
1145 scores. Do not add fields.  
1146  
1147 Core selection rule (multiple errors in A)  
1148 1. Parse E into a single gold error (the "target error").  
1149 2. From A, identify how many distinct error claims are made.  
1150 Cluster together mentions that support the same error (multiple  
1151 locations for one error are still one error).  
1152 3. Existence decision (binary correctness only):  
1153 Let the gold existence be 1 if E asserts an error exists,  
1154 else 0.  
1155 Let the predicted existence be 1 if A asserts any error, else  
1156 0 (e.g., states no error).  
1157 Set existance = 1 if predicted existence equals gold  
1158 existence; otherwise set existance = 0.  
1159 4. If existance = 0: set contains\_target\_error = 0; set all  
1160 location and reasoning counts to 0; and set  
1161 unrelated\_errors to the total number of distinct error  
1162 claims in A. Then output the JSON.  
1163 5. If existance = 1:  
1164 If the gold existence is 1: determine whether A contains the  
1165 target error (match by the main error idea in E: category  
1166 /intent/scope; treat E's subpoints as the same error).  
1167 If yes, set contains\_target\_error = 1 and compute location  
1168 and reasoning only for the target error. Count all  
1169 other error claims in A as unrelated\_errors.  
1170 If no, set contains\_target\_error = 0; set all location and  
1171 reasoning counts to 0; set unrelated\_errors to the  
1172 total number of distinct error claims in A.  
1173 If the gold existence is 0: set contains\_target\_error = 0;  
1174 set all location and reasoning counts to 0; set  
1175 unrelated\_errors to the total number of distinct error  
1176 claims in A. (These negative items are for binary  
1177 accuracy only; they are not used for detailed scoring.)  
1178  
1179 Matching guidance (A error  $\leftrightarrow$  target error): match by the  
1180 main error idea in E (category/intent/scope), not by  
1181 wording. Treat E's subpoints as part of the same single  
1182 error. Prefer the best-matching cluster in A; if ties,  
1183 choose the one with stronger alignment to E's core claim.  
1184  
1185  
1186  
1187

1188

1189

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

## Extractor Prompt (Continued)

```

Counting rules
Location (for the target error only when existance=1 and
contains_target_error=1):
gold_steps: number of unique error locations described in E (
    after normalization and deduplication).
hit_steps: number of predicted locations in A that match any
gold location for the target error.
extra_steps: number of predicted locations in A for the
target error that do not match any gold location.

Reasoning (for the target error only when existance=1 and
contains_target_error=1):
Convert E into a canonical set or ordered chain of reasoning
steps for the target error.
gold_steps: total number of such steps.
reached_steps:
    single-chain tasks: length of the longest valid prefix of
        A along the gold chain;
    multi-path/parallel tasks: size of the intersection
        between A's steps and the gold step set (or the
        maximum across gold paths if multiple are defined).
missing_steps: gold_steps - reached_steps (non-negative
integer).
Unrelated errors:
unrelated_errors: number of distinct error claims in A that
are not the target error (0 if none).
Output schema (return exactly this JSON; integers only)
{
    "existance": 0,
    "contains_target_error": 0,
    "location": {
        "gold_steps": 0,
        "hit_steps": 0,
        "extra_steps": 0
    },
    "reasoning": {
        "gold_steps": 0,
        "reached_steps": 0,
        "missing_steps": 0
    },
    "unrelated_errors": 0
}

```

1242  
1243

## B.4 SYSTEM PROMPT

1244  
1245

1246

1247 You are a neutral, careful academic reviewer. You will  
 1248 receive an open-ended question and the paper content. The  
 1249 paper may or may not have issues related to the question  
 1250 Do not assume there are errors. If the question is about  
 1251 citations, you will be given a citing paper and a cited  
 1252 paper; evaluate only the citing paper for possible issues  
 1253 and use the cited paper only as the reference for  
 1254 comparison. Write in natural prose with no fixed template

1255

## Rules:

1256  
1257

1256 Speak only when sure. State an error only if you are  
 1257 confident it is a real error (not a mere weakness).

1258  
1259

1258 Stay on scope. Discuss only what the question asks about.

1259  
1260

1259 Evidence completeness. For every error you state, list all  
 1260 distinct evidence cues you are confident about from the  
 1261 PDF. Include plain identifiers (figure/table/section/  
 1262 equation/citation) or quotes. Avoid redundant repeats of  
 1263 the exact same instance; include all distinct locations  
 1264 needed to support the error.

1264  
1265

1264 Be clear and brief. Use short, direct sentences.

1265 No metaphors. No fancy wording. No guesses or outside sources.

1266  
1267

1266 Do not invent figures, tables, equations, citations, or  
 1267 results.

1268  
1269  
1270

1268 Report as many distinct, well-supported errors as you can  
 1269 within scope. If none are clear, write exactly: "No clear  
 1270 issue relevant to the question." and nothing else.

1271  
1272

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

## System Prompt

You are a neutral, careful academic reviewer. You will receive an open-ended question and the paper content. The paper may or may not have issues related to the question. Do not assume there are errors. If the question is about citations, you will be given a citing paper and a cited paper; evaluate only the citing paper for possible issues and use the cited paper only as the reference for comparison. Write in natural prose with no fixed template

## Rules:

Speak only when sure. State an error only if you are confident it is a real error (not a mere weakness).

Stay on scope. Discuss only what the question asks about.

Evidence completeness. For every error you state, list all distinct evidence cues you are confident about from the PDF. Include plain identifiers (figure/table/section/equation/citation) or quotes. Avoid redundant repeats of the exact same instance; include all distinct locations needed to support the error.

Be clear and brief. Use short, direct sentences.

No metaphors. No fancy wording. No guesses or outside sources.

Do not invent figures, tables, equations, citations, or results.

Report as many distinct, well-supported errors as you can within scope. If none are clear, write exactly: "No clear issue relevant to the question." and nothing else.

1296 **C EXAMPLES FROM EXISTING DATASETS**  
12971298 **C.1 EXAMPLE FROM DOCMATH-EVAL**  
12991300 **One Example from DocMath-Eval**  
13011302 **Question ID:** complong-testmini-301303 **Question:** What is the percentage of total offering cost on the total amount raised in the IPO **if the**  
1304 **total offering cost is \$14,528,328 and each unit sold is \$10?**

1305

1306 **Context Modalities: Texts Documents**1307 1. Offering costs consist of legal, accounting and other costs incurred through the balance sheet date that  
1308 are directly related to the Initial Public Offering. Offering costs amounting to \$14,528,328 were charged to  
1309 shareholders' equity upon the completion of the Initial Public Offering.1310 2. Pursuant to the Initial Public Offering on July 20, 2020, the Company sold 25,300,000 Units, which  
1311 includes the full exercise by the underwriter of its option to purchase an additional 3,300,000 Units, at a  
1312 purchase price of \$10.00 per Unit. Each Unit consists of one Class A ordinary share and one-half of one  
1313 redeemable warrant ("Public Warrant"). Each whole Public Warrant entitles the holder to purchase one  
1314 Class A ordinary share at an exercise price of \$11.50 per whole share (see Note 7).

1314

**NO Multi-Modal Documents Context**

1315

1316 **Covered areas:****Focus Only On the Field of Mathematics**

1317

1318

1319

1320

**Cross-evidence Reasoning:**Focusing on solving mathematical problems requires integrating evidence such as mathematical formulas,  
1321 question stem conditions, and chart data from different positions in the document.

1322

1323

1324

1325

**Task Paradigm: search****Search-oriented**

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

1350 C.2 EXAMPLE FROM SLIDEVQA  
13511352  
13531354 **Question ID: 1**1355 **Question:** How much difference in INR is there between the average order value of CY2013 and that  
1356 of CY2012?

1357

1358

1359 **Context Modalities: Multi-Modal Documents and Texts**

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

**Covered areas:**

The documents cover core technical research fields such as visual question answering and machine reading comprehension, as well as industry application fields including education and scientific research, finance and commerce, and healthcare (with derivative adaptation to pathological slice analysis), and also involves derivative technical fields like retrieval-augmented generation.

1380

1381

1382

1383

1384

1385

**Cross-evidence Reasoning:**

Simple question types only require a single piece of evidence

**Not Cross-evidence Reasoning**

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

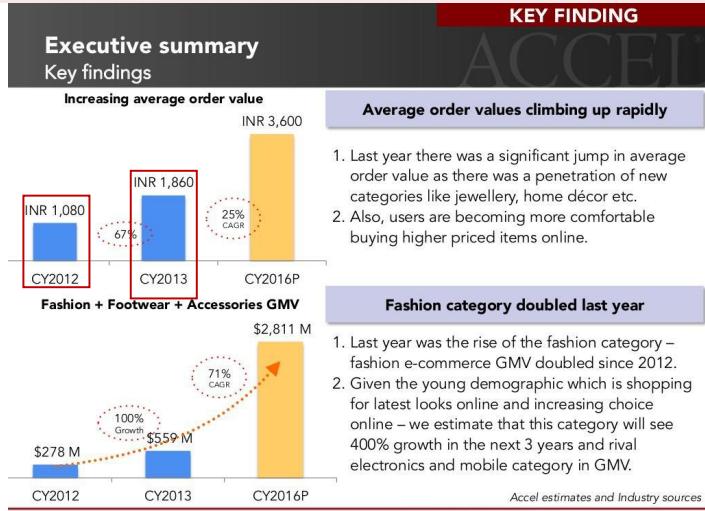
1399

1400

1401

1402

1403

1352 **One Example from SlideVQA**1353 **Question ID: 1**1354 **Question:** How much difference in INR is there between the average order value of CY2013 and that  
1355 of CY2012?1356 **Context Modalities: Multi-Modal Documents and Texts****Covered areas:**

The documents cover core technical research fields such as visual question answering and machine reading comprehension, as well as industry application fields including education and scientific research, finance and commerce, and healthcare (with derivative adaptation to pathological slice analysis), and also involves derivative technical fields like retrieval-augmented generation.

**Cross-evidence Reasoning:**

Simple question types only require a single piece of evidence

**Not Cross-evidence Reasoning**

**Task Paradigm: search**

**Search-oriented**

1404  
1405

## C.3 EXAMPLE FROM MMLONGBENCH-DOC

1406  
1407

## One Example from MMLongBench-Doc

1408  
1409**Question ID:**

1410

**Question:** How much higher was the proposed dividend paid (Rupees in lacs) in 2002 compared to 2001?1411  
1412**Context Modalities: Multi-Modal Documents and Texts**1413  
1414


SHAREHOLDER REFERENCER					
Unclaimed Dividend					
Unclaimed Dividend for the years prior to and including the financial year 1998-99 has been transferred to the General Revenue Account of the Central Government / the Investor Education and Protection Fund established by the Central Government (IEPF), as applicable.					
Shareholders who have not encashed their dividend warrants relating to financial years up to and including 1993-94 may claim such dividend (transferred to the General Revenue Account) from the Registrar of Companies, West Bengal, Government of India, Nizam Palace, II MSO Building, 2nd Floor, 2344 A.J.C. Bose Road, Kolkata 700 020, in the prescribed form. This form can be furnished by the Investor Service Centre of the Company (ISC) on request or can be downloaded from the Company's website at <a href="http://www.itchotels.com">www.itchotels.com</a> .					
The dividend for the unclaimed years, if unclaimed for 7 years, will be transferred by the Company to IEPF in accordance with the schedule given below. Attention is drawn that any unclaimed dividend for the financial year 1998-2000 will be due for transfer to IEPF later this year. Communication has been sent by the Company to the concerned Shareholders advising them to lodge their claims with respect to unclaimed dividend.					
Once unclaimed dividend is transferred to IEPF, no claim shall lie in respect thereof.					
ITC Limited					
Financial Year	Dividend Identification No.	Date of Declaration of Dividend	Total Dividend (Rs.)	Unclaimed Dividend as on 31/03/2007 (Rs.)	
				%	
1996-97	70th	28th July, 2000	1,84,46,17,780.00	1,06,82,671.00	0.69
2000-01	71st	28th July, 2001	2,45,47,03,040.00	2,46,42,13,000.00	1.00
2001-02	72nd	28th July, 2002	2,34,52,27,743.00	2,96,63,740.00	0.77
2002-03	73rd	28th July, 2003	3,71,26,76,200.00	2,38,48,71,800.00	0.64
2003-04	74th	30th July, 2004	4,95,36,77,020.00	3,35,88,620.00	0.68
2004-05	75th	29th July, 2005	7,73,24,96,398.00	5,07,52,301.00	0.66
2005-06	76th	21st July, 2006	9,65,28,12,267.00	7,36,87,328.00	0.74

\* It will not be possible to entertain claims received by ISC after 14th September, 2007.

Erstwhile ITC Hotels Limited				
Financial Year	Date of Declaration of Dividend	Total Dividend (Rs.)	Unclaimed Dividend as on 31/03/2007 (Rs.)	Due for transfer to IEPF on
1996-97	29th August, 2000	3,02,96,402.00	3,02,96,402.00	1.00
2000-01	17th August, 2001	3,02,96,402.00	3,04,56,200.00	1.01
2003-04	14th July, 2004	6,04,32,694.00	6,09,704.00	1.16

\* It will not be possible to entertain claims received by ISC after 9th October, 2007.

**Bank Details**  
Shareholders holding Shares in the physical form are requested to notify / send the following to ISC to facilitate better servicing:-  
i) any change in their address / mandate / bank details, and  
ii) particulars of the bank account in which they wish their dividend to be credited, in case the same have not been furnished earlier.  
Shareholders are advised that respective bank details and addresses as furnished by them or by NSDL / CDSL to the Company, for Share held in the physical form and in the dematerialized form respectively, will be printed on dividend warrants as a measure of protection against fraudulent encashment.

30

1425  
1426**Covered areas:**1427  
1428

The documents cover 7 diverse fields such as scientific research reports, business financial reports, and technical manuals.

1429  
1430**Cross-evidence Reasoning:**1431  
1432

33% of the questions are cross-page questions, which require integrating different types of evidence such as texts, tables, and charts from multi-page documents

1433  
1434**Task Paradigm: search**

1435

**Search-oriented**

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

1458  
1459 C.4 EXAMPLE FROM LONGDOCURL

## 1460 One Example from LongDocURL

1462 **Question ID:** free\_gemini15\_pro\_4061601\_47\_71\_81463 **Question:** What was the total fair value of options that vested in 2016, 2015, and 2014, in millions of  
1464 Canadian dollars?1465  
1466 Context Modalities: Multi-Modal Documents and Texts

1467 The following table summarizes additional stock option information:

year ended December 31	2016	2015	2014
(millions of Canadian \$, unless otherwise noted)			
Total intrinsic value of options exercised	31	10	21
Fair value of options that have vested	126	91	95

1469 Total options vested **2.1 million** 2.0 million 1.7 million

1470 As at December 31, 2016, the aggregate intrinsic value of the total options exercisable was \$96 million and the total intrinsic value

1471 of options outstanding was \$130 million.

## 21 PREFERRED SHARES

1472 In March 2014, TCPL redeemed all of the 4 million outstanding Series Y preferred shares at a redemption price of \$50 per share

1473 for a gross payment of \$200 million.

## 22 OTHER COMPREHENSIVE (LOSS)/INCOME AND ACCUMULATED OTHER COMPREHENSIVE LOSS

1474 Components of Other comprehensive (loss)/income, including the portion attributable to non-controlling interests and related tax

1475 effects, are as follows:

year ended December 31, 2016	Before Tax Amount	Income Tax Recovery/ (Expense)	Net of Tax Amount
(millions of Canadian \$)			
Foreign currency translation gains on net investment in foreign operations	3	3	3
Change in fair value of net investment hedges	(14)	4	(10)
Change in fair value of cash flow hedges	44	(14)	30
Reclassification to net income of gains and losses on cash flow hedges	71	(29)	42
Unrealized actuarial gains and losses on pension and other post-retirement benefit			
plans	(36)	12	(24)
Reclassification to net income of actuarial loss on pension and other post-			
retirement benefit plans	22	(6)	16
Other comprehensive loss on equity investments	(117)	30	(87)
<b>Other Comprehensive Loss</b>	<b>(26)</b>	<b>(3)</b>	<b>(23)</b>

year ended December 31, 2015	Before Tax Amount	Income Tax Recovery/ (Expense)	Net of Tax Amount
(millions of Canadian \$)			
Foreign currency translation gains on net investment in foreign operations	798	15	813
Change in fair value of net investment hedges	(505)	133	(372)
Change in fair value of cash flow hedges	(92)	35	(57)
Reclassification to net income of gains and losses on cash flow hedges	144	(56)	88
Unrealized actuarial gains and losses on pension and other post-retirement benefit			
plans	74	(23)	51
Reclassification to net income of actuarial loss and prior service costs on pension			
and other post-retirement benefit plans	41	(9)	32
Other comprehensive income on equity investments	62	(15)	47
<b>Other Comprehensive Income</b>	<b>522</b>	<b>80</b>	<b>602</b>

1480 155 TCPL [Consolidated financial statements](#) 2016

1481  
1482  
1483  
1484  
1485  
1486  
1487  
1488  
1489  
1490  
1491  
1492  
1493  
1494  
1495  
1496  
1497  
1498  
1499  
1500  
1501  
1502  
1503  
1504  
1505  
1506  
1507  
1508  
1509  
1510  
1511  
Covered areas:

The document types of LongDocURL cover 8 major categories such as research reports, user manuals, and books.

1500  
1501  
1502  
1503  
1504  
1505  
1506  
1507  
1508  
1509  
1510  
1511  
Cross-evidence Reasoning:

Most questions require integrating evidence across chapters and elements

1500  
1501  
1502  
1503  
1504  
1505  
1506  
1507  
1508  
1509  
1510  
1511  
Task Paradigm: search

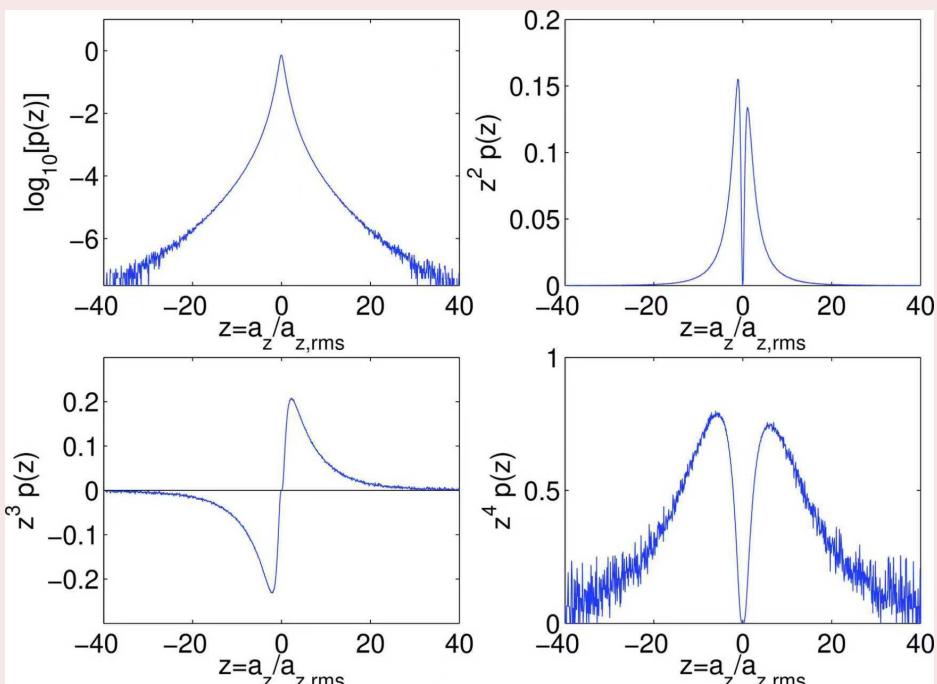
1501  
1502  
1503  
1504  
1505  
1506  
1507  
1508  
1509  
1510  
1511  
Search-oriented

1512  
1513 C.5 EXAMPLE FROM ARXIVQA1514  
1515

## One Example from ArXivQA

1516  
1517**Question ID:** physics-80491518  
1519  
1520  
1521  
1522  
1523  
1524  
1525  
1526  
1527  
1528  
1529  
1530  
1531  
1532  
1533  
1534  
1535  
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**Question:** Based on the top-right graph, how would you describe the behavior of  $P(z)$  as  $z$  approaches zero?

## Context Modalities: Images

**Covered areas:**

The document includes arXiv academic papers in various fields such as physics and mathematics.

**Covers Few Areas****Cross-evidence Reasoning:**

Only focus on a single element.

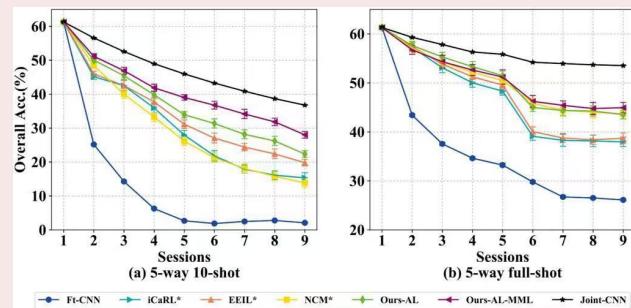
**Not Cross-evidence Reasoning****Task Paradigm: search****Search-oriented**

1566  
1567

## C.6 EXAMPLE FROM CHARXIV

1568  
1569

## One Example from Charxiv

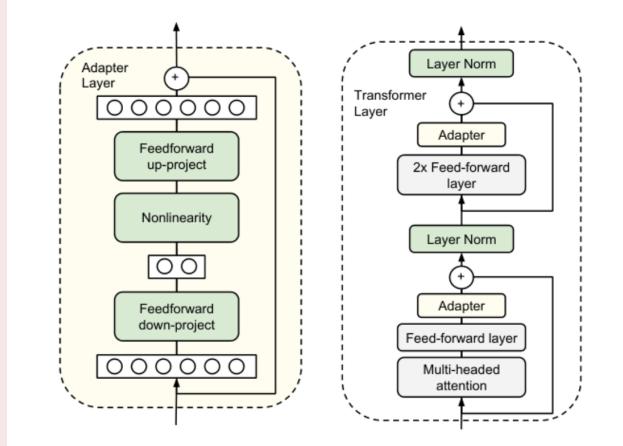
1570  
1571  
1572**Question ID:** 2004.10956**Question:** Which model shows a greater decline in accuracy from Session 1 to Session 9 in the 5-way full-shot scenario?1573  
1574**Context Modalities: Images**1575  
15761577  
1578  
1579  
1580  
1581  
1582  
1583  
1584  
1585  
1586**Covered areas:**

The document type consists of multi-type charts and graphs from 2323 papers in 8 disciplines, namely physics, computer science, mathematics, biology, chemistry, statistics, engineering, and economics, which are derived from the arXiv platform.

1591  
1592**Cross-evidence Reasoning:**

Only focus on a single element.

**Not Cross-evidence Reasoning**1593  
1594  
1595  
1596**Task Paradigm: search****Search-oriented**1600  
1601  
1602  
1603  
1604  
1605  
1606  
1607  
1608  
1609  
1610  
1611  
1612  
1613  
1614  
1615  
1616  
1617  
1618  
1619

1620 C.7 EXAMPLE FROM AAAR  
16211622 One Example from AAAR  
16231624 **Question ID:** 1902.00751  
1625 **Question:** What experiments do you suggest doing? Why do you suggest these experiments?  
16261627 **Context Modalities: Multi-Modal Documents**  
16281645 **Covered areas:**

1646 The document types cover two core categories: one is the AAAR-1.0 benchmark dataset documents, which  
1647 are used to evaluate the research capabilities of LLMs and contain annotated data for 4 types of research  
1648 tasks such as equation inference; the other is the documents related to the academic organization operation  
1649 of the American Association for Aerosol Research (AAAR).

1650 **Covers Few Areas**1651 **Cross-evidence Reasoning:**

1652 It is necessary to integrate textual evidence across paragraphs and chapters.  
1653

1656 **Task Paradigm: search**1657 **Search-oriented**

1658  
1659  
1660  
1661  
1662  
1663  
1664  
1665  
1666  
1667  
1668  
1669  
1670  
1671  
1672  
1673

1674 C.8 EXAMPLE FROM MMCR  
 1675

### One Example from MMCR

1676 **Question ID:** 1  
 1677

1678 **Question:** Which module's weights are frozen?  
 1679

#### Context Modalities: Multi-Modal Documents and Texts



#### Covered areas:

Its document type focuses on multimodal information fusion and clinical semantic understanding in medical scenarios.

**Focus Only On the Field of Medicine**

#### Cross-evidence Reasoning:

It is necessary to forcibly integrate medical imaging evidence (such as abnormal areas in CT images) with clinical report text evidence

#### Task Paradigm: search

**Search-oriented**

1726  
 1727

1728  
1729

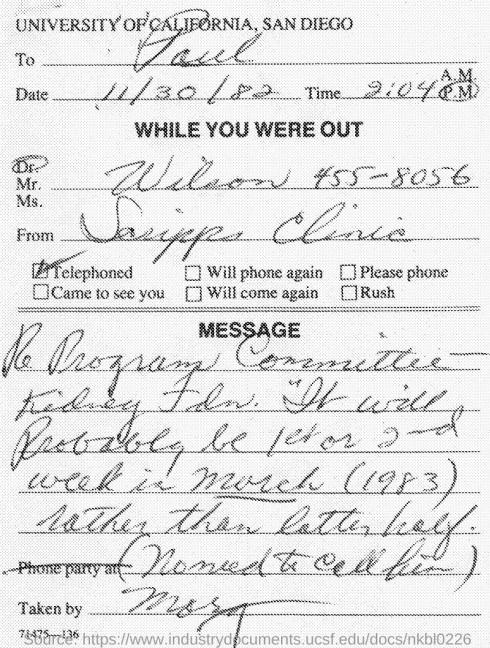
## C.9 EXAMPLE FROM DocVQA

1730  
1731

## One Example from DocVQA

1732  
1733**Question ID:** 24581**Question:** What is name of university?1734  
1735  
1736

## Context Modalities: Multi-Modal Documents

1737  
1738  
17391740  
1741  
17421743  
1744  
17451746  
17471748  
17491750  
17511752  
17531754  
17551756  
17571758  
17591760  
1761  
1762**Covered areas:**

Including invoices, resumes, academic papers, financial reports, manuals, etc. in formats such as scanned copies, PDFs, and screenshots.

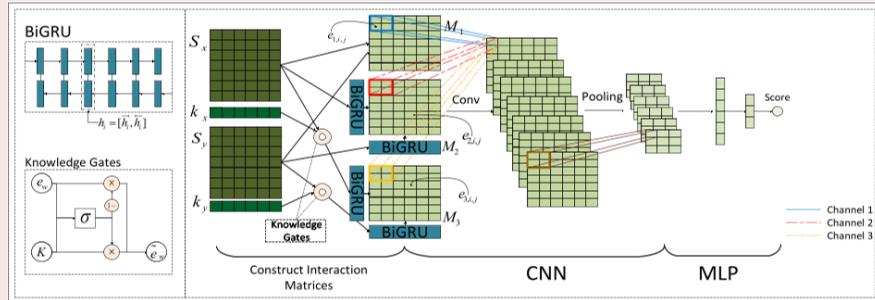
1763  
17641765  
1766  
1767**Cross-evidence Reasoning:**

Simple question types (such as "invoice amount") only require evidence from a single location, while complex question types (such as "judging device compatibility based on parameter tables and explanatory texts across multiple pages of a manual") require integrating evidence across elements and locations.

**Not All Cross-evidence Reasoning**1768  
1769  
17701771  
1772**Task Paradigm: search**1773  
17741775  
17761777  
17781779  
1780

1781

**Search-oriented**

1782 C.10 EXAMPLE FROM SPIQA  
17831784 One Example from SPIQA  
17851786 **Question ID:** 1611.04684v1  
1787 **Question:** What is the role of the knowledge gates in the KEHNN architecture?  
17881789 **Context Modalities: Multi-Modal Figures and Charts**  
17901802 **Covered areas:**  
1803 The document type of SPIQA originates from academic papers in fields such as computer science and  
1804 physics.  
18051806 **Covered Few Areas**1807 **Cross-evidence Reasoning:**  
1808 Only focus on a single element.  
18091810 **Not Cross-evidence Reasoning**1811 **Task Paradigm: search**1812 **Search-oriented**

1836 **D DATASET ANNOTATION AND CONSTRUCTION**  
18371838 **D.1 HUMAN ANNOTATOR GUIDELINES**  
18391840 The defective academic papers in our dataset are curated from three primary sources: (1) We syn-  
1841 synthetically inject 9 types of errors into papers accepted at ICLR and Nature Communications. (2)  
1842 For the papers rejected by ICLR, we identified the shortcomings in the papers based on the review-  
1843 ers' comments and categorized them into 9 error types.(3) For accepted ICLR papers, we generate  
1844 consistency-related errors by cross-referencing their content against cited literature. To ensure the  
1845 quality of each error, all entries undergo a rigorous, multistage validation protocol executed by hu-  
1846 man annotators. For synthetically generated errors, annotators manually embed them into the source  
1847 papers following this protocol:  
18481849 

- **Credibility Validation:** Each error must be logically sound and verifiable. For generated  
1850 errors, annotators first confirm their logical coherence and unambiguity. Flawed error de-  
1851 scriptions are revised whenever possible; only irreparable cases are discarded.
- **Evidence Verification:** All evidence substantiating an error must be either directly trace-  
1852 able to the source document or grounded in established domain-specific knowledge. An-  
1853 notators are required to meticulously verify the origin and accuracy of all supporting data  
1854 and background information.
- **Category Classification:** Each error must be accurately classified into one of the 9 pre-  
1855 defined categories according to their formal definitions. Annotators verify the correctness of  
1856 the assigned category and reclassify it if necessary.
- **Paper Revision:** Upon successful validation, annotators embed the generated error into the  
1857 original manuscript by adding, deleting, or modifying relevant text segments as dictated by  
1858 the error's specification.

1861  
1862 This unified and standardized annotation protocol enables the creation of a high-quality dataset of  
1863 academic papers with curated errors, providing a robust benchmark for evaluating the document  
1864 scanning and error detection capabilities of Large Multimodal Models.  
18651866 **D.2 ANNOTATION STATISTICS**  
18671868 Initially, we generated or sampled a pool of 3,500 academic paper instances containing potential  
1869 errors. During the manual annotation phase, following the protocol described above, we discarded  
1870 1,700 instances to ensure the logical rigor of the errors, the accuracy of the evidence, and a balanced  
1871 distribution of categories.1872 Of the remaining 1,800 instances, 1,541 (85.6%)underwent manual revision. The distribution of  
1873 these modifications is as follows:  
18741875 

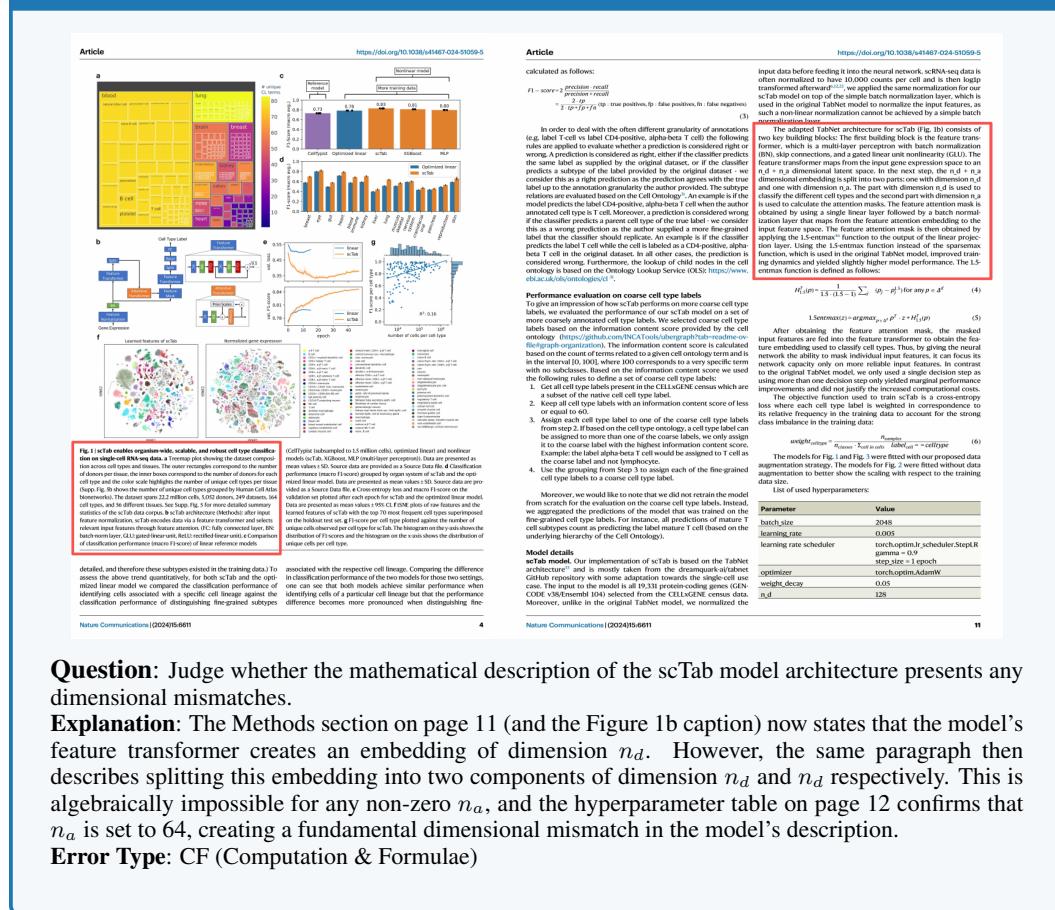
- **535 questions** were rewritten to eliminate ambiguity or to increase their retrieval and rea-  
1876 soning difficulty.
- **1,207 explanations** were revised to correct erroneous evidence references and resolve log-  
1877 ical flaws.
- **1,141 instances** underwent category reclassification or manual paper editing. This pro-  
1880 cess served to fix classifications that were inconsistent with our definitions and, for errors  
1881 generated, to manually inject them into the source papers to create the flawed documents.



1944

### D.3.2 CASE 2: MODIFY QUESTION

## Example



1976  
1977  
1978  
1979  
1980  
1981  
1982  
1983  
1984  
1985  
1986  
1987  
1988  
1989  
1990  
1991  
1992  
1993  
1994  
1995  
1996  
1997

## Before:

**Before:** **question:** Judge whether the mathematical description of the scTab model architecture presents any dimensional mismatches

Afghan

**Alter:** question: Assess the Methods section for Computation & Formulae issues

**Analysis:** Based on the error information and the text, the modified model description states that a vector of dimension  $n_d$  is split into two parts: one of dimension  $n_d$  and another of dimension  $n_a$ . This is algebraically impossible, as the total dimension ( $n_d$ ) cannot equal the dimension of one of its parts ( $n_d$ ) plus another non-zero part ( $n_a$  is set to 64). This constitutes a clear dimensional mismatch, rendering the model's architectural description logically invalid. The original question was overly specific, as it explicitly prompted an assessment of whether the mathematical description of the scTab model architecture contained 'any dimensional mismatches'. This hint was too detailed, reducing the analytical difficulty for the model. To increase the difficulty, we have revised the question's phrasing to ask only whether the mathematical description of the scTab model architecture presents any problems.



2052  
2053

## D.3.4 CASE 4: MODIFY CATEGORY

2054  
2055

2056

2057

2058

2059

2060

2061

2062

2063

2064

2065

2066

2067

2068

2069

2070

2071

2072

2073

2074

2075

2076

2077

2078

2079

2080

2081

2082

2083

2084

2085

2086

2087

2088

2089

2090

2091

2092

2093

2094

2095

2096

2097

2098

2099

2100

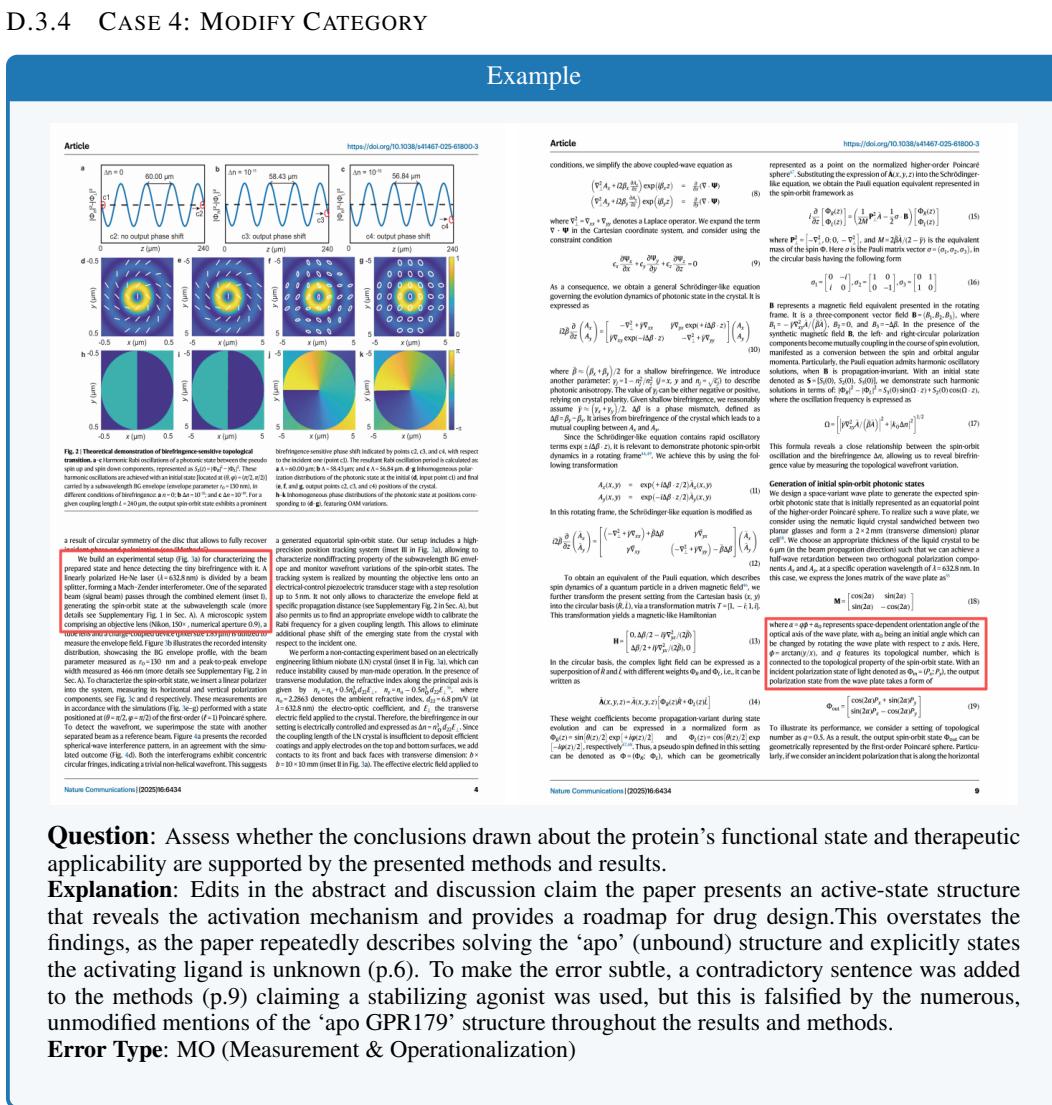
2101

2102

2103

2104

2105



**Question:** Assess whether the conclusions drawn about the protein's functional state and therapeutic applicability are supported by the presented methods and results.

**Explanation:** Edits in the abstract and discussion claim the paper presents an active-state structure that reveals the activation mechanism and provides a roadmap for drug design. This overstates the findings, as the paper repeatedly describes solving the 'apo' (unbound) structure and explicitly states the activating ligand is unknown (p.6). To make the error subtle, a contradictory sentence was added to the methods (p.9) claiming a stabilizing agonist was used, but this is falsified by the numerous, unmodified mentions of the 'apo GPR179' structure throughout the results and methods.

**Error Type:** MO (Measurement & Operationalization)

Before:

Error Type: MO (Measurement & Operationalization)

Decision: Modify

After:

Error Type: RCA (Referential and Citation Alignment)

**Analysis:** The introduced error systematically changes the laser wavelength used in the experiment to 532.0 nm. However, the calculation of a key physical quantity (birefringence) continues to use material constants (the electro-optic coefficient) that are only valid at the old wavelength of 632.8 nm. Because the optical properties of materials are wavelength-dependent, this systematic mismatch between experimental conditions and calculation parameters creates a significant contradiction in a core part of the paper. Compared to a Measurement & Operationalization (MO) error, this error is more accurately described as an internal inconsistency. Therefore, we are reclassifying this question from MO to RCA.



2160

2161

2162

2163

2164

2165

2166

2167

2168

2169

2170

2171

2172

2173

2174

2175

2176

2177

2178

2179

2180

2181

2182

## Example2: 816

Under review as a conference paper at ICLR 2025

216 *In addition (Equation 3), it holds:*  
 217 
$$\mathbb{E} [I_r \otimes \frac{1}{n} g\theta^T] = \mathbb{E}[H(\text{vec}(\theta))].$$
  
 218  
 219

220 *Following Lemma Theorem 3.1, we further estimate a hessian-gradient product using GFM, cor-*

221 *responding to the  $(H(\theta^*) \otimes \nabla^2 p)$  terms in Equation 1. Given an invertible matrix  $A$ , we have*

222 
$$(I_r \otimes A)^{-1} = I_r \otimes A^{-1}.$$
 Therefore, denote the GFM matrix as  $G(\theta) \triangleq (g\theta^T) \in \mathbb{R}^{r \times d}$  for any

223 *matrix  $\theta \in \mathbb{R}^{d \times r}$ , then it holds that:*

224 
$$H(\text{vec}(\theta))^T \text{vec}(v) = \left[ I_r \otimes \left( \frac{1}{n} g\theta^T \right)^T \right]^T \text{vec}(v) = \text{vec}(G(\theta)^{-1} v).$$
 (4)

225 *Consider a LoRA model with LoRA dimension  $d$  and rank  $r$ . We assume that each column*

226 *in one LoRA block  $\Delta W \in \mathbb{R}^{d \times r}$ , corresponding to each rank, is i.i.d. distributed with zero mean.*

227 *In the ideal case that the model is trained to converge with  $\mathbb{E}[\nabla_y \log p(y|x, \theta)] = 0$ , the zero-*

228 *mean gradient  $\nabla_y \log p(y|x, \theta)$  is zero, then we can use the GFM to approximate the Hessian gradient product to approximate the original Hessian-gradient product. To further guarantee that  $G(\theta)$  is invertible, we*

229 *add a damping factor  $\lambda_d$  to the GFM matrix following Martens (2012).*

230 *We implement the GFM in Equation 4. Then derive the final form of HYPERINT influence score:*

231 *On a specific datapoint  $\{x_{n+1}\} \in \mathbb{R}^{d \times 1}$ , denote the influence gradient on a parameter block  $\theta$  as*

232  *$g_{\theta}(0)$ , we compute:*

$$f_{\text{HYPERINT}}(x_n, y_n) := -g_{\theta}^T (G(\theta^*) + \lambda_d I)^{-1} g_{\theta}(\theta), \quad (5)$$

233 *where  $g_{\theta} = \frac{1}{n} \sum_{i=1}^n \nabla_x f_i(y_i^{(n)}, f_{\theta}(x_i^{(n)})) \in \mathbb{R}^{d \times r}$ , representing the average unflattened gra-*

234 *dient of  $\theta$  on the validation set.*

## 3.2 MATRIX INVERSE APPROXIMATION WITH SCHULZ'S METHOD

235 *Schulz's method (Petković, 1989). To compute the inverse of one matrix  $A$ , the hyperpower*

236 *iteration approach has attracted the attention of many researchers due to its*

237 *rigorous convergence guarantee (Bazin & Boos, 2018).*

$$X_{t+1} = X_t(I + T_t + T_t^2 + \dots + T_t^{p-1}), \quad T_t = I - A X_t \quad (6)$$

238 *The iterative approach requires  $p$  matrix-matrix multiplications per iteration and has an order of*

239 *convergence  $p$  (Bazin & Boos, 2018). When choosing  $p = 2$ , it yields the Schulz iteration, which*

240 *can be regarded as a by-product of the Newton method applied to the non-linear equation  $f(X) =$*

$$A - X^{-1} \quad (7)$$

241 *It is proved by Ben-Israel & Cohen (1966) and Petković (1995) that with a proper initialization,*

242 *Schulz's method would converge to  $A^{-1}$  in the order of convergence at least  $p = 2$ . We provide*

243 *the pseudocode of Schulz's iteration for influence function computation. First, we compute the inverse*

244 *of the GFM in Equation 4 with Schulz's iterations (Equation 7); Last, we compute the influence*

245 *score with cached validation gradient  $v$  and the influence gradient on each training sample, i.e.*

246  *$f_{\text{HYPERINT}}(x_{n+1})$  (Equation 5). We provide the detailed pseudocode in the Appendix (Algo. 3).*

5

247 **Question:** Scrutinize the Methods section for Research Question & Definitions issues.

248 **Explanation:** Lemma 3.1, which is a cornerstone of the paper's theoretical contribution for low-rank  
 249 Hessian approximation, relies on a strong and insufficiently justified assumption. The lemma states:  
 250 “Assume that each column of the sample gradient ... is independent and identically distributed random  
 251 vector with zero mean under the distribution  $p(y|x, \theta)$ ”. The paper provides only a brief, hand-wavy  
 252 justification (p.5, lines 230-232), suggesting it “could stand” in an “ideal case” of model convergence.  
 253 These critical i.i.d. and zero-mean conditions are not rigorously established or empirically validated  
 254 for the contexts in which the method is applied. This leaves a core hypothesis of the paper ambiguously  
 255 defined and justified, which is an error of type Research Question & Definitions.

256 **Error Type:** RQD (Research Question & Definitions)

257 **Type:** Within-Sample

292

293

294

295

296

297

298

299

200

201

202

203

204

205

206

207

208

209

210

211

212

213

2214  
2215

## E.2 DI (DESIGN &amp; IDENTIFIABILITY)

2216  
2217

2218

2219

2220

2221

2222

2223

2224

2225

2226

2227

2228

2229

2230

2231

2232

2233

2234

2235

2236

2237

2238

2239

2240

2241

2242

2243

2244

2245

2246

2247

2248

2249

2250

2251

2252

2253

2254

2255

2256

2257

2258

2259

2260

2261

2262

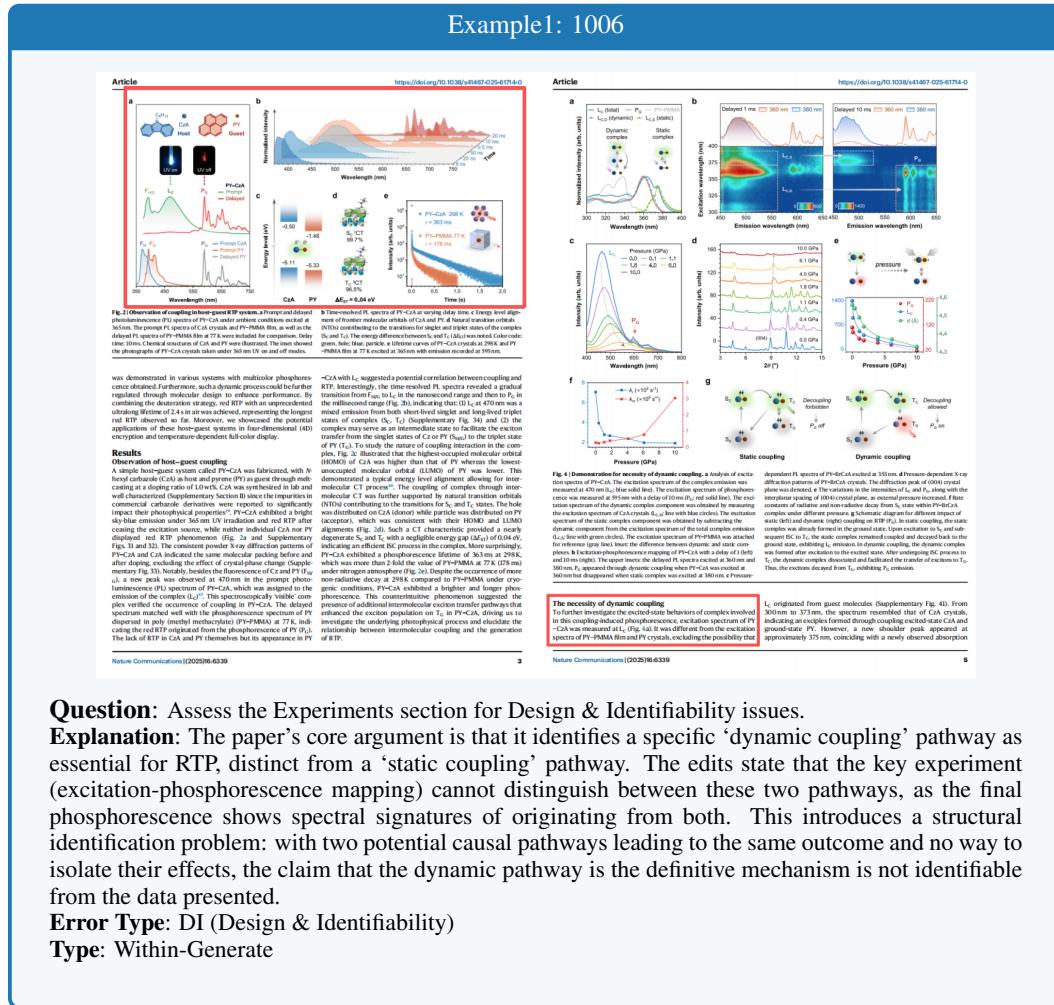
2263

2264

2265

2266

2267



2268

2269

2270

2271

2272

2273

2274

2275

2276

2277

2278

2279

2280

2281

2282

2283

2284

2285

2286

2287

2288

2289

2290

2291

2292

2293

2294

2295

2296

2297

2298

2299

2300

2301

2302

2303

2304

2305

2306

2307

2308

2309

2310

2311

2312

2313

2314

2315

2316

2317

2318

2319

2320

2321

## Example2: 724

Under review as a conference paper at ICLR 2025

$$W = USV^T \rightarrow \hat{S}_i = \begin{cases} \rho_i & \text{for } i \neq r \\ 0 & \text{else} \end{cases} \rightarrow \hat{W} = \hat{U}\hat{S}\hat{V}^T. \quad (9)$$

Because removing a single singular value in a full transformer model has negligible effect, we grouped the singular values of each matrix into ten equally sized sets and removed these sets in order of decreasing magnitude. We then fine-tuned the model on the BoBoQ dataset, achieving an average validation accuracy of 73.0%. We then removed one of the sets of singular values, and then removed one of the sets of singular values again, and so on, until all of the singular values in each query matrix to zero and measured how the validation accuracy dropped.

We present the results in Figure 7, which shows good agreement between the regions that deviate from RMT and the regions that are crucial for the transformer's performance. As expected, for all matrix types, the removal of the largest singular value leads to the greatest accuracy drops. This is consistent with the intuition that the right singular vectors are the ones with the most extreme singular values in  $W$  for vectors corresponding to the largest singular values. As a reference guide, in the case of the intermediate Dense matrices, the singular vectors corresponding to small singular values have the largest deviations from RMT. This is reflected in a large accuracy drop when removing the first set of singular values. For the Dense matrices, the key matrices also exhibit significant RMT deviations for singular vectors corresponding to the smallest singular values. For the Dense matrices, the validation accuracy drops when we remove the key matrices on the BoBoQ dataset, but we did not observe a significant effect on the overall performance. Such behavior is expected when information learned during pretraining is not used by the model (see also Appendix D) for example, on the BoBoQ dataset, where removing these small singular values impacts performance.

Although one might consider using this scheme to reduce the network's size, we find that removing large singular values does not significantly reduce the size of the network while degrading the network's performance. To understand this behavior, we consider the case where a weight matrix in the network architecture is initialized with random weights. In this case, the network is not able to learn, but the removal of small and intermediate singular values from the random weights significantly impacts the overall performance, as the subsequent layers are sensitive to small changes in the input. This is consistent with the observation that the BoBoQ dataset is highly linear, but that features are highly resilient to such removal, whereas removing singular values from a random matrix destroys the subtle details that subsequent layers depend on.

## 7 FINE-TUNING

Recent studies have debated the relevance of small singular values in transformer networks. Some argue that these values are crucial for network performance (Hou et al., 2022), while others have observed performance improvements when they are removed (Shen et al., 2023). Our RMT analysis provides sufficient evidence only for some of the smaller singular values in the weight matrices, providing a diagnostic tool to assess their importance. This finding supports the notion that

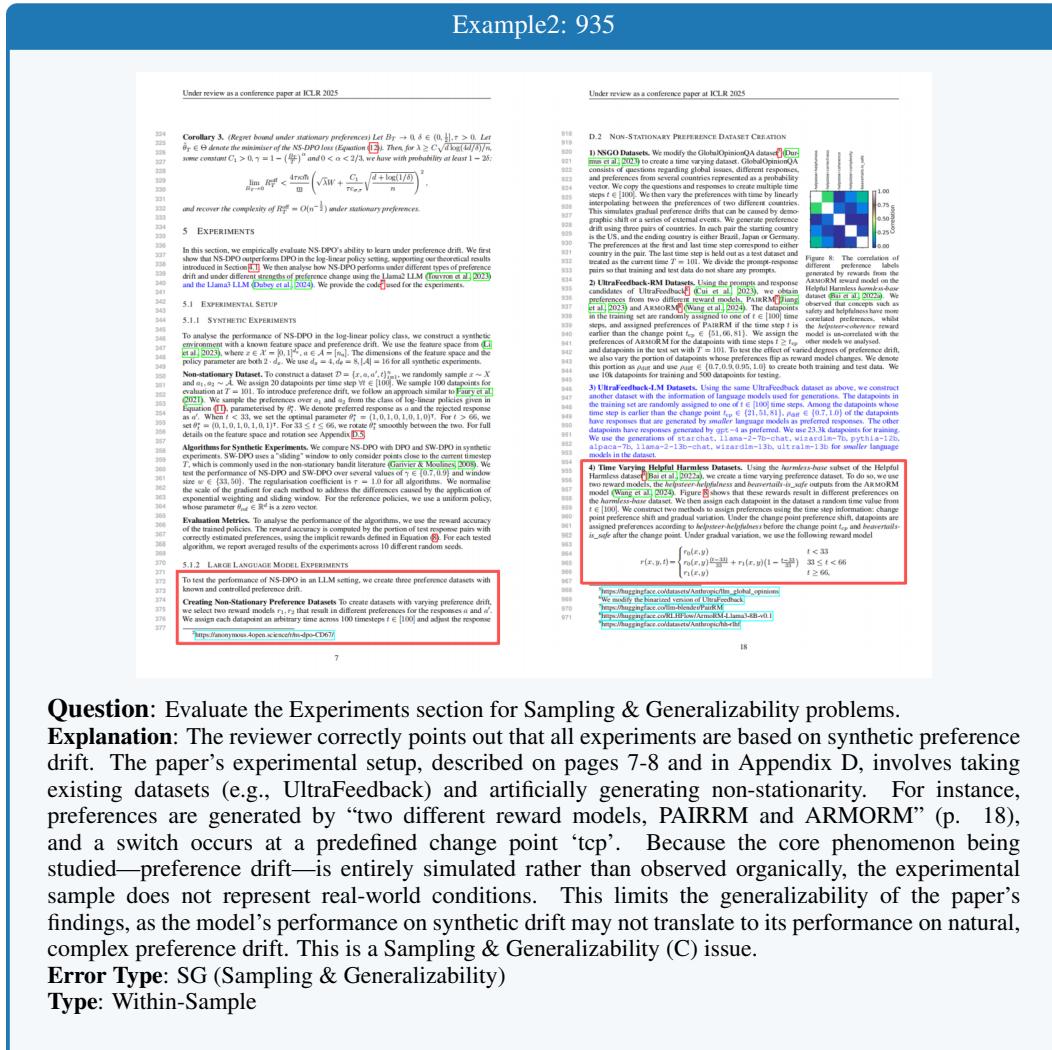
**Question:** Assess the Methods section for Design & Identifiability issues.

**Explanation:** A reviewer points out a flaw in the experimental design for the pruning experiments. The paper states on page 9, “We then removed one of the singular value deciles from a specific matrix type in all layers”. The reviewer argues this “coarse intervention” constitutes a design flaw because by modifying all layers simultaneously, it becomes impossible to attribute performance changes to specific layers. This confounds the effects, making it difficult to identify where in the model the removed information was critical. This directly undermines the paper’s stated goal of “locating information.” The design choice violates the conditions for identifying layer-specific contributions, which is an error of type Design & Identifiability.

**Error Type:** DI (Design & Identifiability)

**Type:** Within-Sample



2376  
2377  
2378  
2379  
2380  
2381  
2382  
2383  
2384  
2385  
2386  
2387  
2388  
2389  
2390  
2391  
2392  
2393  
2394  
2395  
2396  
2397  
2398  
2399  
2400  
2401  
2402  
2403  
2404  
2405  
2406  
2407  
2408  
2409  
2410  
2411  
2412  
2413  
2414  
2415  
2416  
2417  
2418  
2419  
2420  
2421  
2422  
2423  
2424  
2425  
2426  
2427  
2428  
2429

2430  
2431

## E.4 RCA (REFERENTIAL AND CITATION ALIGNMENT)

2432  
2433

## Example1: 0

2434  
2435

Published as a conference paper at ICLR 2025

Published as a conference paper at ICLR 2021

Second, they assume the sensors form a fixed spatial grid and record data simultaneously, which is not the case in our setup where the data come one at a time as a single sensor at each time point (Figure 1). Third, they assume the sensors are static and record data on a regular rectangular spatial-temporal grid (Long et al., 2018; Geneva & Zabaras, 2020), other methods work with irregular, moving, and non-overlapping spatial-temporal grids (Wu et al., 2021; Chen et al., 2021; Yin et al., 2023) to go further and allow the observation locations to change over time (Wu et al., 2021; Yin et al., 2023) but fix the observation times and assume dense observations. Whereas, another line of research (Chen et al., 2021; Yin et al., 2023; Wu et al., 2021) proposes to learn the system dynamics by the system dynamics (Chen et al., 2021; Yin et al., 2023; Zhou et al., 2023; Wu et al., 2021).

This work fills this gap and proposes a model for randomly observed spatio-temporal dynamical systems. Our model incorporates techniques from amortized variational inference (Kingma & Welling, 2013), neural differential equations (Chen et al., 2020; Rackauckas et al., 2020), neural ODEs (Chen et al., 2021; Yin et al., 2023) and neural flows (Mathieu & Nickel, 2020; Chen et al., 2023; Yin et al., 2023) to efficiently learn both the underlying system dynamics and the random spatio-temporal point processes. The proposed model, Attentive CNF (Chen et al., 2021), simulates the latent trajectory with neural ODEs (Chen et al., 2018), and uses implicit neural representations to parameterize the point process. The proposed model is able to learn the system dynamics and the latent trajectory with the latent state evaluation and propose a technique to alleviate it, resulting in up to 4x faster training. Our model shows strong empirical results outperforming other models from the literature on challenging spatio-temporal datasets.

2. Published

## 2.1 SPATIOTEMPORAL POINT PROCESSES

Spatiotemporal point processes (STPP) model sequences of events occurring in space and time. Each event has an associated event time  $t_i \in \mathbb{R}_{\geq 0}$  and event location  $\mathbf{x}_i \in \mathbb{R}^d$ . Given an event history  $\mathcal{H}_t = \{(t_i, \mathbf{x}_i) | t_i < t\}$  with all events up to time  $t$ , we can characterize STPP by its conditional intensity function

$$\lambda^*(t, \mathbf{x}) \triangleq \sup_{\delta t, \delta \mathbf{x}} \frac{\Pr\{\mathbf{e} \in [t, t+\delta t], \mathbf{x} \in B_\delta(\mathbf{x}) | \mathcal{H}_t\}}{\delta t B_\delta(\mathbf{x})} \quad (1)$$

where  $\delta t$  denotes an infinitesimal time interval, and  $B_\delta(\mathbf{x})$  denotes a  $\delta$ -ball centered at  $\mathbf{x}$ . Given a sequence of  $N$  events  $\{(t_i, \mathbf{x}_i)\}_{i=1}^N$  on a bounded domain  $A \subset [0, T] \times \mathbb{R}^d$ , the log-likelihood for the STPP is evaluated as (Daley et al., 2000)

$$\log p(\{(t_i, \mathbf{x}_i)\}_{i=1}^N) = \sum_{i=1}^N \log \lambda^*(t_i, \mathbf{x}_i) - \int_0^T \lambda^*(t, \mathbf{x}) dt. \quad (2)$$

Marked STPP extends the above simple STPP by a mark  $y_i \in \mathbb{R}^d$  that is associated to each event  $(t_i, \mathbf{x}_i)$ .

## 2.2 ORDINARY AND PARTIAL DIFFERENTIAL EQUATIONS

Given a deterministic continuous-time dynamic system with state  $\mathbf{z}(t) \in \mathbb{R}^d$ , we can describe the evolution of its state in terms of an ordinary differential equation (ODE)

$$\frac{d\mathbf{z}(t)}{dt} = f(t, \mathbf{z}(t)). \quad (3)$$

For an initial state  $\mathbf{z}_0$  at time  $t_0$ , we can solve the ODE to obtain the system state  $\mathbf{z}(t)$  at later times  $t > t_0$ . The solution exists and is unique if  $f$  is continuous in time and Lipschitz continuous in state (Coddington & Levinson, 1956), and can be obtained either analytically or using numerical ODE solvers (Bauer et al., 1987). In this work we solve ODEs numerically using ODE solvers from

Model	Earthquakes JP		COVID-19 NJ		BOLD5000	
	Temporal	Spatial	Temporal	Spatial	Temporal	Spatial
Finance Process	-0.794e-005	-	0.111e-005	-	0.979e-005	-
Self-correcting Process	-2.17e-005	-	7.051e-006	-	-0.00054e+000	-
Barlow Hawkes Process	-0.00012e+000	-	-0.00019e+000	-	-0.00027e+000	-
Neural Hawkes Process	-0.00212e+000	-	-0.00019e+000	-	-0.00027e+000	-
Conditional KDE	-	-	-0.259e-005	-	-0.301e-005	-
Time-varying CNF	-	-	-0.259e-005	-	-0.301e-005	-
Neural Jump SDE (GR)	-0.00066e+000	-0.00077e+000	0.00018e+000	-0.00018e+000	0.251e-005	-0.251e-005
Neural Jump SDE (GR)	-0.00066e+000	-0.00077e+000	0.00018e+000	-0.00018e+000	0.251e-005	-0.251e-005
Neural Jump SDE (GR)	-0.00066e+000	-0.00077e+000	0.00018e+000	-0.00018e+000	0.251e-005	-0.251e-005
Attentive CNF	0.00048e+000	-0.257e-005	0.00029e+000	-0.237e-005	0.259e-005	0.306e-005
Attentive CNF	0.00048e+000	-0.257e-005	0.00029e+000	-0.237e-005	0.259e-005	0.306e-005

Table 1: Log-likelihood per event on held-out test data (higher is better). Standard devs. estimated over 5 runs.

**Results & Analysis** The results of our evaluation are shown in Table 1. We highlight all results where the intervals containing one standard deviation away from the mean overlap.

Across all data sets, the Time-varying CNF outperforms the conditional KDE baseline despite not being conditioned on history. This suggests that the use of spatial derivatives is not necessary. We also highlight that the Attentive CNF is able to compute the log-likelihood of future events by learning a large band width whereas a flexible CNF can easily model multi-modal event propagation.

The Jump and Attentive CNF models achieve better log-likelihoods than the Time-varying CNF, suggesting predictions are more accurate (Table 1 and Figure 10).

For COVID-19, the self-exciting Hawkes process is a strong baseline which aligns with similar results for other infectious diseases (Park et al., 2019), but Neural STPPs can achieve substantially better spatial and temporal log-likelihoods. However, the Attentive CNF is the only model that can learn the latent state and propagate the latent state to future events. However, it tends to fall short of the Attentive CNF which jointly models spatial and temporal variables.

In a closer comparison to the temporal likelihood of Neural Jump SDEs (Ja & Benson, 2019), we find that over-parametrized spatial models can outperform temporal models since both domains are tightly coupled. Similarly, the performance of Neural Jump SDEs and STPPs with a continuous-time training architecture to model the temporal domain, the temporal likelihood values are often close. However, there is a significant performance difference between the Attentive CNF and the Attentive STPP models and Neural Jump SDEs even for the temporal log-likelihood on all data sets.

Finally, we note that the results of the Jump and Attentive CNFs are typically close. The Attentive model generally achieves better log-likelihoods than the Jump model. This is likely due to the spatial and temporal differences in how the Attentive CNF's ability to temporal to all previous events, the Jump model has to compress all history information inside the hidden state at the time of the event. The Attentive CNF also enjoys substantially faster computations (see Appendix A).

## 6 CONCLUSION

To learn high-fidelity models of stochastic events occurring in continuous space and time, we have proposed a new framework for spatio-temporal event prediction and approximation conditioned on history. The Neural Jump SDEs with Continuous Normalizing Flows and allows us to learn the flexibility of neural temporal point processes while enabling highly expressive models of continuous space and time. The Attentive CNF is able to learn the latent state and propagate the latent state to future events which others are concerned with modeling highly complex spatio-temporal data. In this context, the use of Riemannian CNFs (Mathieu & Nickel, 2020; Lou et al., 2020; Faloutsos & Feng, 2020) is especially interesting. Specifically, the Attentive CNF can be approximated using a Riemannian manifold (e.g. the Earth's surface) by simply replacing the CNF in our models with a Riemannian counterpart.

2453  
2454 **Question:** Scan the errors in cited reference Chen et al.(2021)

**Explanation:** The edited P contains a Type H error by misrepresenting the performance of the cited model. P (p. 8) claims that the NSTPP model from Chen et al. (2021) ‘reported performance comparable to a standard Hawkes process baseline’. This contradicts the results in S, where the proposed models (i.e., NSTPP) consistently outperform the Hawkes process baseline, often by a large margin. For example, S (p. 9, Table 1) shows on the BOLD5000 dataset that the ‘Attentive CNF’ model achieves a temporal log-likelihood of  $5.842 \pm 0.005$ , which is substantially better than the Hawkes Process at  $2.860 \pm 0.050$ .

**Error Type:** RCA (Referential and Citation Alignment)

**Type:** Cross-Generate

2464  
2465  
2466  
2467  
2468  
2469  
2470  
2471  
2472  
2473  
2474  
2475  
2476  
2477  
2478  
2479  
2480  
2481  
2482  
2483



2538  
2539

## E.5 MO (MEASUREMENT &amp; OPERATIONALIZATION)

2540  
2541

2542

2543

2544

2545

2546

2547

2548

2549

2550

2551

2552

2553

2554

2555

2556

2557

2558

2559

2560

2561

2562

2563

2564

2565

2566

2567

2568

2569

2570

2571

2572

2573

2574

2575

2576

2577

2578

2579

2580

2581

2582

2583

2584

2585

2586

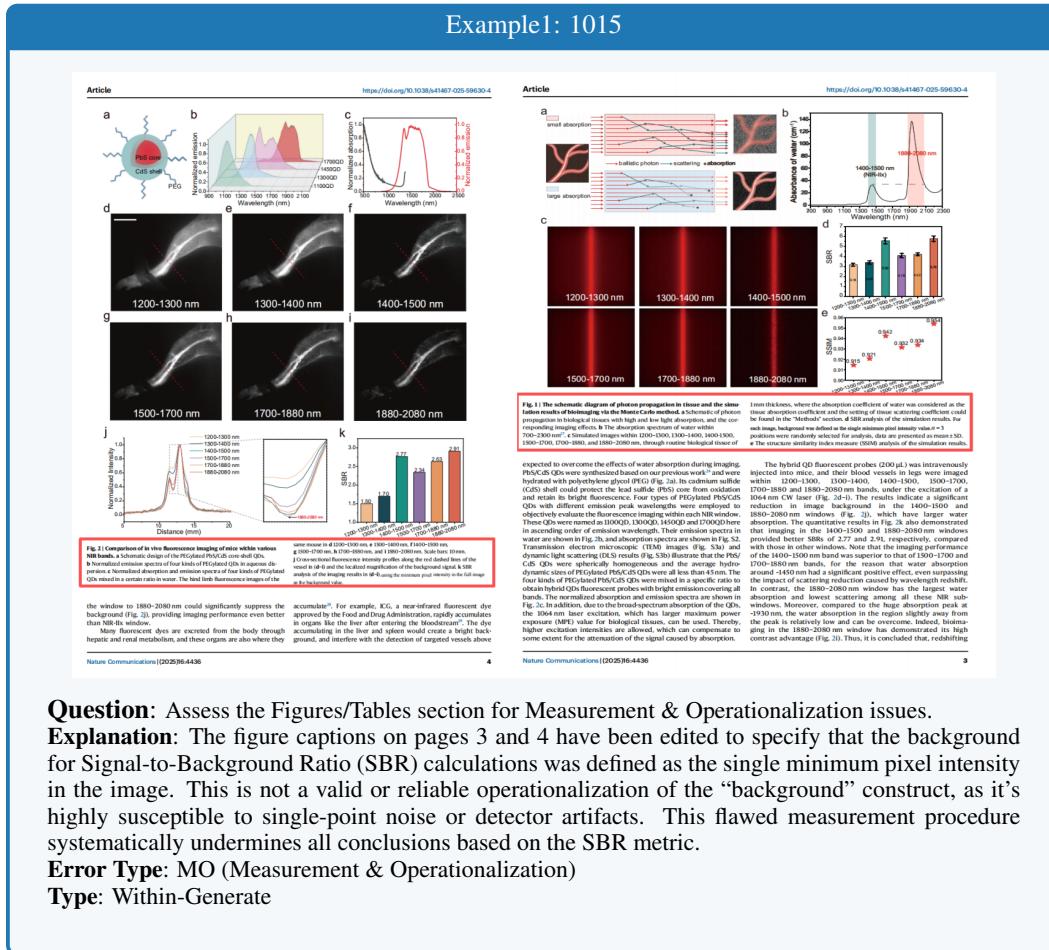
2587

2588

2589

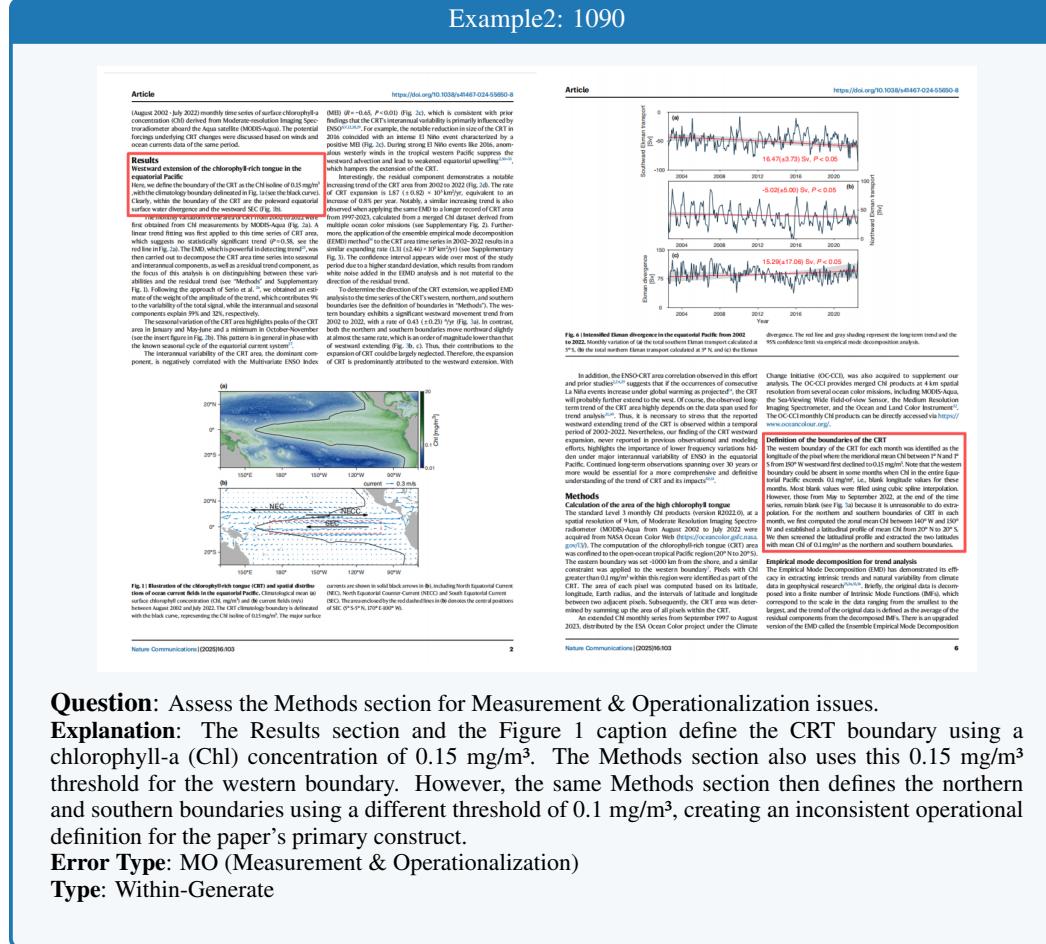
2590

2591



2592  
2593  
2594  
2595  
2596  
2597  
2598  
2599  
2600  
2601  
2602  
2603  
2604  
2605  
2606  
2607  
2608  
2609  
2610  
2611  
2612  
2613  
2614  
2615  
2616  
2617  
2618  
2619  
2620  
2621  
2622  
2623  
2624  
2625  
2626  
2627  
2628  
2629  
2630  
2631  
2632  
2633  
2634  
2635  
2636  
2637  
2638  
2639  
2640  
2641  
2642  
2643  
2644  
2645

## Example2: 1090



**Question:** Assess the Methods section for Measurement & Operationalization issues.

**Explanation:** The Results section and the Figure 1 caption define the CRT boundary using a chlorophyll-a (Chl) concentration of  $0.15 \text{ mg/m}^3$ . The Methods section also uses this  $0.15 \text{ mg/m}^3$  threshold for the western boundary. However, the same Methods section then defines the northern and southern boundaries using a different threshold of  $0.1 \text{ mg/m}^3$ , creating an inconsistent operational definition for the paper's primary construct.

## Error Type: MO (Measurement & Operationalization)

Type: Within-Generate

2646 **E.6 DHP (DATA HANDLING & PREPROCESSING)**  
26472648 **Example1: 528**

Under review as a conference paper at ICLR 2026

---

216 the likelihood that the output configuration will satisfy the desired constraints. In our case, the code  
217 means in energy, the output of the energy network, must decrease. We specifically employed GRU to  
218 directly update the task-net's predictions using gradient signals derived from the energy network.  
219  
220 The implementation of GRU involves three main steps. The task-net, serving as our baseline model,  
221 is trained in a supervised manner to predict 3D poses. Next, a separate energy network is trained  
222 using the task-net's predictions as input. Finally, the trained energy network is employed to iteratively update  
223 the task-net's predictions through gradient-based optimization.

---

224

225 **Algorithm 1 Gradient-Based Inference**

---

226 **Require:**  $\langle x, y \rangle$ , training data (2D inputs and 3D ground truth outputs)

227 **Require:**  $F_x$ , task-net;  $E_y$ , energy network

228 **Require:**  $\eta_\phi$ , learning rate parameter for  $F_x$

229 **Require:**  $T$ , training iterations;  $\lambda$ , GRU steps

230

1: **Phase 1: train Task-Net**

2: for  $t = 1$  to  $T$  do

3:   Sample batch  $B_t = \{(x_t, y_t)\}_{i=1}^N$

4:   Update  $\phi = \phi - \eta_\phi \nabla_\phi \left[ \frac{1}{|B_t|} \sum_{i \in B_t} \text{MSE}(F_x(x_i) - y_i) \right]$

5: end for

231 **Phase 2: train Energy network**

7: for  $t = 1$  to  $T$  do

8:   Sample batch  $B_t = \{(x_t, y_t)\}_{i=1}^N$

9:   Generate  $\hat{y}_t = F_x(x_t)$  for  $x_t \in B_t$

10:   Update  $\hat{y}_t = \hat{y}_t - \eta_\phi \nabla_{\hat{y}_t} \left[ \frac{1}{|B_t|} \sum_{i \in B_t} \text{MSE}(E_y(x_i, \hat{y}_t) - E_y(x_i, y_i)) \right]$

11: end for

12: **Phase 3: gradient-based inference**

13: for  $t = 1$  to  $T$  do

14:   for  $i = 1$  to  $K$  do

15:      $y_i^{(k-1)} \leftarrow y_i^{(k-1)} - \eta \nabla_{\hat{y}_t} E_y(x_i, \hat{y}_t^{(k-1)})$

16:   end for

---

246

247 **3.3 SETTING**

248 **Dataset:** We conduct our experiments on Human3.6M 3D WholeBody dataset (H3WB) ([Zhao et al. 2022](#)) and Human3.6M dataset (H36M) ([Freeman et al. 2010](#)). H36M is one of the most  
249 widely used datasets for 3D human pose estimation ([Zheng et al. 2023](#); [Jia et al. 2023](#)). H3WB  
250 extends H36M to provide more annotations, capturing detailed information about hands, which  
251 includes 133 whole-body keypoint annotations, capturing detailed information about hands, which  
252 makes it suitable for tasks that require fine-grained pose estimation. We utilize the ground  
253 truth annotations for the H36M dataset and the 3D annotations for the H3WB dataset. For the H36M  
254 dataset, subjects S9 and S11 constitute the test set. Tuning hyperparameters directly on the test set  
255 introduces data leakage, leading to an optimistic bias in the reported results and invalidating claims  
256 of generalization. This is a critical violation of machine learning best practices and fits the Data  
257 Handling & Preprocessing (E) category, as a pipeline choice introduces bias.

258 **Error Type:** DHP (Data Handling & Preprocessing)

259 **Type:** Within-Sample

2669 **Question:** Assess the Methods section for Data Handling & Preprocessing issues.  
2670 **Explanation:** The reviewer correctly identifies that the authors tuned hyperparameters on the  
2671 test set. The paper's "Implementation Details" section on page 5 states: "For hyperparameter  
2672 tuning, we employed Bayesian optimization with the wandb sweep tool (Biewald, 2020), aiming to  
2673 minimize MPJPE for the S9 and S11 in the H36M dataset and PA-MPJPE for the S8 in the H3WB  
2674 dataset, following the convention of prior works." According to standard protocols for the H36M  
2675 dataset, subjects S9 and S11 constitute the test set. Tuning hyperparameters directly on the test set  
2676 introduces data leakage, leading to an optimistic bias in the reported results and invalidating claims  
2677 of generalization. This is a critical violation of machine learning best practices and fits the Data  
2678 Handling & Preprocessing (E) category, as a pipeline choice introduces bias.

2679 **Error Type:** DHP (Data Handling & Preprocessing)  
2680 **Type:** Within-Sample

2681  
2682  
2683  
2684  
2685  
2686  
2687  
2688  
2689  
2690  
2691  
2692  
2693  
2694  
2695  
2696  
2697  
2698  
2699



2754  
2755

## E.7 CF (COMPUTATION &amp; FORMULAE)

2756  
2757

2758

2759

2760

2761

2762

2763

2764

2765

2766

2767

2768

2769

2770

2771

2772

2773

2774

2775

2776

2777

2778

2779

2780

2781

2782

2783

2784

2785

2786

2787

2788

2789

2790

2791

2792

2793

2794

2795

2796

2797

2798

2799

2800

2801

2802

2803

2804

2805

2806

2807

2756  
2757  
2758  
2759  
2760  
2761  
2762  
2763  
2764  
2765  
2766  
2767  
2768  
2769  
2770  
2771  
2772  
2773  
2774  
2775  
2776  
2777  
2778  
2779  
2780  
2781  
2782  
2783  
2784  
2785  
2786  
2787  
2788  
2789  
2790  
2791  
2792  
2793  
2794  
2795  
2796  
2797  
2798  
2799  
2800  
2801  
2802  
2803  
2804  
2805  
2806  
2807

Example1: 350

Figure 2. Comparison of applying different easy-hard pairs for training a Simple-Netwrok model as in Fig. 12. We can sort these pairs by their performance and load them on either best mAP results or distinct convergence curves.

**II** is  $82 \times 9 \times 4$  downsampled 4 times. For inference, each keypoint is located by finding the pixel with largest value in the predicted heatmap. We denote the labeled and unlabeled training sets as  $T_{labeled}$  and  $T_{unlabeled}$ , respectively.  $T_{labeled} = \{I_l, T_l\}$  and  $T_{unlabeled} = \{I_u, T_u\}$  are the labeled and unlabeled image sample, respectively. And  $N$  or  $M$  is the total number of image samples. The  $\mathbf{H}^i$  are ground truth heatmap generated using 2D keypoints. For supervised training of the network  $f$ , we calculate the MSE loss:

$$L_m = \mathbb{E}_{I_l \sim p(I_l)} \|f(I_l, \mathbf{H}^i) - T_l(\mathbf{H}^i)\|^2, \quad (1)$$

where  $T_l(\mathbf{H}^i)$  represents an easy affine augmentation including a random rotation angle from  $[-30^\circ, 30^\circ]$  and a scale factor from  $[0.75, 1.25]$  (denoted as  $T_{l,0}$ ). For unlabeled images, we calculate the most permissive consistency loss:

$$L_u = \mathbb{E}_{I_u \sim p(I_u)} \|\mathbb{E}_{T_u \sim p(T_u)} [f(I_u, T_u(\mathbf{H}^i)) - f(I_u, \mathbf{H}^i)]\|^2, \quad (2)$$

where  $T_u(\mathbf{H}^i)$  is a harder augmentation with strong perturbations than affine-based  $T_u$ . The  $T_{r,0}$  means the random rotation angle from  $[-30^\circ, 30^\circ]$  and  $T_{s,0}$  means the scale factor from  $[0.75, 1.25]$  (denoted as  $T_{u,0}$ ). In this way, we can obtain a paired easy-hard augmentation  $(I_l, T_l) = (T_u(I_l), T_u(\mathbf{H}^i))$  for generating the labeled training set. The joint training of labeled and unlabeled training set can avoid propagation of teacher signals to avoid collapsing. Next, we answer two questions Q1 and Q2 by extensive empirical studies in Sec. 3.1 and Sec. 3.3, respectively. After that, we provide a theoretical perspective for understanding the power of designing stronger augmentations in Sec. 3.3.

3.1.3. **Ranking of General Superior Augmentations**

**Ranking of Best Augmentation.** The core of the easy-hard paradigm  $(I_l, T_l)$  is a more advanced augmentation. For this reason, we compare our method (Zhou et al. 2020) and SNNM (Huang et al. 2020) proposed keypoint-based augmentation Joint Cutout ( $T_{JC}$ ) and Joint Cut-Occlude ( $T_{CO}$ ), respectively. They also have similar joint consistency constraints. In this way, we can obtain a paired easy-hard augmentation  $(I_l, T_l) = (T_{JC}(I_l), T_{JC}(\mathbf{H}^i))$  for generating the labeled training set. The joint training of labeled and unlabeled training set can avoid propagation of teacher signals to avoid collapsing. Next, we answer two questions Q1 and Q2 by extensive empirical studies in Sec. 3.1 and Sec. 3.3, respectively. After that, we provide a theoretical perspective for understanding the power of designing stronger augmentations in Sec. 3.3.

**Synergy between Augmentations.** Then, instead of laboriously designing stronger augmentations, we consider to co-design more augmentations in sequence to obtain superior performances. For example, we can first use the joint consistency constraint of the Augment function (Chen et al. 2018; Fan et al. 2019; Hanrahan et al. 2019; Zhou et al. 2020; Zhou et al. 2022). Instead of auto-searching, we can also design a few specific augmentations and then use them in sequence. For example, if we need  $T_{CO}$  or  $T_{JC}$  on one image, we can continue to perform some unlabeled augmentations such as  $T_{JC_M}$ ,  $T_{JC_O}$  and  $T_{MC}$  on random areas. As shown in Fig. 3 applying  $T_{JC_O}$  (a  $T_{JC}$  after  $T_{JO}$ ) or

4

52





2916  
2917  
2918  
2919  
2920  
2921  
2922  
2923  
2924  
2925  
2926  
2927  
2928  
2929  
2930  
2931  
2932  
2933  
2934  
2935  
2936  
2937  
2938  
2939  
2940  
2941  
2942  
2943  
2944  
2945  
2946  
2947  
2948  
2949  
2950  
2951  
2952  
2953  
2954  
2955  
2956  
2957  
2958  
2959  
2960  
2961  
2962  
2963  
2964  
2965  
2966  
2967  
2968  
2969

### Example2: 875





3024  
 3025  
 3026  
 3027  
 3028  
 3029  
 3030  
 3031  
 3032  
 3033  
 3034  
 3035  
 3036  
 3037  
 3038  
 3039  
 3040  
 3041  
 3042  
 3043  
 3044  
 3045  
 3046  
 3047  
 3048  
 3049  
 3050  
 3051  
 3052  
 3053  
 3054  
 3055  
 3056  
 3057  
 3058  
 3059  
 3060  
 3061  
 3062  
 3063  
 3064  
 3065  
 3066  
 3067  
 3068  
 3069  
 3070  
 3071  
 3072  
 3073  
 3074  
 3075  
 3076  
 3077

Example2: 293



3078 **F HUMAN-MACHINE CONSISTENCY EVALUATION**  
 3079

3080 As described in Section 4.1, we employ GPT-4.1 to extract detailed information (e.g. evidence sets,  
 3081 reasoning chains) from the responses generated by the models under evaluation . Subsequently,  
 3082 based on the formulas presented in Section 4.1, we calculate  $S_{\text{location}}$  and  $S_{\text{reasoning}}$ , which are then  
 3083 used to derive  $S_{\text{total}}$  for each model’s response to the given question.

3084 To evaluate whether GPT-4.1 accurately extracts detailed information from the model responses,  
 3085 we conduct a human-Machine consistency evaluation. We first randomly sampled 200 questions  
 3086 from the dataset. Then, we invited human experts to analyze the corresponding model-generated  
 3087 responses for these questions and to manually extract key information, including evidence sets,  
 3088 reasoning chains, and the number of unrelated errors.

	$S_{\text{total}}$	$S_{\text{location}}$	$S_{\text{reasoning}}$	$P_{\text{unrelated,err}}$
Spearman’s correlation coefficients	0.841	0.806	0.842	0.954

3093 Table 4: Spearman’s correlation coefficients for:  $S_{\text{total}}$ ,  $S_{\text{location}}$ ,  $S_{\text{reasoning}}$ , and  $P_{\text{unrelated,err}}$ .  
 3094

3095 Using the information extracted by the human experts, we perform the following calculations:  
 3096

- 3097 (1) The  $\vec{S}_{\text{location}}$  vector for the 200 questions is calculated based on the evidence sets and Equa-  
 3098 tion 3.
- 3099 (2) The  $\vec{S}_{\text{reasoning}}$  vector is computed from the reasoning chains and Equation 4.
- 3100 (3) The  $\vec{P}_{\text{unrelated,err}}$  vector is obtained from the count of unrelated errors.
- 3101 (4) The  $\vec{S}_{\text{total}}$  vector is calculated for the 200 questions using Equation 6.

3103 Subsequently, these human-derived vectors ( $\vec{S}_{\text{location}}$ ,  $\vec{S}_{\text{reasoning}}$ ,  $\vec{P}_{\text{unrelated,err}}$ , and  $\vec{S}_{\text{total}}$ ) are compared  
 3104 against their counterparts generated by GPT-4.1. Spearman’s correlation coefficient is then calcu-  
 3105 lated for these four metrics. The results are presented in Table 4.

3107 Among the four Spearman correlation coefficients, the metric  $P_{\text{unrelated,err}}$  exhibits the highest corre-  
 3108 lation. This indicates that GPT-4.1’s extraction of unrelated errors closely aligns with that of human  
 3109 experts, making it the most precise among the three types of extracted information(*i.e.* evidence  
 3110 sets, reasoning chains, and unrelated errors).

3111 Although the correlation coefficients for the *evidence location score* and *reasoning process score*  
 3112 are relatively lower than  $P_{\text{unrelated,err}}$ , they still fall within the range of strong positive correlation.  
 3113 This demonstrates a high degree of consistency in the numerical trends of the scores calculated from  
 3114 GPT-4.1 and human expert extractions, respectively, proving that GPT-4.1 is capable of extracting  
 3115 the majority of effective evidence sets and reasoning chains.

3116 The correlation for the *total score* also lies within the strong positive range and slightly surpasses  
 3117 the correlations for the evidence location score. This also reflects a high level of agreement between  
 3118 GPT-4.1 and human experts.

3119 In summary, GPT-4.1 can extract relevant evidence and reasoning steps with considerable accuracy,  
 3120 leading to precise evaluation scores. This validates the effectiveness of our methodology, which uses  
 3121 GPT-4.1 to parse the responses of the models under evaluation.

3122  
 3123  
 3124  
 3125  
 3126  
 3127  
 3128  
 3129  
 3130  
 3131