

SAGASCALE: A REALISTIC, SCALABLE, AND HIGH-QUALITY LONG-CONTEXT BENCHMARK BUILT FROM FULL-LENGTH NOVELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Large Language Models (LLMs) have shown significant progress, but understanding long and complex documents remains challenging. Many long-context benchmarks have been proposed, but they face several limitations, including task realism, data scalability, and data quality. We introduce **SagaScale, a realistic, scalable, and high-quality long-context benchmark built from full-length novels**. The entire benchmark is constructed using an automated data collection pipeline that utilizes **external resources** (e.g., Wikipedia pages) to curate question-answer pairs. Critically, these external resources are provided only for benchmark construction and not during evaluation, which allows LLMs to curate complex questions that go beyond what they can answer during evaluation. SagaScale is also bilingual and offers the largest context length to date, with average token counts exceeding 250K for English novels and 320K for Chinese novels. Our evaluation across 12 frontier LLMs and three long-context methods — Naïve RAG, Agentic RAG, and Long Context — yields key insights, including: (1) Directly supplying the full context to the LLM can outperform other methods by a large margin; (2) Most LLMs still struggle with lengthy contexts, but Gemini-2.5-Pro stands out as an exception; and (3) Agentic RAG effectively addresses the retrieval bottleneck in Naïve RAG. Finally, we publicly release the SagaScale benchmark and our data collection codebase to facilitate future research.

1 INTRODUCTION

Large Language Models (LLMs) have achieved remarkable success across a variety of natural language processing (NLP) tasks (OpenAI et al., 2024; Grattafiori et al., 2024). A key remaining challenge, however, is long context understanding: the ability to process, reason over, and synthesize information from lengthy and complex documents. To guide progress in this area, the community has increasingly focused on creating dedicated long-context benchmarks (Zhang et al., 2024; Hsieh et al., 2024; Yuan et al., 2024; Bai et al., 2025; Yen et al., 2025).

However, existing long-context benchmarks have several limitations. For ease of construction and evaluation, many benchmarks rely on synthetic tasks (Kamradt, 2024; Hsieh et al., 2024). Such synthetic tasks, while useful for initial evaluation, often fail to capture the complexity of real-world scenarios. To improve realism and maintain quality, other benchmarks (e.g., Bai et al., 2025) rely on heavy human annotation. While this approach yields high-quality data, it is often prohibitively costly and time-consuming, limiting its scalability for adaptation to new domains or training set curation. To automate data annotation, recent efforts like LaRA (Li et al., 2025) attempt to generate QA pairs from isolated document chunks. However, this leads to simpler, locally focused questions and potential factual errors. Additionally, LaRA relies on manually tuned seed examples and prompts for each task type, which also limits its scalability and question diversity.¹

Therefore, existing long-context benchmarks (Kamradt, 2024; Kuratov et al., 2024; Wang et al., 2025; Bai et al., 2025; Li et al., 2025) face a fundamental conflict among task realism, data scalability, and data quality. To this end, we introduce **SagaScale, a realistic, scalable, and high-quality long-context benchmark built from full-length novels**.

¹For example, over 75% of questions in LaRA’s 128K book reasoning task start with "why".

Benchmark	Realistic Task	Data Scalability	Multilingual	Context Length	Evaluation Task
RULER (Hsieh et al., 2024)	✗	✓	✗	—	Synthetic
Counting-Stars (Song et al., 2025)	✗	✓	✓	—	Synthetic
Loogle (Li et al., 2024a)	✓	✗	✓	<32K	Multitask
Longbench-v2 (Bai et al., 2025)	✓	✗	✗	<128K	Multitask
LaRA (Li et al., 2025)	✓	⦿	✗	<128K	Single-doc QA
NovelQA (Wang et al., 2025)	✓	✗	✗	<200K	Fiction QA
SagaScale (Ours)	✓	✓	✓	228K	Fiction QA

Table 1: Comparison of SagaScale with existing benchmarks. Context lengths represent median token counts. The half-filled circle (⦿) denotes limited data scalability, where humans must manually tune seed examples and prompts for each task type even when using LLMs for direct data generation (Li et al., 2025).

To achieve scalability without sacrificing quality, we introduce an automated, novel data collection pipeline that leverages LLMs to generate and filter QA pairs. In the generation phase, we provide the LLM with both the novel’s text and **external resources** (e.g., Wikipedia articles) for broader context. Crucially, such asymmetry allows the LLM to generate questions that are much more complex than those it can answer during evaluation, where access is restricted to the novel’s text alone. After generation, we employ a multi-stage filtering process that discards factually incorrect, unrealistic, or contaminated QA pairs (Balloccu et al., 2024), where external knowledge sources (e.g., Google Search) still play key roles. By systematically validating our benchmark against six quality criteria, we demonstrate that our data collection pipeline consistently produces high-quality QA pairs.

To the best of our knowledge, SagaScale is the first long-context benchmark that combines high data quality with full realism and scalability. SagaScale is also bilingual and offers one of the largest context lengths to date, with average token counts exceeding 250K for English novels and 320K for Chinese novels² (Figure 5).

Finally, we assess three representative methods on SagaScale, each embodying a distinct strategy for handling long contexts: (1) Long Context, where a long-context language model directly processes the entire document and answers the question in a single pass (Liu et al., 2025); (2) Naïve Retrieval-Augmented Generation (RAG), which retrieves a fixed set of text chunks with the highest embedding similarity to the query, then generating an answer based on these chunks (Zhao et al., 2022; Gao et al., 2024); and (3) Agentic RAG, an extension of Naïve RAG that empowers an agent to iteratively refine retrieval across multiple rounds before producing the final answer (Singh et al., 2025; Liang et al., 2025; Jin et al., 2025). We summarize our core contributions as follows:

- **Dataset:** We present a novel, automated data collection pipeline that leverages external resources to generate QA pairs, providing a foundation for at-scale evaluation and training set construction.
- **Benchmark:** We introduce SagaScale, a realistic, scalable, and high-quality bilingual long-context benchmark featuring one of the largest context lengths to date.
- **Analysis:** We evaluate three representative long-context approaches using a wide range of LLMs, and provide an in-depth analysis of their performance, yielding key insights.

2 RELATED WORK

Long-context Approaches. Methods for addressing long contexts broadly fall into two categories: (1) direct long-context processing, where a long-context language model (LCLM) processes the entire input; and (2) workflow-based approaches, which employ external workflows to manage long contexts (Liu et al., 2025). In the first category, significant progress has been recently made in

²All token counts in this paper are calculated using the tokenizer from <https://huggingface.co/deepseek-ai/DeepSeek-R1-0528>, unless stated otherwise.

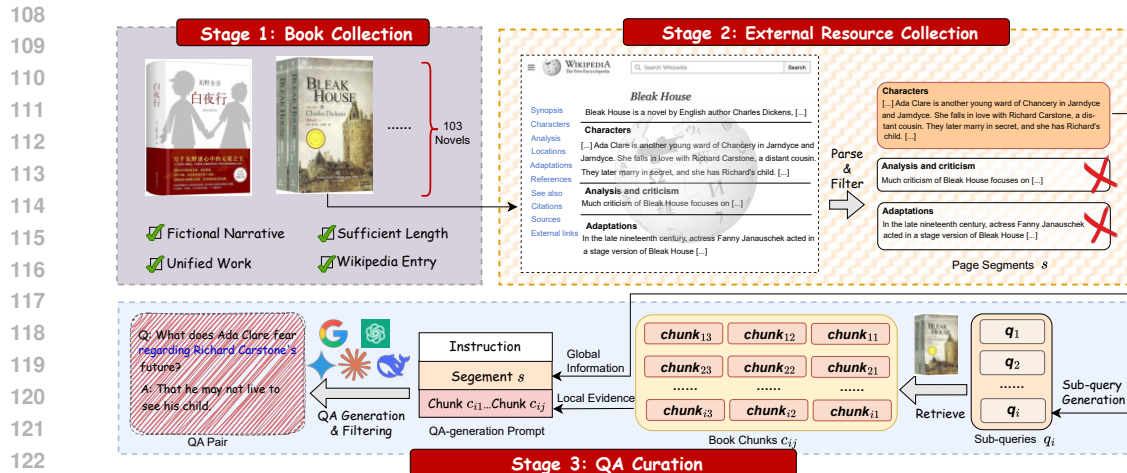


Figure 1: Data collection pipeline of SagaScale.

extending the context length of LLMs, with models now supporting context lengths up to 128K tokens (OpenAI et al., 2024a; DeepSeek-AI et al., 2025b), 200K tokens (Anthropic, 2025; AI et al., 2025), and even 1M tokens (Comanici et al., 2025; GLM et al., 2024). In the second category, retrieval-augmented generation (RAG) has been widely adopted (Gao et al., 2024; Fan et al., 2024). Based on the number of retrieval rounds in the workflow, RAG methods can be further divided into two groups: (1) single-round RAG (Ma et al., 2023; Wang et al., 2023; Shi et al., 2024), where a fixed set of text chunks is retrieved and used to generate the answer; and (2) iterative RAG (Singh et al., 2025; Jiang et al., 2025; Jin et al., 2025), where the model perform multiple rounds of retrieval to refine its retrieved data before generating the final answer. In this paper, the three evaluated methods correspond to: (1) direct long-context processing (Long Context); (2) single-round RAG (Naïve RAG); and (3) iterative RAG (Agentic RAG).

Long-context Benchmarks. Long-context benchmarks can be divided into two categories: (1) synthetic benchmarks, which are artificially constructed to isolate and measure specific long-context capabilities; and (2) real-world benchmarks, which are derived from naturally occurring data and reflect practical applications (Liu et al., 2025). While synthetic benchmarks are widely used for initial evaluation (Tay et al., 2021; Kamradt, 2024; Levy et al., 2024; Hsieh et al., 2024; Kuratov et al., 2024; Song et al., 2025; Ling et al., 2025), they often fail to predict downstream performance (Yen et al., 2025). To bridge this gap, many real-world long-context benchmarks have been developed (An et al., 2024; Wang et al., 2025; Bai et al., 2025), but they still face several limitations. First, the threshold for "long context" has increased rapidly due to advancements in LLMs, making some existing benchmarks inadequate for evaluating current models (Kočíský et al., 2018; An et al., 2024; Bai et al., 2024). Second, real-world benchmarks typically rely on heavy human annotation (Karpinska et al., 2024; Wang et al., 2025; Zhang et al., 2024; Bai et al., 2025), which is unscalable due to prohibitively high costs and time consumption. While recent efforts like LaRA (Li et al., 2025) have attempted to automate data annotation, their scalability and data quality (e.g., diversity) remain limited. We summarize these limitations and highlight our benchmark’s strengths in Table 1.

3 DATA

3.1 DATA COLLECTION

Figure 1 provides an overview of our data collection pipeline for SagaScale, which consists of three stages: Book Collection, External Resource Collection, and QA Curation.

Stage 1: Book Collection. We collect 103 full-length novels, comprising 77 in English and 26 in Chinese. The collection includes public domain books primarily sourced from Project Gutenberg³

³<https://www.gutenberg.org/>

and copyrighted books obtained from the web. Each book is selected based on four criteria: (1) the book must be a fictional narrative to reduce contamination risk, (2) it must be sufficiently long, typically exceeding 100K tokens, (3) it must be a single, unified work rather than an anthology of loosely related stories, and (4) it must have an informative Wikipedia (or Baidu Baike) entry. Together, the collected books feature complex storylines and substantial lengths (Table 9), enhancing the benchmark’s overall difficulty and realism.

Stage 2: External Resource Collection. Existing long-context benchmarks are built primarily from the source documents themselves, while neglecting valuable *external resources* such as community-written summaries. These resources distill key information, provide deep insights, and broaden the overall context—all while being significantly shorter than the original documents. These characteristics make them ideal candidates for crafting high-quality QA pairs.

To collect external resources, we first pair each novel with its Wikipedia (or Baidu Baike) page. Then, we parse each page to extract self-contained segments (e.g., the sections titled "Synopsis"). To ensure questions derived from these segments are answerable using only the book text, we employ an LLM to discard any segment that contains extrinsic information (e.g., adaptations) or subjective interpretations (e.g., criticism).

Stage 3: QA Curation. To obtain high-quality QA pairs, we decompose QA curation into two sub-stages: generation and filtering. In generation, we design a specialized RAG pipeline to ground external resources to the book text, and then use rigorous prompts to generate QA pairs. After generation, we apply three distinct filters to enforce correctness, realism, and non-contamination.

Stage 3.1: QA Generation. For each page segment s obtained from Stage 2, we start with retrieving relevant fixed-length chunks from the book. However, due to the lengthy and complex nature of the page segments, the original segment s cannot serve directly as a query. To address this, we adopt a multi-query generation approach (Kostric & Balog, 2024; Li et al., 2024b). First, an LLM generates concise sub-queries $q = \{q_1, q_2, \dots, q_n\}$ from s , each targeting a distinct aspect of s . Then, we use each q_i to retrieve the three most relevant book chunks $\{\text{chunks}_{i1}, \text{chunks}_{i2}, \text{chunks}_{i3}\}$. Importantly, this approach enhances retrieval efficiency and diversity, while also adaptively scaling the total number of retrieved chunks based on the complexity of the page segment.

After the retrieval step, we concatenate each chunk list $\{\text{chunks}_{i1}, \text{chunks}_{i2}, \text{chunks}_{i3}\}$ with the original page segment s . This forms a comprehensive context that includes two layers of granularity: the broader, global information from the page segment (s) and the specific, local evidence from the book chunks (chunks_{ij}). We pass this context to DeepSeek-R1 (DeepSeek-AI et al., 2025a) to generate QA pairs (refer to Appendix A.4.1 for prompts).

Stage 3.2: QA Filtering. Our QA filtering process consists of three steps: Correctness Verification, Realism Assurance, and Contamination Filtering.

Correctness Verification. Since QAs are generated by LLMs using only page segments and book chunks, errors can arise from model mistakes, flawed segments, or insufficient context. To address this, we use GPT-4o Search Preview (OpenAI, 2025) with web search capability to verify each QA pair. Given each QA pair and the name of its corresponding novel, the LLM searches online, attempts to find relevant information, and returns a binary judgment (correct/incorrect) or states no relevant information is found. Adopting a conservative approach, we retain only QA pairs with confirmed correct answers.

Importantly, this filtering step does not make the benchmark inherently easier. Verifying a QA pair is fundamentally simpler than solving one, since the model is given the correct answer for validation. The model can also benefit from unrestricted web access, a knowledge source far exceeding that of the source books and Wikipedia.

Realism Assurance. During the QA generation sub-stage, the model is unaware of the real-world popularity of the questions it generates, which can result in unnatural or rare questions. Consequently, an additional filtering step is required to ensure the realism of our benchmark.

To estimate a question’s popularity, we use an LLM to extract its keywords and then perform Google search with these keywords. The popularity score is defined as the total number of search results, which serves as an effective proxy for real-world user interests. Then, to account for differences in online presence across novels, we apply filtering on a per-book basis: for each novel, we keep




Figure 2: Distribution of questions across question types.

228
229
230
231
232
233
234

Range	<128K	128K–256K	256K–512K	512K–1M	Total
QA Pairs	165	386	252	321	1124
Novels	20	41	21	21	103

Table 2: Distribution of QA pairs and novels across different token count ranges.

235
236
237
238
239

only questions with a score at least 0.01% of the highest-scoring question for that book. Refer to Appendix A.2.1 for a case study of this step.

240
241
242
243

Contamination Filtering. The novels, along with their related online resources, are often seen by LLMs during training (Zhu et al., 2015; Rae et al., 2020; Gao et al., 2020). This poses a risk of data contamination, where models may answer questions by solely recalling memorized knowledge rather than demonstrating true long-context understanding.

244
245
246
247
248
249

To mitigate this issue, a straightforward filtering step is implemented. We select four state-of-the-art LLMs with extensive world knowledge: GPT-4.1 (OpenAI, 2025), DeepSeek-R1 (DeepSeek-AI et al., 2025a), Claude-Opus-4 (Anthropic, 2025), and Gemini-2.5-Pro (Comanici et al., 2025). Then, we prompt each LLM to answer all questions without access to the books. If any of these models can answer a question correctly, we remove that question from the benchmark.

250 3.2 DATA VERIFICATION

251
252
253
254
255
256
257

SagaScale is built using the data collection pipeline described above (see A.2.2 for examples). To ensure benchmark quality, we validate it against six key criteria: Diversity (questions cover a wide range of question types), Non-contamination (models struggle to answer correctly using only parametric knowledge), Difficulty (questions are challenging), Answerability (questions are solvable using only the book text), Objectivity (Questions are objective and have verifiable answers), and Correctness (no factual errors in questions or answers).

- 258
259
260
261
262
263
264
265
266
267
268
269
- **Diversity:** Following Kočiský et al. (2018), we classify questions into types based on narrative theory. As shown in Figure 2, the questions are well-distributed across these types, indicating a balanced and varied set of questions. For details on question types and the full classification procedure, refer to Appendix A.3.
 - **Non-contamination:** We evaluate LLMs in a closed-book setting (i.e., without access to the books) to test for data contamination, and find that the highest accuracy achieved is only 11.7% (Table 4). This demonstrates the effectiveness of our contamination filtering step and confirms that high performance requires genuine comprehension rather than pure memorization (Balloccu et al., 2024).
 - **Difficulty:** We confirm the benchmark’s challenging nature in Section 4 through experiments. While state-of-the-art models demonstrate strong capabilities, their performance remains well below perfect accuracy, leaving a substantial margin for improvement.

- **Answerability & Objectivity & Correctness:** We randomly sample 100 instances and ask our authors to manually verify them against the remaining criteria. All sampled questions are found to be answerable and objective, and 95 out of 100 instances are fully correct.

3.3 STATISTICS

QA Curation Statistics. Through our rigorous data collection pipeline, we ultimately collect 1,124 high-quality QA pairs, containing 613 for English novels and 511 for Chinese novels. Initially, the QA generation step produces 11,584 QA pairs (7,760 in English and 3,824 in Chinese), most of which are filtered out during contamination filtering.

Token Count Distribution. We categorize all QA pairs and novels into four length ranges based on book token counts: less than 128K, 128K–256K, 256K–512K, and 512K–1M. As shown in Table 2, both QA pairs and novels are distributed relatively evenly across these ranges. For detailed statistics for each novel, refer to Table 9.

4 EVALUATION

4.1 SETUP

Methods. We assess three representative methods on SagaScale: Long Context, Naïve Retrieval-Augmented Generation (RAG), and Agentic RAG. Below, we provide a detailed description of each method. Refer to Appendix A.4.2 for the specific prompts.

- **Long Context (LC).** For each question q_i , we create the model input by concatenating the instruction, the corresponding book text t_i , and the question q_i . When the input fits within the context window, the model generates an answer. Otherwise, we directly record a failure without any truncation attempts.
- **Naïve RAG (NR).** We first index each book by segmenting t_i into 512-token chunks and embedding them with mE5-large (Wang et al., 2024). For each question q_i , we retrieve the top 3 chunks with the highest embedding similarity to q_i . Using these chunks as the context, we prompt each model to either generate an answer or flag the question as unanswerable.
- **Agentic RAG (AR).** This method builds upon Naïve RAG by employing an agentic workflow (Singh et al., 2025; Liang et al., 2025; Jin et al., 2025). Starting from the question q_i , the model repeatedly issues search queries to retrieve text chunks from the book. Once it deems the gathered information sufficient, a final answer is generated. We cap the process at a maximum of 8 retrieval requests. For a fair comparison, we use the same retrieval setup as in Naïve RAG.

Models. We evaluate 12 powerful LLMs, including open-source models like Qwen3-235B-A22B (Yang et al., 2025), DeepSeek-V3 (DeepSeek-AI et al., 2025b), DeepSeek-R1 (DeepSeek-AI et al., 2025a), and proprietary models such as GPT-4o (OpenAI et al., 2024a), GPT-4.1 (OpenAI, 2025), o1 (OpenAI et al., 2024b), o4-mini (OpenAI, 2025), Claude-4 (Anthropic, 2025), and Gemini-2.5 (Comanici et al., 2025). All open-source models are deployed natively without context window extension (e.g., no YaRN (Peng et al., 2024)), while the proprietary models are accessed via their respective APIs. We set the default decoding temperature to 0.0 for all models.

LLM-as-a-Judge. Following Yen et al. (2025), we employ GPT-4o (OpenAI et al., 2024a) to evaluate model-generated answers. To confirm its reliability, we manually annotate 100 samples produced during evaluation⁴ and compare our annotations with GPT-4o’s judgements. We find a Cohen’s Kappa of 0.92, indicating strong agreement.

4.2 RESULTS & ANALYSIS

Table 3 presents the main evaluation results. Below, we analyze the results from multiple perspectives, with particular focus on Long Context (LC) and Agentic RAG (AR).

⁴All QA pairs sampled for LLM-as-a-Judge verification are confirmed to be completely correct.

Model	Max Len	Naïve RAG	Agentic RAG	Long Context
Qwen3-235B-A22B [†]	32K	44.8	50.4	0.0 [§]
Qwen3-235B-A22B [‡]	32K	46.0	46.3	0.0 [§]
DeepSeek-V3-0324 [†]	128K	49.4	52.5	9.3 [§]
DeepSeek-R1-0528 [‡]	128K	48.8	40.3	7.8 [§]
GPT-4o-2024-11-20 [†]	128K	45.8	53.3	7.6 [§]
o1-2024-12-17 [‡]	200K	<u>51.4</u>	61.1	14.5 [§]
o4-mini-2025-04-16 [‡]	200K	52.3	60.4	13.9 [§]
Claude-Sonnet-4-20250514 [†]	200K	47.6	61.9	10.8 [§]
Claude-Opus-4-20250514 [†]	200K	47.2	<u>63.5</u>	13.0 [§]
Claude-Sonnet-4-20250514 [‡]	200K	48.3	63.1	10.3 [§]
Claude-Opus-4-20250514 [‡]	200K	47.6	64.8	13.9 [§]
GPT-4.1-2025-04-14 [†]	1M	50.7	58.1	72.1
GPT-4.1-mini-2025-04-14 [†]	1M	48.8	50.1	49.6
Gemini-2.5-Flash-2025-06-17 [†]	1M	43.5	52.9	68.0
Gemini-2.5-Flash-2025-06-17 [‡]	1M	47.1	53.6	<u>72.5</u>
Gemini-2.5-Pro-2025-06-17 [‡]	1M	50.1	60.6	79.6
Average	—	48.1	55.8	27.7

[†] Non-reasoning models.

[‡] Reasoning models.

[§] Some problems exceed the model’s maximum context length. These processing failures are included in the reported score.

Table 3: Model performance (%) across three tasks: Naïve RAG (NR), Agentic RAG (AR), and Long Context (LC). The highest score for each task is highlighted in **bold**, and the second highest is underlined.

Overall comparison of methods. As shown in Table 3, Agentic RAG (AR) outperforms Naïve RAG (NR) across nearly all models, demonstrating the advantage of agentic retrieval over static, single-pass retrieval. When comparing Long Context (LC) with NR and AR, we observe no advantage for LC due to context window limits. However, when restricted to problems that fit within each model’s context window (Table 5), LC surpasses NR and AR by 10.9% and 4.4% on average, respectively. The superiority of LC is most pronounced for the top-performing 1M-context models (GPT-4.1 and Gemini-2.5 Series), where LC performance consistently surpasses even the highest AR score (64.8%). Therefore, we conclude that the simple approach of directly supplying the entire context as model input can easily outperform multi-step retrieval-based workflows, provided that the LLM can effectively process contexts of such length.

All methods show declining performance as context length increases, with AR exhibiting the smallest drop. Figure 3a plots the cumulative accuracy of each method at different book length thresholds. For all three methods, accuracy consistently declines as the context length increases. To further quantify this effect, we compute the gap in cumulative accuracy between problems up to 128K context and all those up to 1M context. The results reveal a 6.8% drop for NR, a 5.5% drop for AR, and a 6.3% drop for LC. Notably, AR exhibits the smallest relative decline, suggesting that iterative, agentic retrieval partially mitigates the impact of noise caused by extensive distracting content. Nevertheless, the magnitude of the decline across all methods remains considerable, highlighting the need for more robust LLMs and retrievers.

Most models fail to maintain performance as context length increases, but Gemini-2.5-Pro succeeds. We analyze LC performance for each 1M-context model, and plot their accuracy at various book length thresholds in Figure 3b. While a clear downward trend is observed for most models (GPT-4.1, GPT-4.1-mini, and Gemini-2.5-Flash), Gemini-2.5-Pro maintains stable accuracy across all lengths, standing out as an exception. This suggests that Gemini-2.5-Pro possesses unique capabilities for effectively processing extremely long contexts, which warrants further investigation.




Figure 3: Accuracy versus book length. **Left (a):** A comparison of accuracy decay for the three evaluated methods. **Right (b):** A breakdown of the Long Context performance for each model.

Some models fail to adhere to formatting instructions in long contexts. While LLMs are generally proficient at following simple formatting instructions (e.g., enclosing the answer in `\boxed{}`; see Appendix A.4.2) in short contexts (Ouyang et al., 2022), this ability can degrade drastically as context length increases. As shown in Table 6, a stark divergence is observed: top performers like Gemini-2.5 and GPT-4.1 maintain near-perfect formatting adherence, whereas models such as o1 and Claude-Sonnet-4 frequently fail. Although improved prompt engineering might mitigate this, the failure itself points to a critical weakness in certain models: formatting requires neither complex reasoning nor aggregation of information, but merely the recall of a simple instruction.

Agentic RAG addresses the retrieval bottleneck in Naïve RAG. In Naïve RAG (NR), model performance is largely limited by the quality of a fixed set of retrieved chunks. Poor retrieval prevents even strong LLMs from answering correctly, forming a *retrieval bottleneck*. This claim is supported by two observations: first, all models perform comparably in NR, with scores clustered between 43.5% and 52.3%; second, outcomes are highly polarized, with 66.7% of instances resulting in either 0% or 100% accuracy across all models (Table 8). Conversely, Agentic RAG (AR) exhibits no such clustering or polarization: AR scores vary widely from 40.3% to 64.8%, and only 25.4% of instances yield 0% or 100% accuracy. Furthermore, we find that 72.6% of the problems entirely unsolvable by NR are solved by at least one model using AR. These findings show that AR overcomes the retrieval bottleneck present in NR, demonstrating the potential of agentic retrieval to solve complex long-context problems.

Low calibration error⁵ is crucial for Agentic RAG performance, but unnecessary and even detrimental for Naïve RAG performance. To investigate the relationship between calibration error rate (measured from NR) and model accuracy, we compute Spearman’s rank correlation between two sets of variables:

- NR calibration error rates (Table 7), sorted in ascending order.
- Model accuracies for NR or AR (Table 3), sorted in descending order.

Interestingly, we find that better calibration in NR is *negatively* correlated with NR performance ($\rho_{NR}^{cal} = -0.5$) but *positively* correlated with AR performance ($\rho_{AR}^{cal} = 0.6$). While seemingly paradoxical, this finding aligns with our theoretical analysis (Appendix A.1). Intuitively, models prone to overconfidence (i.e., high calibration error rate) are penalized in AR because they tend to prematurely terminate the agentic retrieval process, whereas in NR, the same trait is beneficial because it can lead to occasional lucky guesses. Therefore, while low calibration error rate is unnecessary and even detrimental for the accuracy on common single-pass tasks (e.g., NR), it proves crucial for building strong search agents. We believe such distinction should be carefully considered when developing future LLMs.

⁵Following Wei et al. (2025), we define calibration error rate as the proportion of incorrect answers restricted to instances where the model attempts an answer. Questions where the model signals inability to solve are excluded.




Figure 4: Model accuracy as a function of tool-call chain length in Agentic RAG. All models show a decline in accuracy when more tool calls are invoked. For statistical validity, data points representing fewer than 10 samples are omitted.

Accuracy of Agentic RAG declines with longer tool-call chains. We evaluate the AR accuracy of each model on problems where exactly k tool calls are invoked, for $k \in \{1, 2, \dots, 8\}$. As shown in Figure 4, when the number of tool calls increases, a drop in accuracy is observed across all models, including both top-performing models like Claude-Opus-4 and weaker ones like DeepSeek-V3. We hypothesize that this decline stems from two factors: problems requiring more tool calls are inherently more difficult, and the larger volume of retrieved chunks increases the risk of the model being misled by irrelevant or confusing information. This finding underscores the need for more robust agents and advanced workflows to maintain reliability over extended tool-call chains.

5 CONCLUSION

This paper introduces SagaScale, a bilingual long-context benchmark built from full-length novels. It addresses the fundamental conflict among realism, scalability, and quality, which existing benchmarks fail to resolve. At its core is a novel, automated data collection pipeline that generates complex QA pairs with **external resources** and applies filtering to ensure correctness, realism, and non-contamination. Our evaluation, conducted on 3 representative approaches across 12 frontier LLMs, offers a wealth of insights into their respective strengths and weaknesses. By releasing SagaScale and its full data pipeline to the community, we aim to set a new standard for long-context evaluation and accelerate the development of next-generation AI systems.

6 LIMITATIONS & DISCUSSION

Limited Dataset Size. SagaScale consists of 1,124 QA pairs—a size sufficient for benchmarking but insufficient for training. This limitation stems from two main bottlenecks. First, our data annotation relies on manually collected novels and their Wikipedia (or Baidu Baike) links. While sophisticated web crawling could help, it is beyond the scope of this work. Second, our filtering process removes the vast majority of generated QA pairs to ensure non-contamination. While crucial for a fair comparison of frontier LLMs’ long-context capabilities, a single round of contamination filtering with one model may suffice for training data construction. Future work could investigate simplifying our filtering process for training purposes.

Limited Task Scope. SagaScale is confined to a single task: question answering on novels. While this is an ideal real-world task for evaluating long-context understanding, expanding to a wider variety of tasks could be beneficial. We note that our core methodology, which incorporates external knowledge sources to generate test data, is not limited to novels — It can also be adapted for other domains, such as understanding large code repositories or even long videos (e.g., movies). Future work could explore these directions.

Limited Language Coverage. SagaScale is currently bilingual, including only novels in English and Chinese. Given the rich availability of novels and encyclopedic knowledge across a wide range of languages, future work could extend SagaScale to include more languages.

486 REPRODUCIBILITY STATEMENT
487

488 To ensure reproducibility, we have publicly released all our code and data. The code repository offers
489 detailed instructions on collecting data and replicating experiments. The released dataset contains the
490 public-domain novels used in our benchmark, the generated question-answer pairs, and all evaluation
491 verdicts (e.g., correct, incorrect) for each model and task. We also provide the evaluation setup and
492 prompts in Section 4.1 and Appendix A.4.2, respectively.
493

494 REFERENCES
495

- 496 01. AI, :, Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Guoyin
497 Wang, Heng Li et al. Yi: Open foundation models by 01.ai, 2025. URL <https://arxiv.org/abs/2403.04652>.
498
- 499 Chenxin An, Shansan Gong, Ming Zhong, Xingjian Zhao, Mukai Li, Jun Zhang, Lingpeng Kong,
500 and Xipeng Qiu. L-eval: Instituting standardized evaluation for long context language models. In
501 Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting*
502 *of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14388–14411,
503 Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/
504 2024.acl-long.776. URL <https://aclanthology.org/2024.acl-long.776/>.
505
- 506 Anthropic. Introducing Claude 4, May 2025. URL <https://www.anthropic.com/news/claude-4>.
507
- 508 Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao
509 Liu, Aohan Zeng, Lei Hou et al. LongBench: A bilingual, multitask benchmark for long context
510 understanding. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the*
511 *62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*,
512 pp. 3119–3137, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi:
10.18653/v1/2024.acl-long.172. URL <https://aclanthology.org/2024.acl-long.172/>.
- 513 Yushi Bai, Shangqing Tu, Jiajie Zhang, Hao Peng, Xiaozhi Wang, Xin Lv, Shulin Cao, Jiazheng
514 Xu, Lei Hou, Yuxiao Dong et al. LongBench v2: Towards deeper understanding and reasoning
515 on realistic long-context multitasks. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and
516 Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for*
517 *Computational Linguistics (Volume 1: Long Papers)*, pp. 3639–3664, Vienna, Austria, July 2025.
518 Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.
519 acl-long.183. URL <https://aclanthology.org/2025.acl-long.183/>.
520
- 521 Simone Balloccu, Patrícia Schmidtová, Mateusz Lango, and Ondrej Dusek. Leak, cheat, repeat: Data
522 contamination and evaluation malpractices in closed-source LLMs. In Yvette Graham and Matthew
523 Purver (eds.), *Proceedings of the 18th Conference of the European Chapter of the Association*
524 *for Computational Linguistics (Volume 1: Long Papers)*, pp. 67–93, St. Julian’s, Malta, March
525 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.eacl-long.5. URL
<https://aclanthology.org/2024.eacl-long.5/>.
526
- 527 Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit
528 Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen et al. Gemini 2.5: Pushing the frontier
529 with advanced reasoning, multimodality, long context, and next generation agentic capabilities,
530 2025. URL <https://arxiv.org/abs/2507.06261>.
- 531 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,
532 Qihao Zhu, Shirong Ma, Peiyi Wang et al. DeepSeek-R1: Incentivizing reasoning capability in
533 LLMs via reinforcement learning, 2025a. URL <https://arxiv.org/abs/2501.12948>.
- 534 DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang
535 Zhao, Chengqi Deng, Chenyu Zhang et al. DeepSeek-V3 technical report, 2025b. URL <https://arxiv.org/abs/2412.19437>.
536
- 537 Wenqi Fan, Yujian Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and
538 Qing Li. A survey on rag meeting llms: Towards retrieval-augmented large language models, 2024.
539 URL <https://arxiv.org/abs/2405.06211>.

- 540 Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang,
541 Horace He, Anish Thite, Noa Nabeshima et al. The Pile: An 800gb dataset of diverse text for
542 language modeling, 2020. URL <https://arxiv.org/abs/2101.00027>.
- 543 Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng
544 Wang, and Haofen Wang. Retrieval-augmented generation for large language models: A survey,
545 2024. URL <https://arxiv.org/abs/2312.10997>.
- 546 Team GLM, :, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego
547 Rojas, Guanyu Feng et al. Chatglm: A family of large language models from glm-130b to glm-4
548 all tools, 2024. URL <https://arxiv.org/abs/2406.12793>.
- 549 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
550 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan et al. The Llama 3 herd of
551 models, 2024. URL <https://arxiv.org/abs/2407.21783>.
- 552 Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekesh, Fei Jia, and
553 Boris Ginsburg. RULER: What’s the real context size of your long-context language models?
554 In *First Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=kIoBbc76Sy>.
- 555 Zhouyu Jiang, Mengshu Sun, Lei Liang, and Zhiqiang Zhang. Retrieve, summarize, plan: Advancing
556 multi-hop question answering with an iterative approach, 2025. URL <https://arxiv.org/abs/2407.13101>.
- 557 Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan O Arik, Dong Wang, Hamed Zamani, and
558 Jiawei Han. Search-r1: Training LLMs to reason and leverage search engines with reinforcement
559 learning. In *Second Conference on Language Modeling*, 2025. URL [https://openreview.net/](https://openreview.net/forum?id=Rwhi91ideu)
560 [forum?id=Rwhi91ideu](https://openreview.net/forum?id=Rwhi91ideu).
- 561 Garrett Kamradt. Needle in a haystack - pressure testing LLMs. [https://github.com/gkamradt/](https://github.com/gkamradt/LLMTest_NeedleInAHaystack)
562 [LLMTest_NeedleInAHaystack](https://github.com/gkamradt/LLMTest_NeedleInAHaystack), 2024. GitHub repository. Accessed: 2025-08-16.
- 563 Marzena Karpinska, Katherine Thai, Kyle Lo, Tanya Goyal, and Mohit Iyyer. One thousand
564 and one pairs: A “novel” challenge for long-context language models. In Yaser Al-Onaizan,
565 Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical*
566 *Methods in Natural Language Processing*, pp. 17048–17085, Miami, Florida, USA, November
567 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.948. URL
568 <https://aclanthology.org/2024.emnlp-main.948/>.
- 569 Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis,
570 and Edward Grefenstette. The NarrativeQA reading comprehension challenge. *Transactions of the*
571 *Association for Computational Linguistics*, 6:317–328, 2018. doi: 10.1162/tacl_a_00023. URL
572 <https://aclanthology.org/Q18-1023/>.
- 573 Ivica Kostrić and Krisztian Balog. A surprisingly simple yet effective multi-query rewriting method
574 for conversational passage retrieval, 2024. URL <https://arxiv.org/abs/2406.18960>.
- 575 Yuri Kuratov, Aydar Bulatov, Petr Anokhin, Ivan Rodkin, Dmitry Igorevich Sorokin, Artyom
576 Sorokin, and Mikhail Burtsev. BABILong: Testing the limits of LLMs with long context
577 reasoning-in-a-haystack. In *The Thirty-eight Conference on Neural Information Processing*
578 *Systems Datasets and Benchmarks Track*, 2024. URL [https://proceedings.neurips.cc/](https://proceedings.neurips.cc/paper_files/paper/2024/file/c0d62e70dbc659cc9bd44cbcf1cb652f-Paper-Datasets_)
579 [paper_files/paper/2024/file/c0d62e70dbc659cc9bd44cbcf1cb652f-Paper-Datasets_](https://proceedings.neurips.cc/paper_files/paper/2024/file/c0d62e70dbc659cc9bd44cbcf1cb652f-Paper-Datasets_)
580 [and_Benchmarks_Track.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/c0d62e70dbc659cc9bd44cbcf1cb652f-Paper-Datasets_).
- 581 Mosh Levy, Alon Jacoby, and Yoav Goldberg. Same task, more tokens: the impact of input length
582 on the reasoning performance of large language models, 2024. URL <https://arxiv.org/abs/2402.14848>.
- 583 Jiaqi Li, Mengmeng Wang, Zilong Zheng, and Muhan Zhang. LooGLE: Can long-context
584 language models understand long contexts? In Lun-Wei Ku, Andre Martins, and Vivek
585 Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational*
586 *Linguistics (Volume 1: Long Papers)*, pp. 16304–16333, Bangkok, Thailand, August 2024a.

- 594 Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.859. URL <https://aclanthology.org/2024.acl-long.859/>.
595
596
- 597 Kuan Li, Liwen Zhang, Yong Jiang, Pengjun Xie, Fei Huang, Shuai Wang, and Minhao Cheng.
598 LaRA: Benchmarking Retrieval-Augmented Generation and long-context LLMs – no silver bullet
599 for LC or RAG routing. In *Forty-second International Conference on Machine Learning, 2025*.
600 URL <https://openreview.net/forum?id=CLF25dahgA>.
- 601 Zhicong Li, Jiahao Wang, Zhishu Jiang, Hangyu Mao, Zhongxia Chen, Jiazhen Du, Yuanxing Zhang,
602 Fuzheng Zhang, Di Zhang, and Yong Liu. DMQR-RAG: Diverse multi-query rewriting for RAG,
603 2024b. URL <https://arxiv.org/abs/2411.13154>.
604
- 605 Jintao Liang, Gang Su, Huifeng Lin, You Wu, Rui Zhao, and Ziyue Li. Reasoning rag via system 1
606 or system 2: A survey on reasoning agentic retrieval-augmented generation for industry challenges,
607 2025. URL <https://arxiv.org/abs/2506.10408>.
- 608 Zhan Ling, Kang Liu, Kai Yan, Yifan Yang, Weijian Lin, Ting-Han Fan, Lingfeng Shen, Zhengyin
609 Du, and Jiecao Chen. Longreason: A synthetic long-context reasoning benchmark via context
610 expansion, 2025. URL <https://arxiv.org/abs/2501.15089>.
611
- 612 Jiaheng Liu, Dawei Zhu, Zhiqi Bai, Yancheng He, Huanxuan Liao, Haoran Que, Zekun Wang,
613 Chenchen Zhang, Ge Zhang, Jiebin Zhang et al. A comprehensive survey on long context language
614 modeling, 2025. URL <https://arxiv.org/abs/2503.17407>.
- 615 Xinbei Ma, Yeyun Gong, Pengcheng He, Hai Zhao, and Nan Duan. Query rewriting in retrieval-
616 augmented large language models. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.),
617 *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*,
618 pp. 5303–5315, Singapore, December 2023. Association for Computational Linguistics. doi: 10.
619 18653/v1/2023.emnlp-main.322. URL <https://aclanthology.org/2023.emnlp-main.322/>.
620
- 621 OpenAI. Introducing GPT-4.1 in the api. Online news announcement, 2025. URL <https://openai.com/index/gpt-4-1/>. Launch of GPT-4.1, GPT-4.1 mini, and GPT-4.1 nano with expanded
622 long-context capability.
623
- 624 OpenAI. GPT-4o search preview, 2025. URL <https://platform.openai.com/docs/models/gpt-4o-search-preview>. Accessed: 2025-08-16.
625
626
- 627 OpenAI. OpenAI o3 and o4-mini system card. Online technical document, 2025. URL <https://cdn.openai.com/pdf/o3-and-o4-mini-system-card.pdf>. Technical report detailing the
628 capabilities, limitations, and safety evaluations of the o3 and o4-mini models.
629
- 630 OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni
631 Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman et al. GPT-4 technical report, 2024.
632 URL <https://arxiv.org/abs/2303.08774>.
633
- 634 OpenAI, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan
635 Clark, AJ Ostrow, Akila Welihinda, Alan Hayes et al. GPT-4o system card, 2024a. URL
636 <https://arxiv.org/abs/2410.21276>.
- 637 OpenAI, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden
638 Low, Alec Helyar, Aleksander Madry, Alex Beutel et al. OpenAI o1 system card, 2024b. URL
639 <https://arxiv.org/abs/2412.16720>.
- 640 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
641 Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray et al. Training language models to follow
642 instructions with human feedback. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and
643 Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems, 2022*. URL <https://openreview.net/forum?id=TG8KACxEON>.
644
645
- 646 Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. YaRN: Efficient context window
647 extension of large language models. In *The Twelfth International Conference on Learning
Representations, 2024*. URL <https://openreview.net/forum?id=wHBfxhZu1u>.

- 648 Jack W. Rae, Anna Potapenko, Siddhant M. Jayakumar, Chloe Hillier, and Timothy P. Lillicrap.
649 Compressive Transformers for long-range sequence modelling. In *International Conference on*
650 *Learning Representations*, 2020. URL <https://openreview.net/forum?id=SylKikSYDH>.
651
- 652 Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Richard James, Mike Lewis, Luke
653 Zettlemoyer, and Wen-tau Yih. REPLUG: Retrieval-augmented black-box language models. In
654 Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Conference*
655 *of the North American Chapter of the Association for Computational Linguistics: Human*
656 *Language Technologies (Volume 1: Long Papers)*, pp. 8371–8384, Mexico City, Mexico, June
657 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.463. URL
658 <https://aclanthology.org/2024.naacl-long.463/>.
- 659 Aditi Singh, Abul Ehtesham, Saket Kumar, and Tala Talaei Khoei. Agentic Retrieval-Augmented
660 Generation: A survey on agentic RAG, 2025. URL <https://arxiv.org/abs/2501.09136>.
- 661 Mingyang Song, Mao Zheng, and Xuan Luo. Counting-stars: A multi-evidence, position-aware,
662 and scalable benchmark for evaluating long-context large language models. In Owen Rambow,
663 Leo Wanner, Marianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, and Steven Schockaert
664 (eds.), *Proceedings of the 31st International Conference on Computational Linguistics*, pp. 3753–
665 3763, Abu Dhabi, UAE, January 2025. Association for Computational Linguistics. URL <https://aclanthology.org/2025.coling-main.253/>.
666
- 667 Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao,
668 Liu Yang, Sebastian Ruder, and Donald Metzler. Long range arena : A benchmark for efficient
669 transformers. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=qVyeW-grC2k>.
670
- 671 Cunxiang Wang, Ruoxi Ning, Boqi Pan, Tonghui Wu, Qipeng Guo, Cheng Deng, Guangsheng Bao,
672 Xiangkun Hu, Zheng Zhang, Qian Wang et al. NovelQA: Benchmarking question answering
673 on documents exceeding 200k tokens. In *The Thirteenth International Conference on Learning*
674 *Representations*, 2025. URL <https://openreview.net/forum?id=uMEsKEiB7J>.
675
- 676 Liang Wang, Nan Yang, and Furu Wei. Query2doc: Query expansion with large language models.
677 In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference*
678 *on Empirical Methods in Natural Language Processing*, pp. 9414–9423, Singapore, December
679 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.585. URL
680 <https://aclanthology.org/2023.emnlp-main.585/>.
681
- 682 Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei.
683 Multilingual E5 text embeddings: A technical report, 2024. URL <https://arxiv.org/abs/2402.05672>.
684
- 685 Jason Wei, Zhiqing Sun, Spencer Papay, Scott McKinney, Jeffrey Han, Isa Fulford, Hyung Won
686 Chung, Alex Tachard Passos, William Fedus, and Amelia Glaese. Browsecomp: A simple yet
687 challenging benchmark for browsing agents, 2025. URL <https://arxiv.org/abs/2504.12516>.
688
- 689 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
690 Gao, Chengen Huang, Chenxu Lv et al. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.
691
- 692 Howard Yen, Tianyu Gao, Minmin Hou, Ke Ding, Daniel Fleischer, Peter Izsak, Moshe Wasserblat,
693 and Danqi Chen. HELMET: How to evaluate long-context models effectively and thoroughly.
694 In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=293V3bJbmE>.
695
- 696 Tao Yuan, Xuefei Ning, Dong Zhou, Zhijie Yang, Shiyao Li, Minghui Zhuang, Zheyue Tan, Zhuyu
697 Yao, Dahua Lin, Boxun Li et al. LV-Eval: A balanced long-context benchmark with 5 length levels
698 up to 256k, 2024. URL <https://arxiv.org/abs/2402.05136>.
699
- 700 Xinrong Zhang, Yingfa Chen, Shengding Hu, Zihang Xu, Junhao Chen, Moo Hao, Xu Han, Zhen
701 Thai, Shuo Wang, Zhiyuan Liu et al. ∞ Bench: Extending long context evaluation beyond 100K
tokens. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd*

702 *Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp.
703 15262–15277, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi:
704 10.18653/v1/2024.acl-long.814. URL <https://aclanthology.org/2024.acl-long.814/>.
705

706 Wayne Xin Zhao, Jing Liu, Ruiyang Ren, and Ji-Rong Wen. Dense text retrieval based on pretrained
707 language models: A survey, 2022. URL <https://arxiv.org/abs/2211.14876>.

708 Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and
709 Sanja Fidler. Aligning books and movies: Towards storylike visual explanations by watching
710 movies and reading books. In *Proceedings of the IEEE International Conference on Computer*
711 *Vision*, pp. 19–27, 2015. URL [https://www.cv-foundation.org/openaccess/content_iccv_](https://www.cv-foundation.org/openaccess/content_iccv_2015/papers/Zhu_Aligning_Books_and_ICCV_2015_paper.pdf)
712 [2015/papers/Zhu_Aligning_Books_and_ICCV_2015_paper.pdf](https://www.cv-foundation.org/openaccess/content_iccv_2015/papers/Zhu_Aligning_Books_and_ICCV_2015_paper.pdf).
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




Figure 5: Violin plots illustrating the distribution of book token counts within the English (left) and Chinese (right) subsets of SagaScale.

A APPENDIX

A.1 THEORETICAL ANALYSIS

Below, we provide a theoretical analysis of how calibration error relates to the accuracy of Naïve RAG (NR) and Agentic RAG (AR).

Denote a^{NR} and e^{NR} as the probabilities of producing a correct or incorrect answer in NR, respectively. Note that $a^{\text{NR}} + e^{\text{NR}}$ may be less than 1, since the model can flag a question as unanswerable. Similarly, let a_i^{AR} and e_i^{AR} be the probabilities of producing correct or incorrect answer in AR after exactly i retrieval requests. Here, $a_i^{\text{AR}} + e_i^{\text{AR}}$ may also be less than 1, as the model can continue to issue more retrievals instead of answering.

Throughout this analysis, we assume that the outcomes of each retrieval step are independent; that is, the probability of any outcome at step i does not depend on the outcomes of previous steps. Therefore, the accuracy of NR is simply a^{NR} , and the accuracy of AR is given by

$$\sum_{i=0}^N \left(\prod_{j=0}^{i-1} (1 - a_j^{\text{AR}} - e_j^{\text{AR}}) \right) a_i^{\text{AR}}$$

where N denotes the maximum number of retrievals allowed.

If accuracy alone is considered, NR does not penalize overconfidence, as increasing e^{NR} alone does not affect the NR accuracy which always remains a^{NR} . In fact, overconfidence can even be implicitly rewarded in NR, as higher e^{NR} may result in increased a^{NR} due to lucky guesses. Conversely, AR is highly sensitive to overconfidence: large values of e_i^{AR} can substantially decrease AR accuracy, since the weighting coefficients for a_i^{AR} decrease exponentially as i increases.

To illustrate this, consider a simplified case where $a^{\text{NR}} = a$, $a_i^{\text{AR}} = a$ for all $i \geq 0$, and $e^{\text{NR}} = e$, $e_i^{\text{AR}} = e$ for all $i \geq 0$. Let $N \rightarrow \infty$, then the accuracy of NR is simply a , while the accuracy of AR becomes the sum of a geometric series:

$$\sum_{i=0}^{\infty} (1 - a - e)^i a = \frac{a}{a + e}$$

A.2 CASE STUDY

A.2.1 CASE STUDY OF REALISM ASSURANCE

To illustrate the effectiveness of our realism assurance step, we present several examples in Figure 6. The questions marked in red are representative of those discarded during the realism assurance step: they either delve into negligible details peripheral to the main narrative (e.g., "direction must travelers take"), or are structured more like an exam quiz than a natural user inquiry (e.g., "In which city ..., the same location where ..."). In contrast, the questions marked in green are more concise, natural, and focus on important characters (e.g., "Dr. Holmes", "Sara Macallan") or plot points (e.g., the defense's failure to prove intent). As a result, they achieve higher popularity scores, indicating a stronger alignment with genuine user interest.

Retained and filtered question examples in realism assurance

Q: What direction must travelers take to use the longer but easier route between Burgdale and Shadowy Vale? (The Roots of the Mountains)

Q: In which city was Gaston de Lancy buried, the same location where his mother resided and was falsely claimed to be ill? (The Trail of the Serpent)

Q: How many children does Dr. Holmes have? (Mrs. Dalloway)

Q: Why was the defense unable to prove Sara Macallan intended to obtain arsenic? (The Law and the Lady)

Figure 6: Examples of questions retained (in green) or filtered out (in red) during the realism assurance step. The retained questions are more likely to be asked by real users, while the filtered ones are often too rare or unnatural.

A.2.2 CASE STUDY OF SAGASCALE

We show representative QA examples from SagaScale in Figure 7, including one in English and one in Chinese.

QA examples in SagaScale

Q: After Toad confesses his identity and actions to the Engine Driver, what specific event occurs that leads the driver to allow Toad to escape?

A: passing through a tunnel
(Oliver Twist)

Q: 云天明赠送给程心的恒星DX3906在哪个纪元被发现拥有行星?

A: 威慑纪元61年
(三体系列)

Figure 7: Examples of questions retained (in green) or filtered out (in red) during the realism assurance step. The retained questions are more likely to be asked by real users, while the filtered ones are often too rare or unnatural.

864 A.3 QUESTION CLASSIFICATION

865 Given the large size of our benchmark, we employ GPT-4o to automatically classify questions. The
866 prompt below provides all category names, definitions, and illustrative examples to guide the model.
867
868

869 Prompt for Question Classification

870 ## Task

871 Given a question, classify it into one of the following categories.

872 Put the category name in a single `\boxed{}` block (e.g., `\boxed{Location}`).

873 ## Category Definitions

874 Location: Questions targeting places.

875 Character: Questions targeting roles in the narrative.

876 Reason: Questions targeting the reasons or motivations.

877 Method: Questions targeting how something was accomplished.

878 Event: Questions targeting specific occurrences or plot points.

879 Entity: Questions targeting specific entities (not characters or locations).

880 Quantity: Questions targeting amounts or numbers.

881 Temporal: Questions targeting when an event occurs or its duration.

882 Quotation: Questions targeting specific phrases or terms in the text.

883 Miscellaneous: Questions that do not fit into the categories above.

884 ## Examples

885 Location: Where did Oliver express his grief after being punished for defending his
886 mother?

887 Character: Who does Valeria privately suspect of intentionally poisoning Sara
888 Macallan with arsenic?

889 Reason: Why did the jury return a 'not proven' verdict instead of 'not guilty' for
890 Eustace Macallan?

891 Method: How did Abel Magwitch escape from the prison ship?

892 Event: What happened to the heiress George Vavasor was engaged to before her money
893 could benefit him?

894 Entity: What item of clothing hides Miserrimus Dexter's physical deformity during
895 Valeria's visit?

896 Quantity: How many children does John Boucher have?

897 Temporal: How long did Orlando serve as ambassador to Constantinople before being
898 elevated to a Duke?

899 Quotation: What phrase described Oliver's trial period as an apprentice with Mr.
900 Sowerberry?

901 Miscellaneous: What is the nickname of the child adopted by Joseph Peters?

902 ## Input

903 `{question}`
904

905
906 Figure 8: Prompt template used for question classification. All category names, definitions, and
907 examples are available in the prompt.
908

909 A.4 PROMPTS

910 A.4.1 DATA COLLECTION PROMPTS

911 Below, we present the prompt template used for each stage of our data collection process. The
912 QA generation step is particularly important, so we provide a detailed explanation of the key rules
913 specified in the prompt (Figure 11):
914
915

- 916 • *The question should focus only on fictional elements without involving any real-life*
917 *information.* This rule, combined with the segment filtering step in Stage 2, prevents

- 918 questions that require outside information, thereby enforcing answerability. Thus, the
 919 collected external resources are used only to provide book-based knowledge rather than all
 920 generic facts.
- 921 • *The question should be highly objective and specific, leading to a unique answer.* This
 922 ensures question objectivity and answer verifiability.
 - 923 • *The question must be multi-hop, requiring in-depth understanding and detailed knowledge.*
 924 *It should span multiple passages or different parts of the same passage to be answered*
 925 *correctly.* This promotes questions that require holistic understanding, detailed knowledge,
 926 and complex reasoning, thereby enhancing the difficulty of our benchmark.
 - 927 • *The answer must be a single number, entity or short phrase that directly addresses the*
 928 *question.* This rule reinforces objectivity and simplifies evaluation. It is also crucial
 929 for answer correctness, as it discourages lengthy or complex answers that may introduce
 930 inaccuracies.

Prompt for filtering article segments

```
934 ## Task
935 Given a Wikipedia text that contains information about the novel {novel_name},
936 determine if it is a summary (e.g., plot summary, character summary) of the novel.
937 Additional Requirements:
938 - The text must be highly objective. Excessive subjective contents, e.g., opinions,
939 speculation, interpretative analysis, subjective commentary, should fail the test.
940 - The text must focus on the novel's fictional elements. Excessive inclusion of
941 real-world topics, e.g., literary study, historical context, author biography,
942 publication history, cultural impact, awards, adaptations, academic analysis, should
943 fail the test.
944 Output \boxed{YES} or \boxed{NO} after careful examination.
945 ## Input
946 {article_segment}
947
```

948
 949 Figure 9: Prompt template used for filtering segments parsed from Wikipedia (or Baidu Baike)
 950 articles. The model is instructed to identify segments that are objective summaries of the novel’s
 951 fictional elements, ensuring relevance and answerability.

Prompt for multi-query generation

```
955 ## Task
956 Given a Wikipedia text for the novel {novel_name}. You should break down the text
957 into fine-grained, specific parts and write a concise summary for every part. Each
958 summary should contain one to two sentences.
959 Your output should be in {novel_language} and adhere to the following format for each
960 summary: "<summary>
961 your_summary
962 </summary>".
963 ## Input
964 {article_segment}
965
```

966
 967 Figure 10: Prompt template used for multi-query generation. The model is instructed to break down
 968 the lengthy article segment into concise summaries, each serving as a query to retrieve relevant book
 969 chunks for QA generation (Stage 3.1).

970
 971

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
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

```
Prompt for generating QA pairs

## Task
Given several passages for the novel {novel_name}. The first passage is from
Wikipedia, while the rest are novel snippets.

Generate a single question-answer pair. Requirements:
1. The question should focus only on fictional elements of the novel (e.g.,
characters, plots) without involving any real-life information (e.g., author,
reviews, writing styles, other novels).
2. The question should be highly objective and specific, leading to a unique answer.
However, any unnecessary hint or context should be excluded.
3. The question must be **multi-hop**, requiring in-depth understanding and detailed
knowledge. It should span multiple passages or different parts of the same passage to
be answered correctly.
4. The answer must be a single number, entity or short phrase that directly addresses
the question.
5. Both the question and the answer should be unambiguous when the passages are not
provided.
6. If some previously generated question-answer pairs are provided, the new pair
should be completely different from those pairs (by focusing on distinct characters,
plots, aspects, etc.).

If such pair is available, put the generated question and answer in two separate
\boxed{} blocks. Both of them should be in {novel_language}.

If no such pair is available, your output should be a single \boxed{END} block. Do
not think too long.

## Input
{passages}
```

Figure 11: Prompt template used for generating QA pairs. The model is instructed to create QA pairs that are answerable (i.e., solvable using only the book text), objective, difficult, and diverse.

```
Prompt for verifying QA pairs

## Task
Given a QA pair for the novel {novel_name}, you should check if it is correct.
Requirements:
- The **answer** must be accurate and must not omit any core information.
- The **question** must not contain any factual errors or misleading content.

Output your decision as \boxed{PASS} or \boxed{FAIL}.

If no relevant information is found on the web, output \boxed{FAIL}.

## Input
{question}
{answer}
```

Figure 12: Prompt template used for verifying the factual accuracy of QA pairs. The model is instructed to ensure the correctness of the QA pair through web search.

```
1026
1027
1028 Prompt for extracting keywords
1029
1030 ## Task
1031 Given a question for the novel {novel_name}, you should identify all keywords in the
1032 question. Requirements:
1033 - All unnecessary words that do not affect the core intent should be removed.
1034 - Each keyword should be kept in its shortest form.
1035
1036 Your output should be several \boxed{} blocks, each containing exactly one keyword.
1037
1038 ## Input
1039 {question}
```

Figure 13: Prompt template used for extracting keywords from a question. The extracted keywords are used to perform a Google search to estimate question popularity.

```
1041
1042 Prompt for Contamination Filtering
1043
1044 ## Task
1045 Given a question for the novel {novel_name}, you should answer the question based on
1046 your knowledge of the novel.
1047
1048 Your output should be in {novel_language} and adhere to the following format for the
1049 final answer: "<answer>
1050 {final_answer_text}
1051 </answer>".
1052
1053 ## Input
1054 {question}
```

Figure 14: Prompt template used for filtering contaminated QA pairs. We provide the model with the novel name and the question, and ask it to craft an answer based solely on its internal knowledge. Note that this prompt is also used in closed-book evaluation.

1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

1080 A.4.2 EVALUATION PROMPTS

1081 Below, we present the prompt template for LLM-as-a-Judge and each of the three evaluation settings.

1082 Prompt for LLM-as-a-Judge

```
1083
1084 ## Task
1085 Given a question for the novel {novel_name}, the corresponding ground truth answer,
1086 and a model-generated answer, you should judge whether the model-generated answer is
1087 correct.
1088
1089 The model-generated answer is correct if and only if:
1090 - It directly addresses the question.
1091 - No part of it contradicts the ground truth answer.
1092 - It captures the essence or key points in the ground truth answer (note that a
1093 lexical exact match is not required, but the semantic meaning should be preserved).
1094
1095 Output your decision as either \boxed{CORRECT} or \boxed{INCORRECT}.
1096
1097 ## Input
1098 Question: {question}
1099
1100 Ground truth answer: {ground_truth_answer}
1101
1102 Model-generated answer: {model_generated_answer}
```

1103 Figure 15: Prompt template used for judging the correctness of a model-generated answer against a
1104 ground truth answer.

1107 Prompt for Long Context Evaluation

```
1108
1109 ## Task
1110 Given a question for the novel {novel_name}, you should answer the question based
1111 solely on the full text of the novel.
1112
1113 Your output should be in {novel_language} and adhere to the following format for the
1114 final answer: "<answer>
1115 final_answer_text
1116 </answer>".
1117
1118 ## Book Text
1119 {book_text}
1120
1121 ## Question
1122 {question}
```

1123 Figure 16: Prompt template used for Long Context evaluation. We provide the model with the novel
1124 name, the full text of the book, and the question, and ask it to craft an answer based solely on the
1125 book text.

```

1134
1135
1136 Prompt for Naïve RAG Evaluation
1137
1138 ## Task
1139 Given a question for the novel {novel_name}, you should answer the question based
1140 solely on the provided contexts. Each context is a direct excerpt from the book, with
1141 its location indicated by the source index ('idx').
1142
1143 If the question is answerable based solely on the provided contexts, you should
1144 provide a final answer. Your output should be in {novel_language} and adhere to the
1145 following format for the final answer: "<answer>
1146 {final_answer_text}
1147 </answer>".
1148
1149 Otherwise, output "<answer>
1150 NONE
1151 </answer>".
1152
1153 ## Contexts
1154 {contexts}
1155
1156 ## Question
1157 {question}
1158

```

Figure 17: Prompt template used for Naïve RAG evaluation. We provide the model with the novel name, contexts retrieved from the book, and the question. The model is instructed to craft an answer based solely on the provided contexts, or respond with "NONE" if the information is insufficient.

```

1159
1160 Prompt for Agentic RAG Evaluation
1161
1162 ## Task
1163 Given a question for the novel {novel_name}, you should answer the question based
1164 solely on retrieving contexts from the novel. To do this, you will serve as an agent
1165 that iteratively request a search system to retrieve contexts, refine your query
1166 based on the retrieved contexts, and ultimately generate the final answer when you
1167 have sufficient information.
1168
1169 Specifically, every time a query is sent, you will receive a list of contexts, where
1170 each context is a direct excerpt from the book with its location indicated by the
1171 source index ('idx'). To return these contexts, the search system works by simply
1172 retrieving the top-3 book chunks with the highest cosine similarities to the query,
1173 so they might be noisy.
1174
1175 In each turn, you must provide either a single query or a single final answer. Your
1176 output should be in {novel_language} and adhere to the following formats:
1177 - For queries, use "<query>
1178 query_text
1179 </query>".
1180 - For the final answer, use "<answer>
1181 final_answer_text
1182 </answer>".
1183
1184 Now, begin with the first turn.
1185
1186 ## Question
1187 {question}

```

Figure 18: Prompt template used for Agentic RAG evaluation. We provide the model with the novel name and the question, and instruct it to iteratively retrieve contexts and refine its query until it can provide a final answer.

Model	Accuracy
Qwen3-235B-A22B [†]	6.5
Qwen3-235B-A22B [‡]	5.5
DeepSeek-V3-0324 [†]	8.2
DeepSeek-R1-0528 [‡]	7.3
GPT-4o-2024-11-20 [†]	5.2
o1-2024-12-17 [‡]	11.7
o4-mini-2025-04-16 [‡]	6.9
Claude-Sonnet-4-20250514 [†]	6.2
Claude-Opus-4-20250514 [†]	4.0
Claude-Sonnet-4-20250514 [‡]	6.4
Claude-Opus-4-20250514 [‡]	10.0
GPT-4.1-2025-04-14 [†]	4.4
GPT-4.1-mini-2025-04-14 [†]	3.8
Gemini-2.5-Flash-2025-06-17 [†]	5.8
Gemini-2.5-Flash-2025-06-17 [‡]	5.2
Gemini-2.5-Pro-2025-06-17 [‡]	10.4

[†] Non-reasoning models.

[‡] Reasoning models.

Table 4: Model performance (%) in the closed-book setting, where the model has no access to the books and must answer questions based on its internal knowledge.

Model	Max Len	NR	AR	LC
Qwen3-235B-A22B [†]	32K	—	—	—
Qwen3-235B-A22B [‡]	32K	—	—	—
DeepSeek-V3-0324 [†]	128K	55.8	53.3	63.6
DeepSeek-R1-0528 [‡]	128K	57.0	46.7	53.3
GPT-4o-2024-11-20 [†]	128K	48.8	53.5	65.9
o1-2024-12-17 [‡]	200K	55.6	64.4	51.7
o4-mini-2025-04-16 [‡]	200K	57.1	64.1	49.5
Claude-Sonnet-4-20250514 [†]	200K	53.8	63.1	53.8
Claude-Opus-4-20250514 [†]	200K	51.1	65.8	64.9
Claude-Sonnet-4-20250514 [‡]	200K	51.4	61.4	55.2
Claude-Opus-4-20250514 [‡]	200K	49.5	64.8	74.3
GPT-4.1-2025-04-14 [†]	1M	50.7	58.1	72.1
GPT-4.1-mini-2025-04-14 [†]	1M	48.8	50.1	49.6
Gemini-2.5-Flash-2025-06-17 [†]	1M	43.5	52.9	68.0
Gemini-2.5-Flash-2025-06-17 [‡]	1M	47.1	53.6	72.5
Gemini-2.5-Pro-2025-06-17 [‡]	1M	50.1	60.6	79.6
Average	—	51.5	58.0	62.4

[†] Non-reasoning models.

[‡] Reasoning models.

Table 5: Model performance (%) across three tasks: Naïve RAG (NR), Agentic RAG (AR), and Long Context (LC). Each model is evaluated on a unique subset of the benchmark which only includes problems that fit within its context window. *Consequently, scores are not directly comparable between different models.* The best-performing task for each model is highlighted in **bold**.

Model	Max Len	Err Rate↓
Qwen3-235B-A22B [†]	32K	—
Qwen3-235B-A22B [‡]	32K	—
DeepSeek-V3-0324 [†]	128K	0.6
DeepSeek-R1-0528 [‡]	128K	13.9
GPT-4o-2024-11-20 [†]	128K	2.3
o1-2024-12-17 [‡]	200K	31.1
o4-mini-2025-04-16 [‡]	200K	15.2
Claude-Sonnet-4-20250514 [†]	200K	26.7
Claude-Opus-4-20250514 [†]	200K	12.4
Claude-Sonnet-4-20250514 [‡]	200K	26.2
Claude-Opus-4-20250514 [‡]	200K	1.4
GPT-4.1-2025-04-14 [†]	1M	2.1
GPT-4.1-mini-2025-04-14 [†]	1M	2.2
Gemini-2.5-Flash-2025-06-17 [†]	1M	0.0
Gemini-2.5-Flash-2025-06-17 [‡]	1M	0.3
Gemini-2.5-Pro-2025-06-17 [‡]	1M	0.3

[†] Non-reasoning models.

[‡] Reasoning models.

Table 6: Parsing error rates in the Long Context (LC) task, where only problems that fit within each model’s context window are considered. *Lower* values indicate stronger format adherence, and the lowest value is highlighted in **bold**.

Model	NR↓	AR↓
Qwen3-235B-A22B [†]	16.3	49.3
Qwen3-235B-A22B [‡]	26.0	53.7
DeepSeek-V3-0324 [†]	19.0	47.5
DeepSeek-R1-0528 [‡]	22.0	59.7
GPT-4o-2024-11-20 [†]	17.6	46.7
o1-2024-12-17 [‡]	18.3	38.7
o4-mini-2025-04-16 [‡]	24.8	34.1
Claude-Sonnet-4-20250514 [†]	11.7	35.6
Claude-Opus-4-20250514 [†]	14.5	33.1
Claude-Sonnet-4-20250514 [‡]	15.7	33.9
Claude-Opus-4-20250514 [‡]	12.3	28.1
GPT-4.1-2025-04-14 [†]	24.3	41.8
GPT-4.1-mini-2025-04-14 [†]	39.7	49.8
Gemini-2.5-Flash-2025-06-17 [†]	12.0	46.9
Gemini-2.5-Flash-2025-06-17 [‡]	20.0	43.8
Gemini-2.5-Pro-2025-06-17 [‡]	17.8	38.9

[†] Non-reasoning models.

[‡] Reasoning models.

Table 7: Calibration error rates for Naïve RAG (NR) and Agentic RAG (AR) tasks. In NR, a problem is treated as unanswered if the model flags it as unanswerable (see Section 4.1), whereas in AR, it is considered unanswered if more than the maximum of 8 search queries are issued. The lowest calibration error rate for each task is highlighted in **bold**.

1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

Methods	Polarization
Naïve RAG (NR)	66.7
Agentic RAG (AR)	25.4
Long Context (LC)	34.6

Table 8: Proportion (%) of polarized outcomes (0% or 100% accuracy across all models) in different methods. Among the three methods, Naïve RAG (NR) exhibits the highest polarization, while Agentic RAG (AR) and Long Context (LC) show significantly reduced polarization.

Title	Language	Tokens	#QAs
<i>A Portrait of the Artist as a Young Man</i>	English	113,002	12
<i>A Tale of Two Cities</i>	English	187,847	4
<i>Across the Zodiac: The Story of a Wrecked Record</i>	English	216,635	1
<i>Adam Bede</i>	English	297,822	11
<i>Anna Karenina</i>	English	489,802	10
<i>Anne of Green Gables</i>	English	143,117	12
<i>Babbitt</i>	English	180,959	1
<i>Barnaby Rudge: A Tale of the Riots of Eighty</i>	English	352,812	13
<i>Bleak House</i>	English	491,790	9
<i>Brave New World</i>	English	87,214	3
<i>Brideshead Revisited</i>	English	145,520	19
<i>Can You Forgive Her?</i>	English	419,072	8
<i>Crime and Punishment</i>	English	288,969	3
<i>David Copperfield</i>	English	497,757	14
<i>Dombey and Son</i>	English	498,783	8
<i>Dracula</i>	English	213,967	7
<i>Emma</i>	English	215,375	2
<i>Gone with the Wind</i>	English	544,052	6
<i>Great Expectations</i>	English	253,785	6
<i>Jane Eyre: An Autobiography</i>	English	257,206	1
<i>Les Misérables</i>	English	808,272	9
<i>Little Dorrit</i>	English	466,180	15
<i>Main Street</i>	English	236,938	5
<i>Man and Wife</i>	English	314,721	8
<i>Mansfield Park</i>	English	210,685	6
<i>Marcella</i>	English	339,614	6
<i>Mary Barton: A Tale of Manchester Life</i>	English	219,279	9
<i>Middlemarch</i>	English	434,792	11
<i>Moby-Dick; or, The Whale</i>	English	309,020	3
<i>Mrs. Dalloway</i>	English	88,704	11
<i>Murder Must Advertise</i>	English	153,295	2
<i>New Grub Street</i>	English	251,977	10
<i>Nicholas Nickleby</i>	English	463,064	5
<i>North and South</i>	English	248,470	13
<i>Nostramo</i>	English	233,197	15
<i>Oliver Twist</i>	English	226,372	2
<i>Orlando: A Biography</i>	English	106,267	2
<i>Paul Clifford</i>	English	252,075	1
<i>Pride and Prejudice</i>	English	170,016	2
<i>Quo Vadis: A Narrative of the Time of Nero</i>	English	287,683	10
<i>Rebecca of Sunnybrook Farm</i>	English	100,226	6
<i>Robinson Crusoe</i>	English	142,413	8
<i>Shirley</i>	English	299,751	8
<i>Sons and Lovers</i>	English	228,329	2
<i>Tess of the d'Urbervilles</i>	English	206,879	6
<i>The Adventures of Huckleberry Finn</i>	English	158,133	6
<i>The Adventures of Tom Sawyer</i>	English	101,736	2
<i>The Angel of the Revolution: A Tale of the Coming Terror</i>	English	188,538	8
<i>The Gilded Age: A Tale of Today</i>	English	213,972	3

	Title	Language	Tokens	#QAs
1350				
1351				
1352	<i>The History of Tom Jones, a Foundling</i>	English	469,959	15
1353	<i>The Invisible Man: A Grotesque Romance</i>	English	69,098	8
1354	<i>The Last Man</i>	English	239,038	19
1355	<i>The Law and the Lady</i>	English	188,599	20
1356	<i>The Long Goodbye</i>	English	151,143	8
1357	<i>The Moon and Sixpence</i>	English	100,210	5
1358	<i>The Moonstone: A Romance</i>	English	265,674	19
1359	<i>The Mysteries of Udolpho: A Romance</i>	English	403,637	21
1360	<i>The Old Wives' Tale</i>	English	299,707	4
1361	<i>The People of the Mist</i>	English	193,417	1
1362	<i>The Picture of Dorian Gray</i>	English	106,865	5
1363	<i>The Razor's Edge</i>	English	148,535	2
1364	<i>The Red and the Black: A Chronicle of 1830</i>	English	257,076	2
1365	<i>The Roots of the Mountains</i>	English	212,977	5
1366	<i>The Scarlet Letter</i>	English	117,841	6
1367	<i>The Sun Also Rises</i>	English	96,280	1
1368	<i>The Trail of the Serpent</i>	English	205,726	14
1369	<i>The Well at the World's End</i>	English	304,222	13
1370	<i>The Wind in the Willows</i>	English	83,485	10
1371	<i>The Woman in White</i>	English	326,156	12
1372	<i>To the Lighthouse</i>	English	93,387	2
1373	<i>Toilers of the Sea</i>	English	198,712	11
1374	<i>Twenty Thousand Leagues Under the Seas</i>	English	205,206	5
1375	<i>Twenty Years After</i>	English	349,814	8
1376	<i>Uncle Tom's Cabin</i>	English	263,615	4
1377	<i>Vanity Fair</i>	English	424,757	10
1378	<i>Villette</i>	English	274,495	21
1379	<i>White Fang</i>	English	97,089	18
1380	三个火枪手	Chinese	275,541	11
1381	三体	Chinese	551,978	55
1382	书剑恩仇录	Chinese	373,300	12
1383	亮剑	Chinese	236,235	19
1384	倚天屠龙记	Chinese	696,632	6
1385	四世同堂	Chinese	480,143	9
1386	围城	Chinese	156,627	46
1387	基督山伯爵	Chinese	536,808	4
1388	孽子	Chinese	93,595	4
1389	孽海花	Chinese	191,706	18
1390	官场现形记	Chinese	472,021	45
1391	尼罗河上的惨案	Chinese	89,222	20
1392	平凡的世界	Chinese	538,468	24
1393	沙丘	Chinese	244,388	2
1394	球状闪电	Chinese	105,150	15
1395	白夜行	Chinese	190,339	25
1396	白鹿原	Chinese	332,728	41
1397	百年孤独	Chinese	155,996	3
1398	解忧杂货店	Chinese	98,577	21
1399	许三观卖血记	Chinese	86,558	8
1400	达芬奇密码	Chinese	166,332	8
1401	连城诀	Chinese	164,682	30
1402	金瓶梅	Chinese	598,581	20
1403	金粉世家	Chinese	648,839	17
	飞狐外传	Chinese	322,264	42
	骆驼祥子	Chinese	94,597	6

Table 9: List of novels in SagaScale. Token counts are calculated using the tokenizer from <https://huggingface.co/deepseek-ai/DeepSeek-R1-0528>.