

LIT-RAGBench: Benchmarking Generator Capabilities of Large Language Models in Retrieval-Augmented Generation

Koki Itai^{1,2}, Shunichi Hasegawa¹, Yuta Yamamoto¹,
Gouki Minegishi^{1,3}, Masaki Otsuki^{1,3}

¹neoAI Inc., Tokyo, Japan

²Tokyo Metropolitan University, Tokyo, Japan

³The University of Tokyo, Tokyo, Japan

{k.itai, s.hasegawa, y.yamamoto, g.minegishi, m.otsuki}@neoai.jp

Abstract

Retrieval-Augmented Generation (RAG) is a framework in which a Generator, such as a Large Language Model (LLM), produces answers by retrieving documents from an external collection using a Retriever. In practice, Generators must integrate evidence from long contexts, perform multi-step reasoning, interpret tables, and abstain when evidence is missing. However, existing benchmarks for Generators provide limited coverage, with none enabling simultaneous evaluation of multiple capabilities under unified conditions. To bridge the gap between existing evaluations and practical use, we introduce LIT-RAGBench (the Logic, Integration, Table, Reasoning, and Abstention RAG Generator Benchmark), which defines five categories: Integration, Reasoning, Logic, Table, and Abstention—each further divided into practical evaluation aspects. LIT-RAGBench systematically covers patterns combining multiple aspects across categories. By using fictional entities and scenarios, LIT-RAGBench evaluates answers grounded in the provided external documents. The dataset consists of 114 human-constructed Japanese questions and an English version generated by machine translation with human curation. We use LLM-as-a-Judge for scoring and report category-wise and overall accuracy. Across API-based and open-weight models, no model exceeds 90% overall accuracy. By making strengths and weaknesses measurable within each category, LIT-RAGBench serves as a valuable metric for model selection in practical RAG deployments and for building RAG-specialized models.

Keywords: Large Language Models, Retrieval-Augmented Generation, Evaluation Methodologies, Question Answering, Corpus (Creation, Annotation, etc.)

1. Introduction

Recent advancements in Large Language Models (LLMs) have significantly enhanced their capabilities across multiple domains (Brown et al., 2020; OpenAI et al., 2024; Minaee et al., 2024). However, several challenges have been reported, including factually ungrounded hallucinations (Cao et al., 2020; Huang et al., 2025), outdated information (He et al., 2022), and limited domain-specific expertise (Li et al., 2023; Huang et al., 2025; Zhu et al., 2025). Retrieval-Augmented Generation (RAG) has emerged as a valuable approach to address these challenges (Lewis et al., 2020; Fan et al., 2024). RAG is a framework in which a Generator, such as an LLM, produces answers based on documents retrieved by a Retriever from an external collection. In practical applications, the Generator is required to accurately extract evidence from context while demonstrating multifaceted abilities, such as referencing and integrating evidence from multiple documents, performing multi-hop reasoning, and interpreting tabular data. Although many benchmarks (Chen et al., 2024; Krishna et al., 2025) have been proposed to evaluate the Generator, they do not adequately cover the diverse capabilities needed in real-world RAG scenarios. Moreover, practical scenarios often require multiple

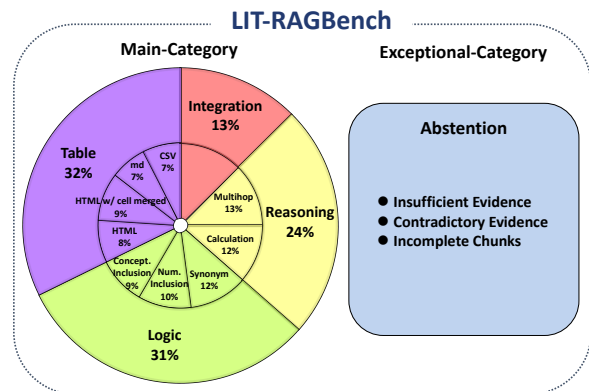


Figure 1: Illustration of the evaluation categories of LIT-RAGBench. These categories reflect the capabilities required of the Generator in RAG based on real-world scenarios.

capabilities simultaneously, yet no existing benchmark systematically evaluates such combinations under unified conditions.

To bridge the gap between existing evaluations and practical use, this study proposes LIT-RAGBench (the Logic, Integration, Table, Reasoning, and Abstention RAG Generator Benchmark), designed to evaluate the Generator independently of retrieval quality. LIT-RAGBench defines funda-

mental Generator capabilities as “evaluation categories”, each comprising detailed “evaluation aspects” derived from real-world RAG use cases.

The overall structure of these categories is illustrated in Figure 1. Specifically, LIT-RAGBench evaluates five core aspects of Generator capability: (1) *Integration*: generating information from multiple sources; (2) *Reasoning*: inferring implicit conclusions from retrieved information; (3) *Logic*: maintaining semantic and deductive consistency; (4) *Table*: comprehending and interpreting tabular data; and (5) *Abstention*: refraining from answering when reliable evidence cannot be established. The evaluation datasets were constructed through a hybrid approach that combines LLM-based synthetic data generation and human curation, followed by manual filtering to ensure quality.

Evaluation experiments on major LLMs, including both API-based and open-weight models, revealed distinct performance patterns across categories. No model exceeded 90% overall accuracy, with performance variations across categories revealing each model’s strengths and weaknesses. These findings demonstrate that LIT-RAGBench serves as a useful metric for model selection in practical RAG deployment and for building RAG-specialized models. To facilitate reproducibility and further research, we release the dataset, the LLM prompts used for dataset construction and evaluation, and the corresponding code¹.

2. Preliminaries

This section formalizes the RAG process for subsequent explanation. Let \mathcal{R} and \mathcal{G} denote the Retriever and Generator components, respectively. \mathcal{R} takes a query x_r as input and outputs a set of related text segments $C = \{c_1, \dots, c_n\}$ from an external data source \mathcal{E} such as a database or the Web. c is a text segment called a chunk, which is created by splitting documents stored in \mathcal{E} into shorter segments for search efficiency. The value n is the number of chunks retrieved by \mathcal{R} , specified by the developer. x_r is typically generated from a user question q by a search query generator f , such as a small language model.

$$C = \mathcal{R}(x_r) \quad x_r = f(q)$$

\mathcal{G} produces an answer y based on the input context x_g . x_g typically includes a task instruction τ , a user query q , and the retrieved chunks C .

$$y = \mathcal{G}(x_g), \quad \text{where } x_g = (\tau, q, C)$$

C is implicitly divided into a relevant chunk set C^+ , which contains evidence for generating an answer to q , and an irrelevant chunk set C^- that does

not contain supporting evidence. From the perspective of retrieval performance, \mathcal{R} does not guarantee retrieving all relevant chunks in C^+ (Muennighoff et al., 2023). Therefore, \mathcal{G} needs to appropriately extract evidence relevant to q from C^+ within C and generate an answer.

3. Related Work

In recent years, various benchmarks have been proposed to systematically evaluate the performance of \mathcal{G} . FRAMES (Krishna et al., 2025) provides an integrated, end-to-end evaluation of both \mathcal{R} and \mathcal{G} capabilities across three aspects—factuality, retrieval, and reasoning. The tasks require multi-document integration and involve temporal and numerical reasoning. This framework highlights how factual consistency and reasoning depth can be jointly measured under controlled retrieval conditions. RAGBench (Friel et al., 2025) proposes the TRACe framework, which assesses \mathcal{G} performance along three interpretable dimensions: Utilization (how much retrieved context is actually used), Adherence (faithfulness and hallucination control relative to the context), and Completeness (coverage of relevant information). It also measures retriever Relevance separately, enabling isolation of retrieval and generation errors. RAGTruth (Niu et al., 2024) provides a dataset for analyzing and detecting hallucinations in RAG, defining four fine-grained types—Evident/Subtle Conflict and Evident/Subtle Baseless Information—to evaluate whether model outputs remain consistent with supporting documents. This benchmark highlights the difficulty of detecting subtle inconsistencies that remain semantically plausible but are factually incorrect. RGB (Chen et al., 2024) evaluates \mathcal{G} along four axes: noise robustness, abstention or negative rejection, multi-document information integration, and counterfactual robustness against misinformation. It tests models under varied, noisy, or conflicting evidence, providing insights into their ability to extract, integrate, or abstain appropriately when uncertainty arises.

Although these benchmarks contribute valuable insights, they primarily address limited aspects of \mathcal{G} ’s behavior or evaluate each skill in isolation. In practical RAG applications, models must often interpret complex tables and perform multi-step reasoning simultaneously—for example, combining numerical computation with multi-hop inference across heterogeneous contexts. Existing work has not yet captured this compound complexity, leaving a gap in evaluating \mathcal{G} ’s real-world robustness. To address this gap, our benchmark systematically assesses \mathcal{G} ’s performance under such intertwined conditions, enabling more realistic and comprehensive evaluation of \mathcal{G} capabilities required for practi-

¹<https://github.com/Koki-Itai/LIT-RAGBench>

cal RAG scenarios.

4. LIT-RAGBench

4.1. Evaluation Framework

4.1.1. Evaluation Categories and Aspects

LIT-RAGBench systematizes the core capabilities of \mathcal{G} into five evaluation categories (Figure 1), with each category subdivided into evaluation aspects based on practical use cases. The five evaluation categories are: (1) *Integration*, (2) *Reasoning*, (3) *Logic*, (4) *Table*, and (5) *Abstention*. The first four categories, collectively referred to as *Main*, represent the core capabilities required for \mathcal{G} to generate a correct answer to q . *Abstention* is defined as an *exceptional category*, distinct from *Main*, as it evaluates \mathcal{G} 's ability to withhold an answer appropriately.

(1) Integration. This category addresses cases where evidence is dispersed across multiple documents, requiring \mathcal{G} to extract and integrate relevant information from each source. This capability has been examined in existing benchmarks, such as Integration in RGB (Chen et al., 2024) and Multiple Constraints in FRAMES (Krishna et al., 2025). This category focuses on *integrating information from multiple sources* ($|C^+| \geq 2$) as an evaluation aspect. Single-source extraction ($|C^+| = 1$) is a fundamental RAG operation that naturally co-occurs with other aspects and is not treated independently. LIT-RAGBench targets integration from $2 \leq |C^+| \leq 3$ sources for simplicity and practicality.

(2) Reasoning. This category assesses \mathcal{G} 's reasoning capabilities across two dimensions. *Multi-hop Reasoning* evaluates \mathcal{G} 's ability to combine information from multiple documents to reach conclusions not explicitly stated in any single source. Questions were created using benchmarks such as HotpotQA (Yang et al., 2018) and JEMHopQA (Ishii et al., 2024). *Numerical Calculation* assesses deriving metrics (totals, averages, profit margins, growth rates) through common-sense arithmetic and business knowledge when formulas are not provided—a challenge for LLMs, compared to language tasks (Li et al., 2025; Yang et al., 2025).

(3) Logic. This category evaluates \mathcal{G} 's ability to interpret logical and linguistic relations between query q and retrieved contexts C despite lexical or semantic discrepancies. Since \mathcal{E} documents are typically unknown to users, phrasing mismatches frequently arise, requiring \mathcal{G} to resolve them through logical understanding across three aspects. *Synonym Interpretation* recognizes equivalent expressions (e.g., "10 thousand yen" and "10,000 yen"),

including multilingual terms or abbreviations, with LIT-RAGBench emphasizing numerical and unit expressions common in practical RAG scenarios. *Numerical Inclusion Interpretation* assesses understanding of numerical conditions—e.g., determining whether a 35-year-old meets "20 or older and under 40" requires boundary-inclusive reasoning. *Conceptual Inclusion Interpretation* evaluates recognizing hierarchical relations, such as identifying "noise-canceling earphones" as "electronic devices" prohibited under device-banning rules.

(4) Table. This category evaluates \mathcal{G} 's ability to interpret and extract information from tabular formats in retrieved contexts C . RAG documents often contain tables in structured formats (HTML, Markdown, CSV) alongside text, requiring \mathcal{G} to understand table structure and identify relevant data—capabilities assessed by benchmarks like MMQA (Wu et al., 2025). *HTML* tables use standard tags requiring \mathcal{G} to map header-data relationships and comprehend hierarchical associations. *HTML with merged cells* employs `rowspan` or `colspan` attributes, creating multi-row/column dependencies that complicate positional relationships, challenging LLMs (Zhao et al., 2023; Sui et al., 2024). *Markdown* tables use pipe delimiters with lightweight syntax but lack explicit typing or hierarchy, requiring \mathcal{G} to infer row-column relationships without structural cues. *CSV* data provide comma/tab-separated values without headers or formatting, necessitating schema and semantic inference from contextual information alone.

(5) Abstention. This category evaluates \mathcal{G} 's ability to refrain from answering when sufficient evidence is unavailable. Withholding answers under inappropriate conditions is essential for response reliability and mitigating hallucinations (Huang et al., 2025), with benchmarks like Abstention-Bench (Kirichenko et al., 2025) and RGB's negative rejection task (Chen et al., 2024). We define this capability across three aspects. *Insufficient Evidence* occurs when retrieved contexts lack necessary information; since \mathcal{R} doesn't guarantee C^+ , when $x_g = (\tau, q, C^-)$, \mathcal{G} should explicitly indicate it cannot answer, as LLMs tend to hallucinate without supporting evidence (Chen et al., 2024). *Contradictory Evidence* arises when documents provide conflicting information, requiring \mathcal{G} to recognize inconsistencies and either abstain or explicitly identify contradictions through cross-document consistency checking. *Incomplete Chunk* represents scenarios where retrieved content is fragmentary due to chunking boundaries splitting semantically connected information; when \mathcal{R} retrieves partial segments—exceptionally long definitions or tables divided across boundaries— \mathcal{G} should refrain from

answering despite overlapping tokens potentially mitigating this issue.

4.1.2. Category Composition

This section formalizes the evaluation categories and aspects defined above. Let Θ denote the set of five evaluation categories treated in this study. The abbreviations $\theta_I, \theta_R, \theta_L, \theta_T$, and θ_A are used for *Integration, Reasoning, Logic, Table*, and *Abstention*, respectively. Collectively, these are expressed as $\Theta = \theta_I, \theta_R, \theta_L, \theta_T, \theta_A$. We define $\Theta_{\text{Main}} = \Theta \setminus \{\theta_A\}$. θ_A is independent and does not co-occur with aspects from Θ_{Main} . Each evaluation category $\theta \in \Theta$ comprises multiple evaluation aspects associated with it. Let Φ denote the set of evaluation aspects, and let Φ_θ represent the subset of aspects belonging to a particular evaluation category θ . Then, Φ is defined as the union of all Φ_θ for $\theta \in \Theta$.

LIT-RAGBench is constructed by combining evaluation aspects belonging to one or two categories. For each evaluation problem q , the set of aspects $\psi(q)$ satisfies:

$$\Psi(q) \subseteq \Phi, \quad 1 \leq |\Psi(q)| \leq 2$$

$$\forall \phi_i, \phi_j \in \Psi(q), \quad \phi_i \in \Phi_{\theta_m}, \phi_j \in \Phi_{\theta_n} \Rightarrow m \neq n$$

Thus, category composition is formalized through the aspect $\Psi(q)$ and the non-overlap constraint across categories. An example of co-occurrence that satisfies these constraints is shown in Figure 2.

Consequently, LIT-RAGBench enables comprehensive and realistic evaluation of the multifaceted capabilities of \mathcal{G} , as well as quantitative analyses across categories and aspects.

4.2. Dataset Construction

The dataset D consists of m samples, each containing a question q_i , its answer a_i , and document sets C_i^+ and C_i^- .

$$D = \{(q_i, a_i, C_i^+, C_i^-, \Psi(q_i)) \mid 1 \leq i \leq m\}.$$

To ensure sufficient context, $|C_i^+| + |C_i^-| \geq 8$ is required.

During evaluation, the task instruction τ specifies two directives: (1) answer based on the given context, and (2) state inability to answer when supporting evidence is unavailable. For Θ_{Main} , answers are generated with $x_g = (\tau, q_i, C_i)$, where $C_i = C_i^+ \cup C_i^-$. We randomize the concatenation order of C_i at input time. This removes confounding from retriever ranking, mitigates position bias (Liu et al., 2024), and evaluates the use of order-invariant evidence and robustness to rank perturbations across retrievers. For the *Insufficient Evidence* evaluation in θ_A , we use $x_g = (\tau, q_i, C_i^-)$, again with randomized order.

Reasoning (Multi-hop)

×

Table (HTML with merged cells)

q : What major research theme was officially adopted by GreenWave in 2024 regarding the marine plastic issue?

C_1^+ : GreenWave concluded a partnership with Seikai University; in Jul 2024, the joint project was selected for the Ministry of the Environment's Blue Innovation Research Grant. . . .

C_2^+ : List of Blue Innovation Research Grant

Univ.	Field	Summary
Seikai Univ.	Marine plastics	Development of biopolymers that decompose microplastics
	Marine ecosystems	Assessment of impact of drifting debris on coastal organisms
Toyo Inst.	Hydrogen energy	Optimization of algae-derived hydrogen production system

Rationale: From C_1^+ , the GreenWave↔Seikai Univ. project was adopted; the merged-cell row in C_2^+ gives the adopted theme. a : Development of a biopolymer that degrades microplastics

Figure 2: Example of a co-occurrence across evaluation categories (θ_R and θ_T)

4.2.1. Dataset Creation Procedure

The dataset was manually created by three native Japanese speakers. They designed question–answer scenarios and constructed relevant document sets following predefined guidelines. For quality assurance, two of the three contributors independently reviewed all items, and only samples that passed both inspections were retained in the final dataset. To enhance the efficiency of the dataset creation process, we additionally used GPT-5 as an auxiliary tool (see Footnote 1 for the prompt templates used).

Decision QA scenarios. QA scenarios were designed based on evaluation aspect patterns, assuming practical RAG use cases. Drawing on the methodology of Kirchenbauer et al. (Kirchenbauer et al., 2025), QA scenarios were created using fictional knowledge entities like company names, product names, and personal names to prevent LLMs from answering based on their pre-trained knowledge when generating answers.

Creation of questions, relevant document sets, and answers. The automatically generated texts were manually reviewed for quality, and only those judged to appropriately reflect the evaluation aspects were adopted. Additionally, the length of each document in C_i^+ was adjusted to approximately 512 tokens using tiktoken² to align with typical chunk lengths.

²<https://github.com/openai/tiktoken>

Creation of irrelevant document sets. Assuming the output of \mathcal{R} , C^- are designed to be related to q_i or C_i^+ , for example, by containing keywords from q_i , even though they do not serve as direct evidence. Following the same procedure as C^+ , manual quality inspection and token length adjustment were conducted.

Human-based filtering. For the obtained q_i , a_i , C_i^+ , C_i^- , and $\Psi(q_i)$, verification was conducted from the following perspectives: (1) whether the question appropriately corresponds to the target evaluation aspect pattern, (2) whether it is possible to derive a_i from C_i^+ as evidence for q_i , (3) whether the problem can not be answered using only the LLM’s pre-trained knowledge, and (4) whether the fictional information contradicts real facts. Only problems that satisfied all criteria were added to D . In cases of non-compliance, the problem was discarded, and the procedure was repeated from scenario design in step 1. The manual verification was conducted by qualified annotators who are native Japanese speakers.

4.2.2. Statistics

Following the creation procedure described above, a 54-question Japanese QA dataset was constructed and included in Θ_{Main} . There are 12 questions with $|\Psi(q_i)| = 1$ and 42 questions with $|\Psi(q_i)| = 2$. From this 54-question Japanese QA dataset, an additional 54 QA questions for the *In-sufficient Evidence* aspect of θ_A were created by converting C^+ to an empty set. For the *Contradictory Evidence* and *Incomplete Chunk* aspects of θ_A , three questions were randomly selected from the constructed Θ_{Main} QA dataset for each aspect, and C^+ was manually edited to construct the corresponding QA data. This process resulted in a Japanese evaluation dataset with a total of 114 questions. An English evaluation dataset of the same scale was constructed by translating the Japanese dataset using the GPT-5 API-based LLM.

For θ_A , the aspect distribution was *In-sufficient Evidence* 34.6%, *Contradictory Evidence* 1.9%, and *Incomplete Chunk* 1.9%. The proportion of *In-sufficient Evidence* is relatively large because these instances can be directly created by removing C^+ from the corresponding Θ_{Main} .

5. Experiments

In this section, we present the evaluation results and analysis of LLMs using LIT-RAGBench. We evaluate both API-based and open-weight models on the Japanese and English datasets of LIT-RAGBench.

Table 1: Evaluated LLMs (API-based and open-weight, † denotes reasoning models)

Model Name	Version	Developer
<i>API Models</i>		
GPT-5†	2025-08-07	OpenAI
GPT-5-mini†	2025-08-07	OpenAI
GPT-5-nano†	2025-08-07	OpenAI
o3†	2025-04-16	OpenAI
o4-mini†	2025-04-16	OpenAI
GPT-4.1	2025-04-14	OpenAI
GPT-4.1-mini	2025-04-14	OpenAI
Gemini-2.5-Flash	2025-05-17	Google
Gemini-2.5-Pro†	2025-06-17	Google
Claude-Sonnet-4	2025-05-14	Anthropic
<i>Open Models</i>		
Gemma-3-27B-Instruct ²	–	Google
Llama-3.1-8B-Instruct ³	–	Meta
Llama-3.3-70B-Instruct ⁴	–	Meta
Qwen3-235B-A22B-Instruct ⁵	–	Alibaba
Qwen3-235B-A22B-Thinking† ⁶	–	Alibaba

5.1. Experimental Settings

Table 1 lists the evaluated API-based and open-weight models, respectively. For Reasoning Language Models (RLMs) such as GPT-5 and o3, the length of reasoning token generation can be specified, and in all cases, the maximum generation length setting is used. For models that support temperature configuration, `temperature` is set to 0.0 and `top_p` to 1.0.

5.2. Evaluation Method

To evaluate the accuracy of answers generated by \mathcal{G} , this study employs automated evaluation using an LLM-as-a-Judge. The prompt used for the LLM-as-a-Judge is provided in the supplementary repository (see Footnote 1). While LLM-based automatic evaluation methods have raised concerns about reliability, our evaluation task is a relatively straightforward closed-ended comparison between the generated answer and the reference answer, which is more amenable to automated judgment. Previous studies have demonstrated a strong correlation between LLM-as-a-Judge evaluations and human judgments (Zheng et al., 2023; Liu et al., 2023; Kim et al., 2024). Let \mathcal{J} denote the LLM-based evaluator. In this study, we implement \mathcal{J} using GPT-4.1 (OpenAI API, version dated 2025-04-14), which is prompted to determine whether the generated answer y_i is semantically consistent with the reference answer a_i . \mathcal{J} outputs a binary

²google/gemma-3-27b-it

³meta-llama/Llama-3.1-8B-Instruct

⁴meta-llama/Llama-3.3-70B-Instruct

⁵Qwen/Qwen3-235B-A22B-Instruct-2507

⁶Qwen/Qwen3-235B-A22B-Thinking-2507

label for each instance:

$$\mathcal{J}(q_i, a_i, y_i) = \begin{cases} 1 & \text{if } y_i \text{ is consistent with } a_i \\ 0 & \text{otherwise} \end{cases}$$

This study uses accuracy as the performance metric. Let Q and Q_θ denote the set of all questions and the subset belonging to category $\theta \in \Theta$, respectively. Accuracy for each category is defined as:

$$\text{Accuracy}(\theta) = \frac{1}{|Q_\theta|} \sum_{q_i \in Q_\theta} \mathcal{J}(q_i, a_i, y_i)$$

The overall accuracy, denoted as $\overline{\text{Accuracy}}$, is obtained by averaging $\text{Accuracy}(\theta)$ over all θ . When a question involves multiple aspects and is answered correctly, it is counted as correct for all applicable aspects.

5.3. Results

The evaluation results for $\overline{\text{Accuracy}}$ on the Japanese and English datasets are shown in Figure 3. Table 2 presents $\text{Accuracy}(\theta)$ for each $\theta \in \Theta$.

As shown in Figure 3, no model achieved an $\overline{\text{Accuracy}}$ of 0.9 in either the Japanese or English evaluations. GPT-5 achieved the highest score of 0.872 in both languages. Among open-weight models, Qwen3-235B-A22B-Instruct and Qwen3-235B-A22B-Thinking showed the best performance, with scores of 0.859 and 0.821, respectively. In contrast, open-weight models with small to medium parameter sizes, such as Llama-3.1-8B-Instruct and Gemma-3-27B-Instruct, showed generally low scores.

5.4. Analysis

5.4.1. Integration

When examining outputs that produced incorrect answers, errors frequently occurred when C^+ lacked explicit lexical cues or contained multiple similar pieces of information differing only in minor details. For example, when extracting prices from rate tables of three meeting room reservation systems—Company A, B, and C—only Company A included additional notes on pricing. Correct extraction required accounting for these notes, but the model returned only the basic fee, as with Companies B and C. These observations indicate that \mathcal{G} often fails when data sources vary in quality or granularity, reaffirming the importance of preprocessing, such as document structuring and normalization, in practical RAG deployments.

5.4.2. Reasoning

In the numerical calculation aspect, o3 solved all tasks correctly. In contrast, o4-mini and GPT-5

generally followed appropriate reasoning steps but made minor arithmetic errors in intermediate or final calculations, consistent with previous findings (Li et al., 2025; Yang et al., 2025). Small-to-medium-scale open-weight and API-based mini models showed limited reasoning capability. These models often failed to identify intermediate entities in the reasoning chain or refused to answer when explicit lexical cues were absent.

For instance, in a question requiring inference of a company's 2024 ranking from its 2025 position and relative improvement, such models recognized both facts but failed to infer the implicit 2024 value, resulting in incorrect responses. This limitation frequently co-occurred with θ_T , suggesting difficulty in integrating implicit relationships or reasoning across multiple documents rather than from explicitly stated facts, consistent with prior work (Zhuang et al., 2024).

5.4.3. Logic

As shown in Table 2, most API-based models achieved high accuracy in both Japanese and English for θ_L . However, some category-specific errors were observed. For instance, when asked about data capacity in GB, models answered "500 MB" from the source instead of the correct "0.5 GB." Although the relevant text was correctly identified, such responses were judged incorrect under the benchmark criteria. These issues can be mitigated through prompts that enforce output-unit alignment, emphasizing the importance of operational rule design. A Japanese-specific error was also found in unit conversion tasks common in business contexts: models often generated "760 million yen" instead of the correct "7.6 billion yen." Addressing these language-dependent hallucinations is practically important, and future work will extend the benchmark to include such cases.

5.4.4. Table

From Table 2, Gemini-2.5-Flash achieved the highest θ_T score in both Japanese and English. Most LLMs understood basic row-column structures but struggled with merged-cell tables, making accurate information retrieval difficult, especially for small-to-medium-scale open-weight models.

When large tables exceeding 512 tokens were split into chunks in C^+ , nearly all models failed to extract relevant data. We anticipated this issue during dataset construction and inserted header information into each chunk during preprocessing. However, because the chunks were intentionally shuffled to simulate real-world disorder, models likely failed to recognize the overall structure and often abstained despite sufficient evidence.




Figure 3: Evaluation results for overall accuracy. Blue bars represent Japanese scores, and red bars represent English scores.

Table 2: Category-wise accuracy comparison across models for Japanese and English evaluations. The best and second-best performances in each column are highlighted in bold and underlined, respectively. In cases of ties, all models achieving the same best or second-best score are highlighted accordingly.

Model	Integration		Reasoning		Logic		Table		Main (avg)		Abstention	
	ja	en	ja	en	ja	en	ja	en	ja	en	ja	en
<i>API Models</i>												
GPT-5	0.833	0.833	0.870	0.739	0.867	0.900	0.839	0.839	0.852	0.828	0.900	0.933
GPT-5-mini	0.917	0.750	0.826	0.696	0.867	0.867	0.774	0.806	0.846	0.780	0.883	0.917
GPT-5-nano	0.750	0.667	0.565	0.739	0.733	0.833	0.677	0.710	0.681	0.737	0.767	0.817
o3	0.833	0.833	0.957	0.870	0.900	0.900	0.839	0.806	0.882	0.852	0.817	0.833
o4-mini	0.917	0.833	0.913	0.783	0.900	0.900	0.871	0.839	0.900	0.839	0.783	0.900
GPT4.1	0.833	1.000	0.739	0.739	0.800	0.800	0.839	0.806	0.803	0.836	0.850	0.933
GPT4.1 mini	0.917	0.833	0.870	0.826	0.800	0.833	0.806	0.871	0.848	0.841	0.783	0.867
Claude-Sonnet-4	0.750	0.917	0.783	0.565	0.700	0.600	0.677	0.516	0.728	0.650	0.950	0.967
Gemini-2.5-Flash	0.917	0.917	0.783	0.826	0.867	0.833	0.871	0.903	0.860	0.870	0.767	0.883
Gemini-2.5-Pro	0.667	0.833	0.870	0.696	0.833	0.833	0.774	0.774	0.786	0.784	0.800	0.883
<i>Open Models</i>												
Gemma-3-27B-Instruct	0.667	0.667	0.348	0.609	0.567	0.733	0.484	0.710	0.517	0.680	0.733	0.850
Llama-3.1-8B-Instruct	0.167	0.500	0.130	0.261	0.367	0.500	0.355	0.323	0.255	0.396	0.600	0.833
Llama-3.3-70B-Instruct	0.750	0.583	0.478	0.609	0.733	0.700	0.677	0.581	0.660	0.618	0.767	0.917
Qwen3-235B-A22B-Instruct	0.917	0.833	0.870	0.783	0.833	0.867	0.839	0.742	0.865	0.806	0.867	0.900
Qwen3-235B-A22B-Thinking	0.917	0.750	0.826	0.870	0.767	0.867	0.774	0.774	0.821	0.815	0.900	0.917

These findings highlight that in practical applications, preprocessing such as table restructuring and reordering chunks is essential before inputting to \mathcal{G} , since large tables are often unintentionally split.

5.4.5. Abstention

From Table 2, Claude-Sonnet-4 achieved notably high θ_A scores in both Japanese and English. In the Insufficient Evidence category, numerous mod-

els produced knowledge that appeared contextually plausible but lacked evidential grounding. Such behavior has been identified in a previous study as a type of hallucination arising from fabricated knowledge (Huang et al., 2025; Kirichenko et al., 2025). For Contradictory Evidence, Claude-Sonnet-4 detected inconsistencies in all tasks and correctly abstained, while many models relied on a single source and failed to cross-reference evidence. In the Incomplete Chunk aspect, most models appropriately abstained, but when faced with ambigu-

ous questions that could be answered using general knowledge, they often hallucinated responses. Such ambiguity frequently arises in practical RAG scenarios, as users expect responses grounded in retrieved evidence \mathcal{E} . These results underscore the importance of prompt design to guide appropriate abstention behavior in RAG systems.

Based on quantitative and qualitative analyses, Claude-Sonnet-4 demonstrates a strong ability to refrain from answering. However, this also reveals a tendency toward Over-Abstention even when a valid response is possible. In this study, instances within θ_{Main} where models abstained despite having sufficient information—resulting in incorrect outcomes—were quantified as the Over-Abstention Rate (Table 3). Based on the average Over-Abstention Rate (*avg* in Table 3), notable disparities are observed across models. Claude-Sonnet-4 shows the highest rate (0.259), followed by Llama-3.1-8B-Instruct (0.213), indicating a strong tendency to over-abstain. Among open-weight models, larger models such as Qwen3-235B-A22B-Thinking (0.120) maintain moderate rates, whereas smaller models like Llama-3.1-8B-Instruct tend to over-abstain, likely due to weaker reasoning capacity in Θ_{Main} . These patterns suggest that underdeveloped generation capabilities may lead models to default to abstention when uncertain, even in answerable cases. In a previous study designed to measure over-abstention tendencies, a strong correlation between safety alignment and usefulness was also observed, indicating a clear trade-off between the two dimensions (Cui et al., 2025). Consistent with our results, the study reported that Claude exhibited the strongest tendency toward over-abstention among the evaluated models. These results indicate that while a higher abstention rate may reflect cautious alignment behavior, it does not necessarily translate into greater overall usefulness.

5.5. Overall Discussion

Experiments using LIT-RAGBench demonstrated that composite-category evaluations can quantitatively assess the multifaceted abilities of \mathcal{G} . Unlike existing benchmarks that assess each category independently, LIT-RAGBench enables absolute comparisons across categories, revealing integrated performance differences and weaknesses among models. The experiments also showed that LIT-RAGBench effectively captures variations in abstention tendencies. While improved abstention behavior enhances response reliability, it may reduce accuracy on answerable questions. These findings indicate that LLMs can benefit from further training and prompt optimization to balance abstention and accuracy.

Table 3: Over-Abstention Rates on the Japanese and English Evaluation Sets for the Main Category. The column *avg* indicates the mean of the Japanese (*ja*) and English (*en*) rates.

Model	ja	en	avg
Claude-Sonnet-4	0.148	0.370	0.259
Llama-3.1-8B-Instruct	0.204	0.222	0.213
Llama-3.3-70B-Instruct	0.093	0.222	0.157
Gemma-3-27B-Instruct	0.167	0.074	0.120
Qwen3-235B-A22B-Thinking	0.130	0.111	0.120
GPT-4.1	0.056	0.111	0.083
Qwen3-235B-A22B-Instruct	0.056	0.111	0.083
GPT-5-mini	0.074	0.074	0.074
GPT-5	0.037	0.093	0.065
Gemini-2.5-Flash	0.074	0.056	0.065
GPT-5-nano	0.074	0.019	0.046
o3	0.019	0.074	0.046
o4-mini	0.019	0.074	0.046
Gemini-2.5-Pro	0.000	0.093	0.046
GPT-4.1 mini	0.037	0.019	0.028

6. Limitations

Our main limitations are the small sample size and the imbalance across aspects. We designed an evaluation framework targeting realistic failure cases of \mathcal{G} , which have often been overlooked in previous studies, and built the accompanying dataset through careful human curation. This process yielded a compact, high-quality dataset that covers the minimum occurrence patterns for each evaluation aspect, but it remains smaller than many existing benchmarks. While maintaining human-verified evaluation quality is important (Maheshwari et al., 2024; Gill et al., 2025), expanding the benchmark to be more diverse and larger in scale is an essential direction for future work.

7. Conclusion

We constructed LIT-RAGBench, a benchmark comprising five evaluation categories designed from practical failure cases in real-world RAG systems. Experiments on major LLMs showed that no model exceeded 90% overall accuracy, and performance gaps were observed across categories. These findings demonstrate that LIT-RAGBench provides a systematic and interpretable framework for assessing generator strengths and weaknesses in RAG. For future work, we plan to extend the dataset and advance towards Agentic RAG (Singh et al., 2025), where LLMs autonomously plan retrieval and reasoning steps. We will release LIT-RAGBench as open source to support the advancement of RAG research and foster more reliable evaluation of large language models across diverse application domains.

Acknowledgments

We thank the anonymous reviewers for their constructive feedback. We also thank neoAI Inc. for supporting this work by providing computational resources and covering experimental costs. We are grateful to Koshiro Terasawa, Yuji Mochizuki, and Ryo Yagi, for their valuable assistance and insightful advice in constructing this benchmark.

Bibliographical References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Meng Cao, Yue Dong, Jiapeng Wu, and Jackie Chi Kit Cheung. 2020. [Factual error correction for abstractive summarization models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6251–6258, Online. Association for Computational Linguistics.
- Jing Chen, Hangyu Lin, Xu Han, and Le Sun. 2024. [Benchmarking large language models in retrieval-augmented generation](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 17754–17762.
- Justin Cui, Wei-Lin Chiang, Ion Stoica, and Chojui Hsieh. 2025. [OR-bench: An over-refusal benchmark for large language models](#). In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 11515–11542. PMLR.
- Wenqi Fan, Yujuan Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and Qing Li. 2024. [A survey on rag meeting llms: Towards retrieval-augmented large language models](#). In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '24, page 6491–6501, New York, NY, USA. Association for Computing Machinery.
- Robert Friel, Masha Belyi, and Atindriyo Sanyal. 2025. [Ragbench: Explainable benchmark for retrieval-augmented generation systems](#).
- Alexander Gill, Abhilasha Ravichander, and Ana Marasović. 2025. [What has been lost with synthetic evaluation?](#)
- Hangfeng He, Hongming Zhang, and Dan Roth. 2022. [Rethinking with retrieval: Faithful large language model inference](#).
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2025. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *ACM Trans. Inf. Syst.*, 43(2).
- Ai Ishii, Naoya Inoue, Hisami Suzuki, and Satoshi Sekine. 2024. [JEMHopQA: Dataset for Japanese explainable multi-hop question answering](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 9515–9525, Torino, Italia. ELRA and ICCL.
- Seungone Kim, Jamin Shin, Yejin Cho, Joel Jang, Shayne Longpre, Hwaran Lee, Sangdoon Yun, Seongjin Shin, Sungdong Kim, James Thorne, and Minjoon Seo. 2024. [Prometheus: Inducing fine-grained evaluation capability in language models](#).
- John Kirchenbauer, Janny Mongkolsupawan, Yuxin Wen, Tom Goldstein, and Daphne Ippolito. 2025. [A fictional qa dataset for studying memorization and knowledge acquisition](#).
- Polina Kirichenko, Mark Ibrahim, Kamalika Chaudhuri, and Samuel J. Bell. 2025. [Abstentionbench: Reasoning llms fail on unanswerable questions](#).
- Satyapriya Krishna, Kalpesh Krishna, Anhad Mohananeey, Steven Schwarcz, Adam Stambler, Shyam Upadhyay, and Manaal Faruqi. 2025. [Fact, fetch, and reason: A unified evaluation of retrieval-augmented generation](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4745–4759, Albuquerque, New Mexico. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Haoyang Li, Xuejia Chen, Zhanchao Xu, Darian Li, Nicole Hu, Fei Teng, Yiming Li, Luyu Qiu, Chen Jason Zhang, Li Qing, and Lei Chen. 2025. [Exposing numeracy gaps: A benchmark to evaluate fundamental numerical abilities in large language models](#). In *Findings of the Association*

- for *Computational Linguistics: ACL 2025*, pages 20004–20026, Vienna, Austria. Association for Computational Linguistics.
- Xianzhi Li, Samuel Chan, Xiaodan Zhu, Yulong Pei, Zhiqiang Ma, Xiaomo Liu, and Sameena Shah. 2023. [Are ChatGPT and GPT-4 general-purpose solvers for financial text analytics? a study on several typical tasks](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 408–422, Singapore. Association for Computational Linguistics.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. [Lost in the middle: How language models use long contexts](#). *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. [G-eval: NLG evaluation using gpt-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Gaurav Maheshwari, Dmitry Ivanov, and Kevin El Haddad. 2024. [Efficacy of synthetic data as a benchmark](#).
- Shervin Minaee, Tomá Mikolov, Narjes Nikzad, Meysam Asgari Chenaghlu, Richard Socher, Xavier Amatriain, and Jianfeng Gao. 2024. [Large language models: A survey](#). *ArXiv*, abs/2402.06196.
- Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. 2023. [MTEB: Massive text embedding benchmark](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2014–2037, Dubrovnik, Croatia. Association for Computational Linguistics.
- Cheng Niu, Yuanhao Wu, Juno Zhu, Siliang Xu, KaShun Shum, Randy Zhong, Juntong Song, and Tong Zhang. 2024. [RAGTruth: A hallucination corpus for developing trustworthy retrieval-augmented language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10862–10878, Bangkok, Thailand. Association for Computational Linguistics.
- OpenAI et al. 2024. [Gpt-4 technical report](#).
- Aditi Singh, Abul Ehtesham, Saket Kumar, and Tala Talaei Khoei. 2025. [Agentic retrieval-augmented generation: A survey on agentic rag](#).
- Yuan Sui, Mengyu Zhou, Mingjie Zhou, Shi Han, and Dongmei Zhang. 2024. [Table meets llm: Can large language models understand structured table data? a benchmark and empirical study](#). In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining, WSDM '24*, page 645–654, New York, NY, USA. Association for Computing Machinery.
- Jian Wu, Linyi Yang, Dongyuan Li, Yuliang Ji, Manabu Okumura, and Yue Zhang. 2025. [Mmqa: Evaluating llms with multi-table multi-hop complex questions](#). In *International Conference on Learning Representations (ICLR)*.
- Haotong Yang, Yi Hu, Shijia Kang, Zhouchen Lin, and Muhan Zhang. 2025. [Number cookbook: Number understanding of language models and how to improve it](#). In *The Thirteenth International Conference on Learning Representations*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Bowen Zhao, Changkai Ji, Yuejie Zhang, Wen He, Yingwen Wang, Qing Wang, Rui Feng, and Xiaobo Zhang. 2023. [Large language models are complex table parsers](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14786–14802, Singapore. Association for Computational Linguistics.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 46595–46623. Curran Associates, Inc.
- Chenghao Zhu, Nuo Chen, Yufei Gao, Yunyi Zhang, Prayag Tiwari, and Benyou Wang. 2025. [Is your LLM outdated? a deep look at temporal generalization](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7433–7457, Albuquerque, New Mexico. Association for Computational Linguistics.

Ziyuan Zhuang, Zhiyang Zhang, Sitao Cheng, Fangkai Yang, Jia Liu, Shujian Huang, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang, and Qi Zhang. 2024. [EfficientRAG: Efficient retriever for multi-hop question answering](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3392–3411, Miami, Florida, USA. Association for Computational Linguistics.