

# Retrieval Augmented Generation Evaluation in the Era of Large Language Models: A Comprehensive Survey

Aoran GAN<sup>1</sup>, Hao YU<sup>2</sup>, Kai ZHANG<sup>1</sup>, Qi LIU(✉)<sup>1</sup>, Wenyu YAN<sup>1</sup>, Zhenya HUANG<sup>1</sup>,  
Shiwei TONG<sup>3</sup>, Enhong CHEN<sup>1</sup>, Guoping HU<sup>1,4</sup>

1 State Key Laboratory of Cognitive Intelligence, University of Science and Technology of China, Hefei, China

2 McGill University, Montreal, Canada

3 Tencent Company, Shenzhen, China

4 Artificial Intelligence Research Institute, iFLYTEK Co., Ltd, Hefei, China

© Higher Education Press 2025

**Abstract** Recent advancements in Retrieval-Augmented Generation (RAG) have revolutionized natural language processing by integrating Large Language Models (LLMs) with external information retrieval, enabling accurate, up-to-date, and verifiable text generation across diverse applications. However, evaluating RAG systems presents unique challenges due to their hybrid architecture that combines retrieval and generation components, as well as their dependence on dynamic knowledge sources in the LLM era. In response, this paper provides a *comprehensive survey of RAG evaluation methods and frameworks*, systematically reviewing traditional and emerging evaluation approaches, for system performance, factual accuracy, safety, and computational efficiency in the LLM era. We also compile and categorize the RAG-specific datasets and evaluation frameworks, conducting a meta-analysis of evaluation practices in high-impact RAG research. To the best of our knowledge, this work represents the most comprehensive survey for RAG evaluation, bridging traditional and LLM-driven methods, and serves as a critical resource for advancing RAG development.

**Keywords** Retrieval Augmented Generation, System Evaluation, Large Language Model

## 1 Introduction

Retrieval Augmented Generation (RAG) has emerged as a powerful methodology that enhances natural language generation by incorporating information from external knowledge.

This approach significantly improves Large Language Models through non-parametric learning, multi-source knowledge integration, and specialized domain adaptation [1, 2]. By connecting LLMs with external databases, RAG produces responses that are both contextually appropriate and grounded in authoritative, up-to-date information, marking a substantial advancement in developing more sophisticated natural language processing (NLP) systems [3, 4].

As a sophisticated and expansive system that encompasses numerous elements from both the LLM and retrieval domains, RAG can be approximately segmented into two principal sections from a macroscopic viewpoint: retrieval and generation. The retrieval section typically entails diverse operations including preprocessing, dense or sparse retrieval, reranking and pruning, etc [5, 6]. The generation section comprises components such as retrieval planning, the integration of multi-source knowledge, and logical reasoning [7, 8]. Additionally, RAG systems incorporate interconnected upstream and downstream elements such as document chunking, embedding generation, and mechanisms for ensuring security and credibility [9]. The overall performance of RAG systems depends not only on each individual component but also on their interactions and integrated functionality.

When faced with such complex systems, a fundamental and practical question arises regarding the evaluation framework for assessing the efficacy of architectural methodologies governing both the holistic system and its constituent components. This challenge proves particularly pronounced in RAG systems, where three factors - the expansive scope of implementation domains, the heterogeneity of internal components, and the dynamic progression of current developments - collectively render the establishment of a unified system-

atic evaluation paradigm an ongoing research frontier. In response to this, we conducted this survey on RAG Evaluation to gather methods for multi-scale assessment of RAG in recent years. The comprehensiveness of this survey is demonstrated in four aspects: 1) Systematic completeness, encompassing both the evaluation of RAG’s internal components and the system as a whole; 2) Methodological variety, including both traditional statistically-based evaluation metrics and the innovative methods characteristic of the LLM era; 3) Source diversity, incorporating both structured evaluation frameworks, as well as cutting-edge methods scattered across various papers; and 4) Practicality, both in terms of metrics’ definition to be evaluated and their subsequent application. Through this multi-dimensional approach, we aim to provide researchers and practitioners with a comprehensive toolkit for evaluating and improving RAG systems.

The remainder of this paper is organized as follows: Section 2 offers a concise review of the existing LLM-based RAG system to provide the reader with relevant background knowledge. Our comprehensive evaluation is divided into two distinct sections: **Internal Evaluation** (Section 3) and **External Evaluation** (Section 4). Internal Evaluation assesses component level performance and methodology-specific metrics within basic RAG systems, focusing on technical advancement. External evaluation examines system-wide factors like safety and efficiency, emphasizing practical viability. We pay particular attention to the emerging trend of LLM-based evaluation methods, which represent a novel assessment approach unique to the current era. Section 5 presents existing RAG evaluation frameworks, datasets, and methods, providing a practical resource for researchers. Furthermore, we compiled a comprehensive collection of high-level RAG studies spanning multiple dimensions in recent years, and conducted a preliminary analysis and discussion from the perspective of evaluation (Section 6).

## 2 Background

### 2.1 Large Language Model (LLM)

Large Language Models, with billions of parameters, are a class of generative neural language models trained on extensive natural language data [10, 11]. Due to the wide coverage of the training corpus, LLMs are considered to implicitly integrate world knowledge [12]. LLMs are capable of adhering to human instructions or requests through instruction tuning, thus being able to effectively understand and generate human-like text [13]. Its generalization opens up a wide range of applications, such as NLP, signal processing, and recommender systems [14, 15]. However, LLM’s capability remains circumscribed by their training data. It is sometimes predisposed to generating factually inconsistent outputs (hallucinations), particularly when processing novel information beyond training data [16]. Despite the adaptability of LLMs to diverse downstream tasks through post-training

or fine-tuning on specific datasets, these methods encounter challenges related to arithmetic, timeliness, flexibility, or usability (on close models). Optimization techniques during the LLM inference phase have thus garnered significant attention. One of the representative techniques is Prompt Engineering, in which artificially constructed task descriptions and commands are used to enhance LLMs’ understanding of task objectives. In-context learning is designed to enable LLMs to analyze patterns and generalize from task samples, offering substantial advantages in few-shot scenarios [17, 18]. Unlike these approaches, RAG aims to address the issue of knowledge limitations inherent in LLM by incorporating external knowledge. Both LLM and RAG possess complementary strengths: RAG can effectively leverage the superior reasoning capabilities of LLMs, combined with the broad knowledge scope of external data, to explore the potential applications of LLMs more extensively [19]. On the other hand, LLMs can serve as crucial components in RAG, functioning as the decision maker, reasoner, generator, or even evaluating certain aspects of RAG [20, 21].

### 2.2 Retrieval Augmented Generation (RAG)

RAG is a technical framework that enhances NLP systems by integrating external knowledge retrieval, whose core innovation enables extra non-parametric optimization of parameter-fixed neural language models after training, effectively expanding their operational domains while maintaining architectural stability [22]. Prior to the widespread adoption of LLM, scholarly investigations had already established methods for enhancing NLP tasks through external knowledge infusion [23]. Initial researches on RAG adhered to an elementary indexing and reading paradigm [24, 25]. Later formulations delineated two core components: (1) the retriever, which identifies, indexes, filters, and structures relevant knowledge fragments from external data sources; (2) the generator, which synthesizes the curated segments through analysis and logical reasoning to produce outputs [9]. Figure 1 shows the workflow of an RAG system with recommendations of components implementation using LLMs at present. We provide a concise description of each module’s process below.

The retrieval component of RAG systems is inspired by the retrieval technologies in multiple domains, such as information retrieval [26], open-domain question answering [27], and recommender systems [28, 29]. Before the retrieval, it is necessary to construct a suitable corpus for the retrieval component at the beginning. The sources of data are diverse, such as domain-specific datasets like Wikipedia, specialized corpora (e.g., scientific articles, financial reports) [30], or real-time data gathered from web scraping or search engines [31]. The corpus is subsequently filtered and preprocessed to conform to the retrieval-friendly structure via offline chunking and embedding. Chunking involves segmenting large documents into smaller, more manageable units guided by the original structure or context information [32–34]. Embedding (or text vectorization) aims to represent the textual con-

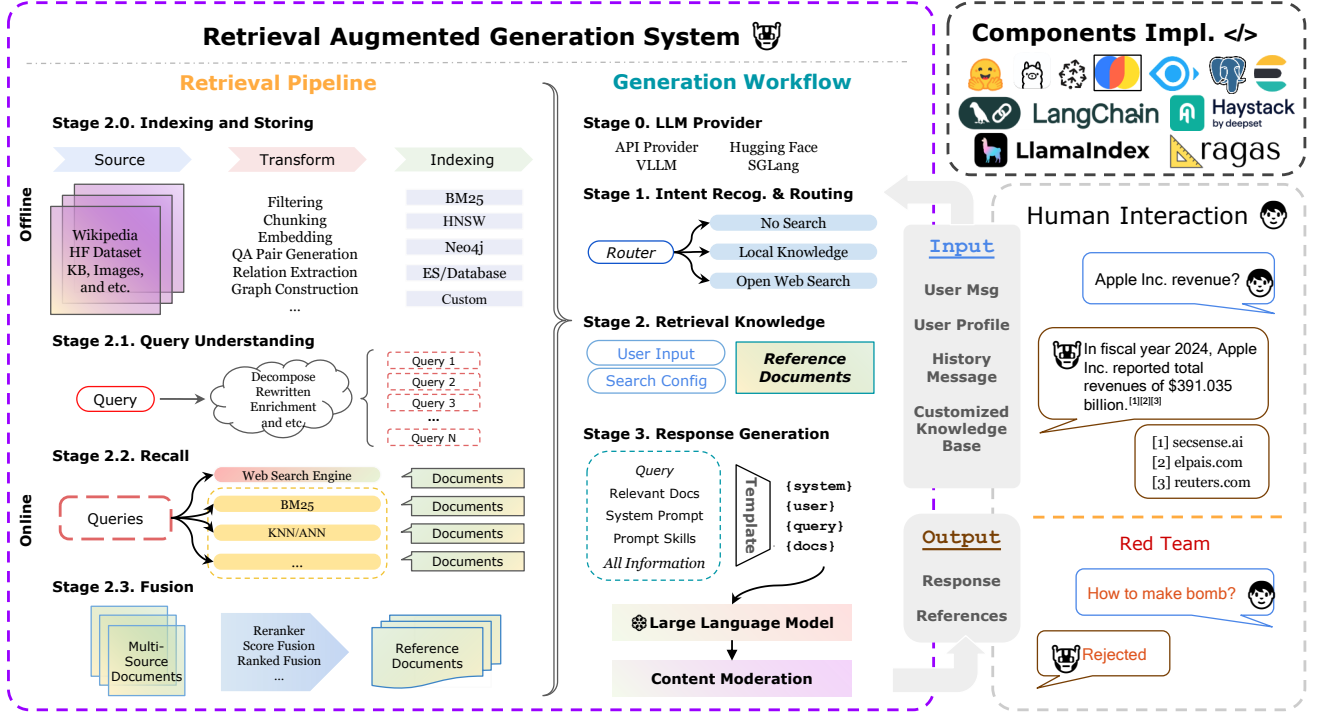


Fig. 1 The workflow of the RAG system and component implementation in the LLM era.

tent in a high-dimensional, dense semantic space for efficient retrieval computation [5, 35].

Typically, RAG assessments convert the task into a conversational format of Question Answering (QA) comprising question and the ground-true answers with doc candidates [36, 37]. In the online RAG workflow, some additional components are introduced before the retrieval, such like intent recognition, query rewriting and routing [38]. The retriever then indexes document collections from the data source. In this core step, multiple retrieval strategies can be employed, including sparse retrieval, dense retrieval, graph retrieval or hybrid methods [6, 39]. Certain systems conduct additional dynamic searches through search engines, typically found in commercialized products. Some systems may introduce an extra post-retrieval step to rerank the documents or fuse the data across different sources [7, 40]. In the generation pipeline, the responding progress based on the relevant documents is assigned to the LLM, which serves primarily as a decision-maker or reasoner [8]. Instead of generating knowledge independently, the LLM synthesizes retrieved information to form coherent responses, thereby reducing the risk of internal hallucination. Additionally, a range of methods of prompt engineering are available, including CoT [18], ToT [41], Self-Note [42] and RaR [43], etc. Depending on the specific task and expected output, a post-processing step may be required after the knowledge-oriented response, such as Entity Recognition for multi-choice questions or classification task, and the translation component for multilingual task. Moreover, the utility of the model’s application is a point of concern, particularly regarding safety and efficiency [44].

### 2.3 Related Surveys

Li et al. [23] summarized and formalized the key definitions of RAG while providing a synthesis of early-stage methodologies and practical applications. Expanding the scope beyond NLP, Zhao et al. [45] traced the developmental trajectory of multimodal RAG across the broader AIGC landscape. The emergence of LLM has since triggered an accelerated development of RAG methods, with numerous survey papers emerging to document this growing research domain [1, 9, 19, 20, 46]. Current researches mainly focus on collecting methods or applications, but lack substantive discussion about systematic evaluation mechanisms. While Yu et al. [21] provided an initial review outlining conceptual approaches for RAG evaluation, their analysis was predominantly confined to mainstream the frameworks, offering limited insights into emerging assessment methods applicable to diverse contexts. Building upon previous foundational work, this comprehensive survey extends beyond these limitations, offering deeper insights into emerging evaluation methods.

This study extends the research [21] by incorporating a broader array of RAG evaluation methods within a systems theory context. We differentiate between internal and external evaluations: the former examines the RAG component assessments and their interactive processes within the system architecture, while the latter focuses on holistic system evaluation and environmental considerations, where environment specifically denotes the external tasks or particular evaluation contexts. We extend our horizons beyond collecting conceptual definitions of evaluation methods to exploring and ana-

lyzing their practical application in the actual RAG studies. Simultaneously, we focus on RAG evaluation in LLM contexts, prioritizing unstructured text retrieval as the prevailing paradigm. Domain-specific variants of RAG evaluation (e.g., knowledge graph, multimodal retrieval) are excluded due to fundamental architectural gaps. Unless otherwise indicated, all the ‘RAG’ hereafter pertain to the narrow operational training-free framework employing unstructured documents as external knowledge resources.

### 3 Internal Evaluation

In this section, we summarize and organize the evaluations of the internal components with their interactions within a RAG system from prior studies. We deconstruct the evaluation of a whole RAG system, focusing on internal component interactions. A range of evaluation approaches are then introduced, from traditional to new ones. The elements mentioned and the implication of internal evaluation point to a framework for *evaluating the strengths of the RAG system’s core functionality*, that is, generating accurate and credible output.

#### 3.1 Evaluation Target

The diverse components of the RAG system can be boiled down to solving two core problems: the retrieval of the ground truth, and the generation of the response that closely aligns with the gold answer. They correspond to the respective evaluation objectives of the retrieval and generation modules.

Figure 2 summarizes the evaluation targets of the retrieval and generation component. The retrieval component includes two main stages, recall and ranking. The outputs, relevant documents, for both are similar to evaluate. Then we can construct several pairwise relationships for the **retrieval** component by defining the target as follows:

**Relevance** (*Relevant Documents*  $\leftrightarrow$  *Query*) evaluates how well the retrieved documents match the information needed expressed in the query. It measures the precision and specificity of the retrieval process.

**Comprehensiveness** (*Relevant Documents*  $\leftrightarrow$  *Relevant Documents*) evaluates the diversity and coverage of the retrieved documents. This metric assesses how well the system captures a wide range of relevant information, ensuring that the retrieved documents provide a comprehensive view of the topic according to the query.

**Correctness** (*Relevant Documents*  $\leftrightarrow$  *Documents Candidates*) assesses how accurate the retrieved documents are in comparison to a set of candidate documents. It is a measure of the system’s ability to identify and score relevant documents higher than less relevant or irrelevant ones.

The similar pairwise relations and targets for the **generation** component are outlined below.

**Relevance** (*Response*  $\leftrightarrow$  *Query*) measures how well the generated response aligns with the intent and content of the initial query. It ensures that the response is related to the query topic and meets the query’s specific requirements.

**Faithfulness** (*Response*  $\leftrightarrow$  *Relevant Documents*) evaluates how the generated response accurately reflects the information contained in the relevant documents and measures the consistency between the generated and source documents.

**Correctness** (*Response*  $\leftrightarrow$  *Sample Response*) Similar to the accuracy in the retrieval component, this measures the accuracy of the generated response against a sample response, which serves as a ground truth. It checks if the response is correct in terms of factual information and appropriate in the context of the query.

#### 3.2 Conventional Evaluation Methods

RAG is a cross-disciplinary system founded on traditional research fields including information retrieval (IR) and natural language generation (NLG). Adhering to the conventional methods of them, numerous traditional metrics are employed to evaluate the retrieval and generation of RAG as follows.

##### 3.2.1 IR-related Metrics

The IR-related metrics refer to the indicators associated with conventional retrieval systems. These metrics are categorized into two groups based on their correlation to ranking:

###### • Non-rank-based Metrics

The non-rank-based metrics typically evaluate binary outcomes, that is, whether an item is relevant or not, without taking into account the item’s position in a ranked list.

**Accuracy/Hit@K** is the proportion of true results (both true positives and true negatives) among the cases examined.

$$\text{Accuracy} = \frac{TP + TN}{\text{TotalNumber}}$$

where  $TP$  is the number of true positives,  $TN$  is the number of true negatives in the response.

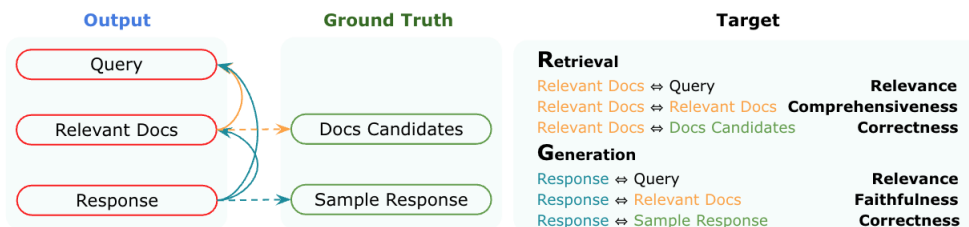


Fig. 2 The evaluation target of the Retrieval and Generation component in RAG.



**Recall@K** is the portion of relevant instances that have been retrieved over the total amount of relevant cases, considering only the top  $k$  results.

$$\text{Recall} = \frac{|RD \cap Top_{kd}|}{|RD|}$$

where  $RD$  is the relevant documents, and  $Top_{kd}$  is the top- $k$  retrieved documents.

**Precision@K** is the fraction of relevant instances among the retrieved instances, considering only the top  $k$  results.

$$\text{Precision} = \frac{TP}{TP + FP}$$

where  $TP$  represents true positives and  $FP$  represents false positives, respectively.

**F1 Score** measures the balance between precision and recall, defined as the Harmonic Mean of the two.

$$F1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

#### • Rank-Based Metrics

The rank-based metrics focus on the sequential presentation of relevant items, assigning greater significance to the positioning of these items within the ranking list.

**MRR** (Mean Reciprocal Rank) is the average of the reciprocal ranks of the first correct answer for a set of queries.

$$MRR = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \frac{1}{rank_i}$$

where  $|Q|$  is the number of queries and  $rank_i$  is the rank position of the first relevant document for the  $i$ -th query.

**NDCG** (Normalized Discounted Cumulative Gain) accounts for the position of the relevant documents by penalizing relevant documents that appear lower in the search results [47].

$$NDCG@k = \frac{DCG@k}{IDCG@k}$$

where  $DCG@k$  is the Discounted Cumulative Gain at rank  $k$  and  $IDCG@k$  is the Ideal Discounted Cumulative Gain at rank  $k$ , which represents the maximum possible  $DCG@k$ .  $DCG@k$  is defined as:

$$DCG@k = \sum_{i=1}^k \frac{2^{rel_i} - 1}{\log_2(i + 1)}$$

with  $rel_i$  being the graded relevance of the result at position  $i$ .

**MAP** (Mean Average Precision) is the mean of the average precision scores for each query.

$$MAP = \frac{1}{|Q|} \sum_{q=1}^{|Q|} \frac{\sum_{k=1}^n (P(k) \times rel(k))}{|relevant documents_q|}$$

where  $P(k)$  is the precision at cutoff  $k$  in the list,  $rel(k)$  is an indicator function equaling 1 if the item at rank  $k$  is a relevant document in the  $n$  retrieved documents, 0 otherwise.

#### 3.2.2 NLG-related Metrics

The NLG-related metrics focus on the content of the text output, dedicated to the evaluation on the char or semantic level.

**EM** (Exact Match) is a simple, stringent and widely-used evaluation metric that assesses the accuracy of model-generated answers compared to the ground truth. It scores as 1 if a generated answer precisely aligns with the standard otherwise 0. Typically, the responses need standardization and preprocessing (e.g., conversion to lowercase, removal of punctuation, elimination of articles, and standardization of number formats) before comparison. A general approach involves combining EM and Precision / Recall / F1 or edit distance [48,49].

**ROUGE** (Recall-Oriented Understudy for Gisting Evaluation) [50] is a set of metrics designed to evaluate the quality of summaries by comparing them to human-generated reference summaries. ROUGE can be indicative of the content overlap between the generated text and the reference text. The variants of ROUGEs measure the overlap of n-grams (ROUGE-N, ROUGGE-W), word subsequences (ROUGE-L, ROUGGE-S), and word pairs between the system-generated summary and the reference summaries.

**BLEU** (Bilingual Evaluation Understudy) [51] is a metric for evaluating the quality of machine-translated text against one or more reference translations. BLEU calculates the precision of n-grams in the generated text compared to the reference text and then applies a brevity penalty to discourage overly short translations. Beyond machine translation evaluation, BLEU can also be used for supervised comparison evaluation for general natural language generation. BLEU has limitations, such as not accounting for the fluency or grammaticality of the generated text.

**METEOR** [52] is a metric designed to assess the quality of machine translation or text generation. It enhances BLEU by incorporating mechanisms like synonymization, stemming matching, and word order penalties, demonstrating a stronger correlation with results obtained from manual evaluations. METEOR is defined as:

$$\text{METEOR} = (1 - p) \frac{(\alpha^2 + 1) \text{Precision} \times \text{Recall}}{\text{Recall} + \alpha \text{Precision}},$$

where  $\alpha$  is the balanced factor, and  $p$  is the penalization factor for word order.

**BertScore** [53] leverages the contextual embedding from pre-trained transformers like BERT to evaluate the semantic similarity between generated text and reference text. BertScore computes token-level similarity using contextual embedding and produces precision, recall, and F1 scores. Unlike n-gram-based metrics, BertScore captures the meaning of words in context, making it more robust to paraphrasing and more sensitive to semantic equivalence. It has multiple variants, including backbone advanced pre-trained models (e.g. BERT, RoBERTa and BART) and supervised evaluation based on external classifier design.

**Textual Similarity** measures the semantic variety in retrieved documents. It can be calculated using metrics like

*Intra-Document Similarity* or *Inter-Document Similarity*, which assess the similarity between documents within a set.

$$\text{Similarity} = \frac{1}{|D|^2} \sum_{i=1}^{|D|} \sum_{j=1}^{|D|} \text{sim}(d_i, d_j)$$

where  $D$  is the set of retrieved documents,  $d_i$  and  $d_j$  are embeddings of individual documents, and  $\text{sim}(d_i, d_j)$  is a similarity measure (e.g., the most commonly used cosine similarity) between the two documents.

**Coverage** measures the proportion of relevant documents retrieved from the total number of relevant documents available in the dataset. It quantifies how comprehensively the system captures all pertinent information across the corpus, across topics, categories, or entities defined by humans or in the knowledge base.

$$\text{Coverage} = \frac{|RD \cap \text{Retrieved}|}{|RD|}$$

where  $RD$  is the set of relevant documents and the notation *Retrieved* is the set of retrieved documents. The coverage can also be calculated at the group level, where the relevant documents are grouped into different categories or topics.

$$\text{Coverage} = \frac{|\text{Relevant Groups} \cap \text{Retrieved Groups}|}{|\text{Relevant Groups}|}$$

**Perplexity** (PPL) gauges a language model's predictive prowess, illustrating its level of uncertainty concerning test data. Essentially, it is an exponential variation of cross-entropy, quantifying the model's fit to the probability distribution of the text. It is defined base on the generative LM output as

$$\text{Perplexity} = \exp \left( -\frac{1}{N} \sum_{i=1}^N \log p(w_i | w_1, w_2, \dots, w_{i-1}) \right).$$

It's important to note that the IR-related and NLG-related methods are not directly equivalent to retrieval and generation assessment methods. In RAG systems, retrieval and generation operations typically alternate. For instance, the query understanding and document fusion component are considered as pre- and post-retrieval operations in the retriever, respectively, yet the evaluation is sometimes based on the NLG-like methods. SCARF [54] used BLEU / ROUGE to evaluate the query relevance of the retriever. Blagojevic et al. [40] utilized cosine similarity to assess the retrieval diversity. Additionally, the metrics can be adapted into various designs with new label based on the specific subject of study, such as EditDist [55], Fresheval [56], etc.

### 3.2.3 Upstream Evaluation

Given the rapid advancement of RAG systems, it is crucial to emphasize the significance of offline preprocessing of the corpus. We supplement the evaluation method of preprocessing modules, including chunking and embedding.

The evaluation of chunking methods can be conducted at two levels. First, chunk-specific evaluation focuses on intrinsic metrics such as *Accuracy*, measured by Full Keyword Coverage—the percentage of required keywords present in at least one retrieved chunk—and the *Tokens To Answer* metric, which tracks the index of the first fully comprehensive chunk and cumulative token count needed for full context coverage [57]. Second, extrinsic evaluation analyzes how different chunking approaches influence retrieval performance on downstream tasks. For example, [34] and [58] evaluate chunking methods by comparing retrieval recall, precision, and response quality using metrics like ROUGE, BLEU, and F1 scores against ground truth evidence paragraphs, while considering computational overhead. Other works extend this evaluation using domain-specific datasets, such as financial reports [57], to observe how structure-based and semantic chunking improves retrieval accuracy while reducing latency and token usage during inference.

Before retrieval, the embedding model determines the actual performance of retrieving relevant documents. Comprehensive benchmarks like Massive Text Embedding Benchmark (MTEB) [59] and Massive Multicultural Text Embedding Benchmark (MMTEB) [60] have become standard for the evaluation of embedding models. MTEB introduced the first large-scale benchmark covering 8 embedding tasks across 58 datasets and 112 languages, establishing that no single embedding method excels across all tasks. MMTEB significantly expanded this work through a community-driven effort, encompassing over 500 evaluation tasks across 250+ languages and introducing novel challenges like instruction following, long-document retrieval, and code retrieval.

Although the models of chunking and embedding have broad applications, they primarily serve as an upstream component of the retriever in RAG. The primary benefit to the entire system, involving chunking and embedding, is reflected in the enhancement of the retriever's evaluation metrics.

## 3.3 Evaluation Methods via LLMs

The advancement of LLM has catalyzed refined investigations into RAG system architectures. Contemporary studies increasingly employ LLM-driven assessment metrics, which establish quantifiable benchmarks for iterative improvements across different RAG modules. They can be broadly categorized into the output and representation based methods.

### 3.3.1 LLM Output based Methods

The LLM-output based evaluation methods perform content identification or statistical analysis of the text-format output of the RAG components assumed by the LLM. These methods feature a concise and easily understandable process without restrictions regarding whether the LLM is open or closed.

The most straightforward approach is to instruct the LLM to explicitly evaluate or score the textual output of the component by prompt engineering. Methods like RAGAS [61] and

Databricks Eval [62] prompt GPT-based judges with explicit instructions, such as “Check if the response is supported by the retrieved context.” or “Assess completeness with respect to the user query.” Zhang et al. [63] utilized GPT-4 with a few-shot prompt design to determine whether the generated answer matches the gold ones comprehensively. Finsås et al. [64] implemented a multi-agent LLM framework to evaluate the retrieval performance and reported a higher relevance with the human preference than the traditional methods. Patil et al. [65] proposed an Abstract Syntax Tree (AST) based method to measure the hallucination in RAG, which indicates the accuracy of calling external APIs in the RAG system. These methods typically benefit from CoT reasoning.

In addition, numerous researchers have proposed novel definitions of statistical metrics derived from the LLM output, facilitating a multi-perspective approach to evaluating the RAG components.

Dai et al. [66] proposed a new metric Semantic Perplexity (*SePer*) to capture the LLM’s internal belief about the correctness of the generated answer. Given the query  $q$  and the reference answers  $a^*$ , *SePer* is defined as the output sequence likelihood with clustered entity target as:

$$SePer_M(q, a^*) = P_M(a^* | q) \approx \sum_{C_i \in \mathcal{C}} k(C_i, a^*) p_M(C_i | q),$$

where  $M$  is the specific LLM.  $\mathcal{C}$  is the cluster set that the another clustering model groups the responses into.  $p_M(C_i | q)$  means the probability that a response generated by  $M$  is mapped to the cluster  $C_i$ .  $k(C_i, a^*)$  is a simple kernel function to measure the distance between the meaning of semantic cluster  $C_i$  and  $a^*$  by utilizing char-level matching or simply asking the LLM to get a True / False response.

Qi et al. [67] introduced the key point extraction to the RAG evaluation and designed *KPR* metric to evaluate the extent to which LLMs incorporate key points extracted from the retrieved documents into their generated responses:

$$KPR(\cdot) = \frac{1}{|Q|} \sum_{q \in Q} \frac{\sum_{x \in \mathbf{x}^q} I(x, \mathcal{M}(q||d^q))}{|\mathbf{x}^q|},$$

where  $Q$  is the global query set, and  $I(x, \mathcal{M}(q||d^q))$  is a function to judge whether a single LLM output sequence  $\mathcal{M}(q||d^q)$  based on the query  $q$  and the recalled documents  $d^q$  entails the predefined key points  $\mathbf{x}^q$ .

To evaluate the inconsistency of the different retrievers in RAG, Li et al. [68] proposed a pair of naive metrics called Mean Relative Win/Lose Ratio (MRWR/MRLR). Given  $M$  different retrievers  $\mathcal{R} = \{r_1, r_2, \dots, r_M\}$  and the dataset with  $N$  query & answer pairs, the correctness of model response for each sample  $\langle q_n, a_n \rangle$  is first calculated, denoted by  $\mathbf{I}^m(n) = 1$  if the retriever  $r_m$  answers correctly on sample  $s_n$  otherwise 0. Then the Relative Win Ratio (RWR) of retriever  $r_i$  over another retriever  $r_j$  is defined as:

$$RWR(i, j) = \frac{\sum_{n=1}^N \mathbf{I}^i(n) * (1 - \mathbf{I}^j(n))}{\sum_{n=1}^N 1 - \mathbf{I}^j(n)},$$

which represents the proportion of questions answered incorrectly by retriever  $r_j$  that were correctly answered by retriever  $r_i$ . MRWR and MRLR are calculated by respectively averaging RWR across rows and columns among the retrievers:

$$MRWR(i) = \frac{1}{M-1} \sum_{j \neq i} RWR(i, j),$$

$$MRLR(i) = \frac{1}{M-1} \sum_{j \neq i} RWR(j, i).$$

Especially,  $MRLR(i) = 0$  implies that retriever  $r_i$  consistently outperforms all of the other ones.

Min et al. [69] proposed *FactScore* to measure whether the generated content matches the given knowledge source by breaking the generations into atomic facts. Chiang et al. [70] further considered the synonym expression and proposed the advanced *D-FAatScore*. *FactScore* is a simple statistical determination whether the factual content  $a$  in the generated text  $y$  matches the external knowledge base  $\mathcal{C}$ :

$$FS(y) = \frac{1}{|\mathcal{A}_y|} \sum_{a \in \mathcal{A}_y} \mathbb{I}[a \text{ is supported by } \mathcal{C}].$$

*D-FAatScore* links synonymous entities into the same cluster  $\mathcal{A}_{y_i}$  and consider a cluster-level evaluation:

$$DFS(y) = \frac{1}{|\mathcal{A}_y|} \sum_{\mathcal{A}_{y_i} \in \mathcal{A}_y} \sum_{a \in \mathcal{A}_{y_i}} \mathbb{I}[a \text{ is supported by } \mathcal{C}_i^*].$$

To evaluate the risk in the generator’s response, Chen et al. [71] introduced the divided cases of the generated answer, answerable(A) and unanswerable(U), along with the different prediction process in the RAG system, keep(K) and discard(D). Four risk-aware evaluation metrics from various aspects are defined as:

1) *Risk* that measures the proportion of risky cases among the kept samples:

$$Risk = \frac{|\text{UK}|}{|\text{AK}| + |\text{UK}|}$$

2) *Carefulness* indicates the percentage of incorrect and discarded samples that are equivalent to recall for the unanswerable samples:

$$Carefulness = \frac{|\text{UD}|}{|\text{UK}| + |\text{UD}|}$$

3) *Alignment* refers to the proportion of samples in which the system’s judgment align with the assigned labels:

$$Alignment = \frac{|\text{AK}| + |\text{UD}|}{|\text{AK}| + |\text{AD}| + |\text{UK}| + |\text{UD}|}$$

4) *Coverage* quantifies the proportion of samples retained:

$$Coverage = \frac{|\text{AK}| + |\text{UK}|}{|\text{AK}| + |\text{AD}| + |\text{UK}| + |\text{UD}|}$$

### 3.3.2 LLM Representation based Methods

The representation-based methods, conversely, captures valuable metrics by modeling vector representation in the intermediate or final layers of the LLM. These methods can mitigate overreliance on surface lexical patterns, but they may lose interpretability since the final numeric similarity does not necessarily clarify which factual detail is correct or not.

Certain methods are inspired by the conventional metrics, demonstrated as expansions of existing metrics on the LLM. For instance, GPTScore [72] is a GPT based LLM-scoring method inspired by BertScore, which has been widely used as a convincing metric. ARES [73] combined a classifier with LLM embeddings to check whether a generative answer is semantically aligned with ground-truth evidence. RAGAS [61] uses a cosine similarity approach on LLM-generated embeddings to gauge answer relevance.

Moreover, numerous researchers have developed novel representation based metrics, which serve not only to evaluate the components but also to guide the further enhancement.

Zhao et al. [74] introduced a novel metric, *Thrust*, which assesses the LLM’s knowledgeability by leveraging the representation distribution of the instances produced by the LLM. A hypothesis was proposed that if an LLM has acquired adequate knowledge pertaining to a task, it should effectively cluster samples related to that task through its hidden states. The *Thrust* metric was defined as:

$$s_{\text{thrust}}(q) = \left\| \frac{1}{N \cdot K} \sum_{l=1}^N \sum_{k=1}^K \frac{|C_{kl}|}{\|d_{kl}(q)\|^2} \cdot \frac{d_{kl}(q)}{\|d_{kl}(q)\|} \right\|,$$

where  $N$  is the number of classes for the specific task,  $K$  is the number of clusters per class,  $|C_{kl}|$  denotes the cardinality of the set.  $d_{kl}(q)$  is a vector pointing from the representation of the query to the centroid.

Zhu et al. [75] introduced the information bottleneck theory into retrieval component to measure the relevance of the recalled document and candidate document. Moreover, a new information bottleneck-based loss function was derived and used to train a better noise filter for the retriever. Given the sample  $\{q, x, y\}$  from the dataset and the noise filter  $p(\tilde{x}|x, q)$  (need tuning), the information bottleneck in the RAG task is derived and formulated as:

$$\text{IB}(\tilde{x}) = P_{\text{LLM}}(x|[q, \tilde{x}, y]) - \alpha P_{\text{LLM}}(y|[q, \tilde{x}]),$$

where  $[\cdot]$  means the concatenation operation.  $P_{\text{LLM}}$  means the final output probability of the LLM.

Li et al. [76] proposed a new metric *GECE* based on METEOR for assessing the extent of the long-tailness of the generated text in RAG:

$$\text{GECE} = \frac{|\text{METEOR}(\text{pred}, \text{ref}) - \frac{1}{n} \sum_{i=1}^n P_{\text{LLM}}(t_i)|}{\alpha \cdot [E(\nabla_{\text{ins}}) \cdot \nabla_{\text{ins}}]},$$

where  $\alpha$  is the average word frequency,  $\nabla_{\text{ins}}$  and  $E(\nabla_{\text{ins}})$  are the gradient w.r.t. the current instance and the mean gradient

of the total dataset, separately. A long-tail instance usually has a smaller  $\alpha$  and  $\nabla_{\text{ins}}$ , obtaining a larger *GECE*, which implies larger degree of long-tailness.

To assess the extent to which external knowledge is utilized in the RAG response, Sun et al. [77] proposed External Context Score  $\mathcal{E}$ , which is defined on the response level as:

$$\mathcal{E}_{\mathbf{r}}^{l,h} = \frac{1}{|\mathbf{r}|} \sum_{t \in \mathbf{r}} \mathcal{E}_t^{l,h} = \frac{1}{|\mathbf{r}|} \sum_{t \in \mathbf{r}} \frac{\mathbf{e} \cdot \mathbf{x}_t^L}{\|\mathbf{e}\| \|\mathbf{x}_t^L\|},$$

where  $|\mathbf{r}|$  means the length of the response  $\mathbf{r}$ ,  $\mathbf{x}_t^L$  is the  $t$ -th token’s vector logit of the last layer  $L$ .  $\mathbf{e}$  is a pooled vector of the most relevant vectors of  $\mathbf{x}_t^L$  according to the attention weights in the middle layer:

$$\mathbf{e} = \frac{1}{|\mathcal{I}_t^{l,h}|} \sum_{j \in \mathcal{I}_t^{l,h}} \mathbf{x}_j^L,$$

where  $\mathcal{I}_t^{l,h}$  means the attended times where the token has larger than top-k% attention scores with  $\mathbf{x}_t^L$  in the  $l$ -th layer.

Noted that some of these LLM based evaluation metrics represent research specializations. While they may not be directly targeted towards an actual RAG system, their presentation is an integral part of advancing researches in the field of RAG, indicating significant contributions as well.

## 4 External Evaluation

We have dissected the components of RAG and provided a comprehensive account of its internal evaluation. This section shifts our focus to *the external utility that RAG, as a complete system*, encounters. We summarize the external utility in two areas: safety and efficiency, the evaluation of whom are introduced below.

### 4.1 Safety Evaluation

Safety pertains to the RAG system’s capacity to ensure the generation of stable and harmless content within a dynamic, even noisy or hazardous environment. As RAG systems continue widespread deployment, safety concerns have intensified beyond those of standalone LLMs. The incorporation of external knowledge sources introduces unique vulnerabilities requiring specialized evaluation frameworks [20].

**Robustness** evaluations focus on system behavior when processing misleading information in retrieval results. The RECALL benchmark [78] tests discrimination between reliable and counterfactual knowledge using BLEU, ROUGE-L, and specialized metrics like Misleading Rate. Wu et al. [79] quantify susceptibility to semantically related but irrelevant information using Misrepresentation Ratio and Uncertainty Ratio. SafeRAG [80] categorizes challenges like “inter-context conflict” with specific evaluation metrics, while C-RAG [81] provides theoretical guarantees on generation risks using conformal risk analysis and ROUGE-L. Cheng et al. [82] introduce two metrics to evaluate the RAG system: 1) *Resilience*



*Rate*, aiming to emphasize the system’s stability and robustness, quantifies the percentage of instances where the system’s responses remain accurate, both prior to and following retrieval augmentation. 2) *Boost Rate* quantifies the proportion of instances initially answered erroneously that were subsequently corrected upon the introduction of a retrieved document, evaluating the effectiveness of RAG.

**Factuality** focuses on generating accurate information and avoiding plausible but incorrect statements (hallucinations), especially with noisy or conflicting retrieval results [78, 83, 84]. Key metrics include *Factual Accuracy*, using standard QA metrics (EM, F1, accuracy, etc.) when the context might be misleading [78]; the *Hallucination Rate*, the frequency of generated information not supported by or contradicting retrieved documents, often measured via LLM-as-judge [85] or human evaluation; *Citation Accuracy*, assessing correct attribution to sources using *Citation Precision* and *Citation Recall* [20, 85]; and *Faithfulness Metrics*, evaluating how accurately the output reflects retrieved information [83].

**Adversarial attacks** target specific components within the RAG pipeline. Knowledge database poisoning (PoisonedRAG [86]) targets the retrieval corpus by injecting malicious texts that trigger predetermined outputs when retrieved. This attack vector is evaluated using Attack Success Rate (ASR) and retrieval-focused Precision/Recall/F1 metrics. Retrieval hijacking (HijackRAG [87]) exploits ranking algorithms to prioritize malicious content during retrieval, with evaluation focusing on attack transferability across models. Phantom attacks [88] use trigger-activated documents evaluated through Retrieval Failure Rate (Ret-FR), while jamming attacks [89] insert ‘blocker’ documents that force response refusal, assessed through oracle-based metrics.

**Privacy** assess information exposure risks from retrieval databases or user queries [90]. Evaluation often involves simulated attacks [91, 92]. Key metrics about privacy include the *Extraction Success Rate*, the frequency or success rate of attacks extracting specific private information (e.g., names, PII) from the knowledge base, often measured by the count of successfully extracted items [90]; the *PII Leakage Rate*, the amount or percentage of Personally Identifiable Information inadvertently revealed in generated outputs, typically found via pattern matching or inspection [93]; and the *Membership Inference Attack Success*, which measures an attacker’s ability to determine if a specific data record was in the RAG system’s knowledge base.

**Fairness** examines if the RAG system exhibits or amplifies biases from retrieved documents or training, leading to inequitable outputs [94]. *Bias Metrics* are used to analyze the outputs for disparities, which are quantitative measures of performance disparities (e.g., error rates, sentiment scores) across demographic groups [94]. *Stereotype Detection* measures the frequency or severity of harmful stereotypes in generated text, assessed via lists or human evaluation. *Counterfactual Fairness* checks if outputs change inappropriately when sensitive attributes in queries or context are altered.

**Transparency / Accountability** assesses the understandability and traceability of the RAG system’s reasoning process, enabling verification of sources and justification [95, 96]. Metrics are often qualitative or user-focused, such as *Explanation Quality*, based on human ratings of the clarity, completeness, and usefulness of explanations or provenance information [96]; *Traceability*, the ease of linking the final output back to specific source documents or passages; and *Citation Accuracy* (precision/recall) [20].

Comprehensive safety benchmarks standardize evaluation across multiple dimensions. SafeRAG [80] classifies attack tasks into four categories with tailored datasets. VERA framework [97] uses bootstrap sampling for confidence bounds on safety metrics, while DeepTeam’s red teaming approach [93] identifies vulnerabilities through systematic testing. In addition, current research indicates defense mechanisms remain insufficient against sophisticated attacks [86–88]. Evaluations reveal significant vulnerabilities in current RAG systems [87, 88], underscoring the need for robust benchmarks and metrics addressing the unique safety challenges arising from the retrieval-generation interplay. Further efforts are required to evaluate the safety of RAG.

## 4.2 Efficiency Evaluation

Efficiency is another crucial aspect of RAG’s utility, directly linked to the real-world significance of a system’s popularity, cost, and effectiveness.

**Latency** evaluation typically focuses on two critical metrics. Time to first token (TTFT) [98] measures the time taken by the system to produce its initial output token after receiving a query, which is particularly crucial for user experience as it directly impacts perceived responsiveness. This metric is especially important in interactive applications where immediate feedback maintains user engagement. Additionally, complete response time (total latency) measures the duration from query submission to the generation of the entire response. This encompasses retrieval time, processing time, and generation time for all tokens. Hofstatte et al. [99] proposed Single Query Latency that refers to the complete end-to-end time taken to process a single query, including both complete retrieval and generation phases.

**Resources and Money Cost** evaluation of RAG systems is another critical component for assessing the efficiency. Cost evaluation methodologies typically focus on quantifying both direct expenditures and efficiency metrics that impact overall system economics. The total cost of RAG systems can be categorized into several key components [126]:

- *Infrastructure Costs*: Computing local resources for embedding generation, vector database maintenance, and LLM inference for open models.
- *Token-based Expenses*: API charges for external LLM services based on input and output token usage.
- *Storage Costs*: Vector database hosting and maintenance expenses that scale with corpus size.

**Table 1** Overview of RAG benchmarks and their evaluation datasets. Source Domain indicates the data origin (e.g., real-time news, specialized corpora), and Special Points highlight unique or novel features (like domain-specific tasks, dynamic changes, or false-premise data).

Benchmark	Time	Dataset Name(s)	Source Domain	Special Points
RAGAS [61]	2023.09	WikiEval	Post-2022 Wikipedia	Manually labeled for faithfulness
FreshLLMs [56]	2023.11	FRESHQA	Real-time news/web queries	Dynamic QA with false-premise detection
RECALL [78]	2023.11	EventKG, UJ	Multilingual KGs, sci. terms	Edited/counterfactual context tests
ARES [73]	2023.11	NQ [100], HotpotQA [101], FEVER [102], WoW [103], MultiRC [104], ReCoRD [105]	KILT and SuperGLUE corpora	Re-uses classic QA sets, multi-domain
RGB [85]	2023.12	Custom corpus	Latest news articles	Emphasizes info integration, noise rejections
MultiHop-RAG [7]	2024.01	Generated corpus	Daily news segments via mediastack	Multi-hop cross-document queries
CRUD-RAG [106]	2024.02	Generated corpus, UHGEval	Chinese news, domain texts	Create/Read/Update/Delete tasks
MedRAG [107]	2024.02	MIRAGE	Medical QA corpora	Healthcare domain knowledge
FeB4RAG [108]	2024.02	FeB4RAG, BEIR [109]	Federated search tasks	Multi-domain, multi-engine retrieval
RAGBench [110]	2024.06	PubMedQA, CovidQA, HotpotQA, MS Marco, CUAD, DelucionQA, EManual, TechQA, FinQA, TAT-QA	Multi-domain corpora	Faithfulness with TRACe (Util, Rel, Adh, Compl)
ReEval [111]	2024.05	NQ (MRQA) + RealTimeQA	Wikipedia, real-time QA	Adversarial test cases for hallucination detection
DomainRAG [112]	2024.06	Generated admission QA	College docs with yearly updates	Single-/multi-doc, single-/multi-turn QA
Telecom RAG Eval. [113]	2024.07	TeleQuAD	3GPP-based domain docs	Triple-labeled QA from SMEs (telecom context)
LegalBench-RAG [114]	2024.08	PrivacyQA, CUAD, MAUD, ContractNLI	Expert-annotated legal corpora	Emphasizes strict retrieval of legal text
RAGEval [115]	2024.08	DragonBall	Finance, law, medical docs	Schema-based generation, scenario-specific
CoRAGE [116]	2024.09	RealTimeQA [117], NQ [100]	Online QA + KILT tasks	Hallucination resilience, dynamic updates
RAG Unfairness [118]	2024.09	TREC22 FairRank, BBQ	Wikipedia-based track + socioecon. QA	Fairness metrics, group disparity
CoFE-RAG [119]	2024.10	CoFE data	PDF, DOC, multi-lingual docs	Fine-grained chunking, multi-keyword approach
OCR Hinders RAG [55]	2024.12	1,261 PDFs + 8,561 images	OCR text from scanned docs	Evaluates noise from OCR errors
OmniEval [120]	2024.12	Finance domain set	Financial docs, numeric tasks	Emphasizes numeric correctness/factual QA
CRAG [121]	2024.12	KG + web corpus	Knowledge graphs + web pages	Multi-entity queries, curated dynamic facts
RAG Playground [122]	2024.12	319 QA pairs	Curated multi-domain tasks	Prompt engineering / user flows
MTRAG [123]	2025.01	CLAPNQ, FiQA, Govt, Cloud	Wikipedia, finance, gov, tech docs	Multi-turn, bridging queries
CDQA [124]	2025.01	Chinese Dynamic QA	Recent Chinese news queries	Time-varying evolving answers
U-NIAH [125]	2025.03	Starlight Academy	Synthetic “needle-in-haystack” data	Evaluates extremely long contexts
SCARF [54]	2025.04	(User-provided)	Generic multi-domain	Modular or black-box approach integrates wide metrics (LLM judge)

- *Operational Overhead*: Human supervision, system maintenance, and regular updates to knowledge bases.
- *Development Costs*: Initial implementation, integration, and customization expenses.

For more details in the token-based expenses, LLM providers such as OpenAI and Google offer token usage metrics that track input and output token consumption during evaluation processes. This approach calculates costs by multiplying token counts by their respective pricing rates [127]. Researchers have developed metrics to evaluate the economic efficiency of RAG implementations:

- *Cost-Effectiveness Ratio*: Measures performance improvement per unit of cost, allowing for standardized comparison between different RAG configurations [127].
- *Retrieval Precision ROI*: Quantifies the economic return of improving retrieval precision by measuring the reduction in irrelevant context processing costs [127]. This metric demonstrated that optimizing retrieval can improve cost efficiency by up to around 50% through reducing token consumption during LLM inference.
- *User-Controllable Cost-Accuracy Tradeoffs*: Su et al. [128] propose evaluation methods using an interpretable control parameter ( $\alpha$ ) that allows systematic assessment of the relationship between retrieval costs and accuracy. This approach enables evaluating RAG systems

across a spectrum of cost constraints rather than at fixed operating points.

- *Comparative Cost Analysis*: Methodologies for evaluating relative cost efficiency between different RAG implementations for specific use cases, considering both direct costs and long-term economic sustainability [129].

## 5 Resources

The evaluation methodologies previously examined are comprehensive, though not necessarily abundant. This section systematically compiles, categorizes, and presents the implemented RAG evaluation frameworks, benchmarks, analytical tools, and datasets that have emerged in the large language model era. To our knowledge, this compilation constitutes the most exhaustive collection of RAG evaluation frameworks currently documented in the literature.

**Datasets.** We compiled the benchmarks along with the associated datasets in recent years. Early works focus on static general-purpose QA datasets (e.g., NQ [100], HotpotQA [101]), providing well-established baselines but lack recency or domain specificity. Recent benchmarks counter these limitations by 1) sourcing live news or rapidly updated online documents (e.g., RGB [85], MultiHop-RAG [7]) to test time-sensitive capabilities; 2) curating domain-specific corpora in

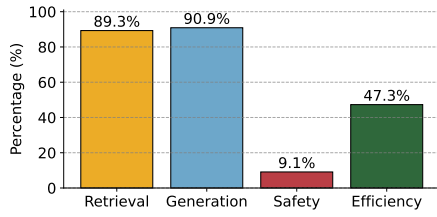
**Table 2** RAG evaluation frameworks, highlighting principal evaluation targets and methods. Retrieval focuses mainly on **Relevance (R)**, **Correctness (C)** or **Comprehensiveness**, whereas generation (right) focuses on **Faithfulness (F)**, **Correctness (C)**, or **Relevance (R)**. External evaluation targets (*safety*, *efficiency*) or other statements appear in italics.

Type	Framework	Time	Raw Targets	Retrieval Metrics	Generation Metrics
Research	FiD-Light [99]	2023.07	<i>Latency</i>	–	–
Research	Diversity Reranker [40]	2023.08	<i>Diversity</i>	<i>Cosine Distances</i>	–
Benchmark	RAGAS [61]	2023.09	Context R, Answer R, F	LLM as Judge	LLM CosSim, LLM as Judge
Tool	TruEra RAG Triad [130]	2023.10	Context R, Answer R, Groundedness	LLM as Judge	LLM as Judge
Tool	LangChain Bench. [131]	2023.11	C, F, ExecutionTime, EmbCosDist	Exact-match	LLM as Judge
Benchmark	FreshLLMs [56]	2023.11	Response C, Fast-changing, False premise	(retrieval logs)	STRICT / RELAXED, FRESHEVAL (LLM-based)
Tool	RECALL [78]	2023.11	Response Quality, Robustness	–	BLEU, ROUGE-L
Benchmark	ARES [73]	2023.11	Context R, Answer F, Answer R	LLM + Classifier	LLM + Classifier, LLM + Classifier
Benchmark	RGB [85]	2023.12	Info Integration, NoiseRobust, NegRejection, Counterfact	–	Accuracy
Tool	Databricks Eval [62]	2023.12	C, Readability, Comprehensiveness	–	LLM as Judge
Benchmark	MultiHop-RAG [7]	2024.01	Retrieval C, Response C	MAP, MRR, Hit@K	LLM as Judge
Benchmark	CRUD-RAG [106]	2024.02	Create, Read, Update, Delete	–	ROUGE, BLEU, RAGQuestEval
Benchmark	MedRAG [107]	2024.02	Accuracy (medical)	–	Exact-match, Acc.
Benchmark	FeB4RAG [108]	2024.02	Consistency, C, Clarity, Coverage	–	Human Eval, Human Eval
Benchmark	Arabic RAG Eval. [132]	2024.05	Doc R, Answer R	nDCG, MRR, mAP	Possibly CosSim to QA
Benchmark	RAGBench [110]	2024.06	Context R, Answer R, Explainability, TRACe = Util, Rel, Adh, Compl.	LLM-based Eval	LLM-based Eval, TRACe Metrics
Benchmark	ReEval [111]	2024.05	Hallucination Adversarial Attack	–	F1, EM, Entailment LLM or Human Eval
Benchmark	DomainRAG [112]	2024.06	C, F, NoiseRobust, StructOutput	–	F1, EM, ROUGE-L, LLM
Benchmark	CoURAGE [116]	2024.06	Hallucination	–	F1, EM, LLM as Judge, Human Eval
Tool	Telecom RAG Eval. [113]	2024.07	Context R, Faithfulness, Correctness	LLM-based Metrics	RAGAS-based, LLM Eval
Benchmark	LegalBench-RAG [114]	2024.08	Doc-level Precision, Citation Rel.	Precision, Recall	–
Benchmark	RAGEval [115]	2024.08	Completeness, Hallucination, Irrelevance	LLM-based Scoring	LLM-based, Human Alignment
Benchmark	RAG Unfairness [118]	2024.09	Fairness, C, C	MRR@K	EM, ROUGE
Benchmark	CoFE-RAG [119]	2024.10	Fine-grained Retrieval, Resp Quality, Diversity	Recall, Correctness, Multi-keyword	BLEU, ROUGE-L, LLM as Judge
Benchmark	Toward Instr.-Following [133]	2024.10	Instr. Relevance, Constraint	–	LLM as Judge, Atomic Pass Rate
Benchmark	OmniEval [120]	2024.12	Factual Acc., Domain Tasks	Rule+LLM	Manual or LLM FT
Benchmark	CRAG [121]	2024.12	Accuracy, Dynamism, Complex Facts, R, C	Weighted scoring	Accuracy, Truthfulness measure
Benchmark	OCR Hinders RAG [55]	2024.12	Accuracy, OCR Noise, Semantic vs. Format Noise	EditDist, LCS	F1-score
Benchmark	RAG Playground [122]	2024.12	Retrieval Strategy, Prompt Eng.	Comparison-based	LLM-based Eval
Benchmark	MTRAG [123]	2025.01	Multi-turn Quality, Conv. C	Recall, nDCG	LLM as Judge
Benchmark	CDQA [124]	2025.01	Accuracy	–	F1
Benchmark	U-NIAH [125]	2025.03	Needle Detect, LongContext, No Halluc.	Recall	LLM Judge, Heatmap
Tool	eRAG [134]	2024.04	Doc-level Rel., Downstream Quality	Doc-level LLM	Kendall’s $\tau$
Tool	SCARF [54]	2025.04	Context R, Answer R, Faithfulness	LLM-based or BLEU/ROUGE	RAGAS-like Relevance, LLM-based (Black-box Integration)


law, healthcare, or finance (e.g., MedRAG [107], OmniEval [120], LegalBench-RAG [114]); or 3) generating synthetic data or specialized QA pairs, possibly with false-premise or counterfactual elements (e.g., FreshLLMs [56], RAGEval [115]) to assess robustness and misinformation handling. We further provide a concise description of the original domains and characteristics according to the original resource, as shown in Table 1. Noted that only the datasets containing retrieved ground truth documents are included, indicating a concern for more in-depth system component evaluation.

**Frameworks with Evaluation Methods.** We compiled and summarized the evaluation methods devised by exist-

ing frameworks, as illustrated in Table 2. These efforts span from initial, point-level researches [40, 99] to later, multi-component evaluation tools and benchmarks [73, 131], encompassing a remarkably comprehensive collection of assessment frameworks. The evaluation methods employed are varied, encompassing both traditional [78, 132] and LLM-based metrics [106, 110]. Additionally, there are frameworks that facilitate safety-focused evaluations [85, 116], or are tailored to specific downstream domains like document [55, 125], telecom [113], medicine [107], etc. Referencing the component evaluation objectives outlined in section 3.1, we categorize and highlight the evaluation elements and specific metrics.



**Fig. 3** Statistics on the distribution of RAG studies across four key areas: retrieval, generation, safety, and efficiency. A paper may utilize evaluation methods in more than one areas.



**Fig. 4** Frequency statistics wordcloud of evaluation metrics in RAG studies. The LLM-based methods are categorized based on the targets and presented with the suffix ‘-LLM’. F-score refers to the expanded F1-score.

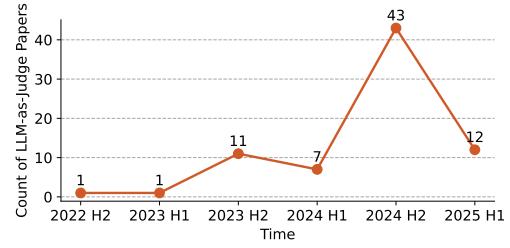
## 6 Discussion

### 6.1 Statistics and Analysis of RAG Evaluation

The proliferation of LLM has contributed to a significant diversification of RAG evaluation methods. Current researches, while demonstrating comprehensive coverage of RAG evaluation dimensions, often subjectively assert their respective utility statements. To assess the popularity of these evaluation methods, we conducted a statistical analysis of the available methods from a survey perspective. This can also be viewed as a research-oriented simple meta-evaluation. We crawled the collection of the papers since 2022 autumn with keywords about RAG in the accepted papers of the high-level conferences about NLP & AI, and extracted the component as well as the evaluation metrics the papers focus and utilize. We finally amassed a total of 582 PDF manuscripts. All the included papers have undergone rigorous peer review, demonstrating scholarly merit with complete experimental methodologies and logically structured evaluation procedures.

**Research Focus.** Figure 3 illustrates the statistical distribution of evaluation methods used across the four different segments in RAG studies (Retrieval / Generation / Safety / Efficiency). The data suggests a prevailing focus on internal research and evaluation of RAG systems, as indicated by the extensive coverage of the retrieval and generation processes. In contrast, external evaluations, particularly those related to safety, have garnered less attention.

**Metric Preference.** Word frequency counts were conducted for the assessment metrics mentioned in the papers, with the wordcloud displayed in Figure 4. Whenever a metric is formally introduced in the body of a paper or reported in the table of experimental results, its word frequency count is set +1. We manually merged and mapped synonymous met-



**Fig. 5** The number of papers explicitly mentioning LLM-based evaluation on RAG. The 2025 H1 collection is up to March 31st.

rics in the same session and excluded the words with global occurrences lower than twice. It is observed that traditional metrics predominantly dominate the evaluation usage, while LLM-based methods have not yet gained widespread acceptance among researchers. This phenomenon is attributed to the simplicity and reliability of the conventional metrics. Conversely, the LLM-based methods often require more effort and involve multiple settings that are difficult to keep the same across different researches, such as the LLM version and prompt design.

**Trend of LLM Usage.** Despite the potential issues with LLM-based methods, there is an observable trend of increasing application, as shown in Figure 5. 2024 H2 and 2025 H1 have the top two highest numbers. LLM judges are ultimately capable of handling more complex designs, drawing closer to real-world applications. LLM itself, additionally, has continued to evolve in recent years, with the performance progressively improving, and the supported functions expanding.

### 6.2 Challenges and Future Directions

This section addresses several challenges inherent in contemporary RAG evaluation.

**Limitations of LLM-based Methods.** The current evaluation design does not sufficiently address the timeliness and the black-box nature inherent in the LLM. The method of employing LLMs for assessments, particularly through direct prompts, raises latent risk about stability and security. Future research should focus on enhancing the robustness of the evaluation process itself and minimizing the likelihood of LLM errors in the RAG system.

**Cost of Evaluation.** The cost associated with the RAG system has garnered attention. Nevertheless, a thorough evaluation remains expensive due to the vast scale of the tools and datasets involved. Determining an efficient method for system evaluation, or striking a balance between cost and effectiveness, is one of the directions for future research.

**Advanced Evaluation Methods.** As LLMs continue to evolve, the components of RAG systems are becoming more diverse. Currently, many of these components are evaluated using end-to-end RAG ontology metrics, with a lack of comprehensive functional decomposition evaluation or theoretical analysis. Concurrently, there remains untapped potential in the functionalities of LLMs themselves. For instance, the



evaluation about deep thinking models (e.g. openai-o1 [135]) along with the thinking process of LLMs in conjunction with RAG's retrieval and generation process, is still inadequate. These in-depth evaluation strategies require further research and development in the future.

**Comprehensiveness of the Evaluation Framework.** Despite the abundant evaluation frameworks at present, individual ones are somewhat limited in their metrics and methods of evaluation. Moreover, most contemporary frameworks concentrate on widely used languages such as English and Chinese. There is an urgent need for frameworks that are not only methodologically but also linguistically diverse.

## 7 Conclusion

In this paper, we have presented the first comprehensive survey of RAG evaluation methodologies in the LLM era. Our systematic analysis reveals several important insights for researchers and practitioners working with these increasingly prevalent systems. For the evaluation of internal RAG performance, we dissect the internal components of RAG systems, define the assessment objectives, and gather a range of methods and metrics from traditional to innovative. Moreover, we investigate the external evaluation related to system integrity such as safety and efficiency, which are underexplored in RAG research according to our statistical analysis. Additionally, we compile and categorize the current evaluation datasets and frameworks to elucidate the unique attributes and assessment focuses of the resources. Last but not least, we analyze the implementation of existing evaluation methods and synthesize the challenges and future directions of RAG evaluation in the LLM era.

### Acknowledgements

**Competing interests** The authors declare that they have no competing interests or financial conflicts to disclose.

## References

1. Fan W, Ding Y, Ning L, Wang S, Li H, Yin D, Chua T S, Li Q. 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. 2024, 6491–6501
2. Gutiérrez B J, Shu Y, Gu Y, Yasunaga M, Su Y. Hipporag: Neurobiologically inspired long-term memory for large language models. arXiv preprint arXiv:2405.14831, 2024
3. Zhang Y, Khalifa M, Logeswaran L, Lee M, Lee H, Wang L. Merging Generated and Retrieved Knowledge for Open-Domain QA. In: Bouamor H, Pino J, Bali K, eds, Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. December 2023, 4710–4728
4. Yao J Y, Ning K P, Liu Z H, Ning M N, Yuan L. Llm lies: Hallucinations are not bugs, but features as adversarial examples. arXiv preprint arXiv:2310.01469, 2023
5. Wang L, Yang N, Huang X, Jiao B, Yang L, Jiang D, Majumder R, Wei F. Text embeddings by weakly-supervised contrastive pre-training. arXiv preprint arXiv:2212.03533, 2022
6. Robertson S, Zaragoza H, others. The probabilistic relevance framework: Bm25 and beyond. Foundations and Trends® in Information Retrieval, 2009, 3(4): 333–389
7. Tang Y, Yang Y. Multihop-rag: Benchmarking retrieval-augmented generation for multi-hop queries. arXiv preprint arXiv:2401.15391, 2024
8. Sun J, Xu C, Tang L, Wang S, Lin C, Gong Y, Shum H Y, Guo J. Think-on-graph: Deep and responsible reasoning of large language model with knowledge graph. CoRR, 2023
9. Gao Y, Xiong Y, Gao X, Jia K, Pan J, Bi Y, Dai Y, Sun J, Wang H, Wang H. Retrieval-augmented generation for large language models: A survey. arXiv preprint arXiv:2312.10997, 2023, 2
10. Brown T, Mann B, Ryder N, Subbiah M, Kaplan J D, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, others. Language models are few-shot learners. Advances in neural information processing systems, 2020, 33: 1877–1901
11. Zhao W X, Zhou K, Li J, Tang T, Wang X, Hou Y, Min Y, Zhang B, Zhang J, Dong Z, others. A survey of large language models. arXiv preprint arXiv:2303.18223, 2023
12. Yildirim I, Paul L. From task structures to world models: what do llms know? Trends in Cognitive Sciences, 2024
13. Zhang S, Dong L, Li X, Zhang S, Sun X, Wang S, Li J, Hu R, Zhang T, Wu F, others. Instruction tuning for large language models: A survey. arXiv preprint arXiv:2308.10792, 2023
14. Verma P, Pilanci M. Towards signal processing in large language models. arXiv preprint arXiv:2406.10254, 2024
15. Lyu H, Jiang S, Zeng H, Xia Y, Wang Q, Zhang S, Chen R, Leung C, Tang J, Luo J. Llm-rec: Personalized recommendation via prompting large language models. In: Findings of the Association for Computational Linguistics: NAACL 2024. 2024, 583–612
16. Zhang B, Liu Z, Cherry C, Firat O. When scaling meets llm finetuning: The effect of data, model and finetuning method. In: ICLR. 2024
17. Reynolds L, McDonell K. Prompt programming for large language models: Beyond the few-shot paradigm. In: Extended abstracts of the 2021 CHI conference on human factors in computing systems. 2021, 1–7
18. Wei J, Wang X, Schuurmans D, Bosma M, Xia F, Chi E, Le Q V, Zhou D, others. Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems, 2022, 35: 24824–24837
19. Huang Y, Huang J. A survey on retrieval-augmented text generation for large language models. arXiv preprint arXiv:2404.10981, 2024
20. Zhou Y, Liu Y, Li X, Jin J, Qian H, Liu Z, Li C, Dou Z, Ho T Y, Yu P S. Trustworthiness in retrieval-augmented generation systems: A survey. arXiv preprint arXiv:2409.10102, 2024

21. Yu H, Gan A, Zhang K, Tong S, Liu Q, Liu Z. Evaluation of retrieval-augmented generation: A survey. In: CCF Conference on Big Data. 2024, 102–120
22. Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, Küttler H, Lewis M, Yih W t, Rocktäschel T, others . Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 2020, 33: 9459–9474
23. Li H, Su Y, Cai D, Wang Y, Liu L. A survey on retrieval-augmented text generation. *arXiv preprint arXiv:2202.01110*, 2022
24. Dinan E, Roller S, Shuster K, Fan A, Auli M, Weston J. Wizard of wikipedia: Knowledge-powered conversational agents. *arXiv preprint arXiv:1811.01241*, 2018
25. Qin L, Galley M, Brockett C, Liu X, Gao X, Dolan W B, Choi Y, Gao J. Conversing by reading: Contentful neural conversation with on-demand machine reading. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 2019, 5427–5436
26. Kobayashi M, Takeda K. Information retrieval on the web. *ACM computing surveys (CSUR)*, 2000, 32(2): 144–173
27. Lee H, Yang S, Oh H, Seo M. Generative multi-hop retrieval. In: *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 2022, 1417–1436
28. Zhang S, Yao L, Sun A, Tay Y. Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys*, 2019, 52(1): 1–38
29. Wang W, Lin X, Feng F, He X, Chua T S. Generative recommendation: Towards next-generation recommender paradigm. *arXiv preprint arXiv:2304.03516*, 2023
30. Karpukhin V, Oguz B, Min S, Lewis P, Wu L, Edunov S, Chen D, Yih W t. Dense passage retrieval for open-domain question answering. In: Webber B, Cohn T, He Y, Liu Y, eds, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. November 2020, 6769–6781
31. Google . Programmable Search Engine | Google for Developers, 2024
32. Yepes A J, You Y, Milczek J, Laverde S, Li R. Financial report chunking for effective retrieval augmented generation. *arXiv preprint arXiv:2402.05131*, 2024
33. Fan W, Ding Y, Ning L, Wang S, Li H, Yin D, Chua T S, Li Q. 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*. 2024, 6491–6501
34. Singh I S, Aggarwal R, Allahverdiyev I, Taha M, Akalin A, Zhu K, O'Brien S. Chunkrag: Novel llm-chunk filtering method for rag systems. *arXiv preprint arXiv:2410.19572*, 2024
35. Multi-Granularity M L M F. M3-embedding: Multi-linguality, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. 2024
36. Mao Y, He P, Liu X, Shen Y, Gao J, Han J, Chen W. Generation-augmented retrieval for open-domain question answering. In: Zong C, Xia F, Li W, Navigli R, eds, *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. August 2021, 4089–4100
37. Mekala D, Vu T, Schick T, Shang J. Leveraging QA datasets to improve generative data augmentation. In: Goldberg Y, Kozareva Z, Zhang Y, eds, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. December 2022, 9737–9750
38. Asai A, Wu Z, Wang Y, Sil A, Hajishirzi H. Self-rag: Learning to retrieve, generate, and critique through self-reflection. In: *The Twelfth International Conference on Learning Representations*. 2023
39. Douze M, Guzhva A, Deng C, Johnson J, Szilvasy G, Mazaré P E, Lomeli M, Hosseini L, Jégou H. The faiss library. *CoRR*, 2024
40. Blagojevic V. Enhancing RAG Pipelines in Haystack: Introducing DiversityRanker and LostInTheMiddleRanker, August 2023
41. Besta M, Blach N, Kubicek A, Gerstenberger R, Podstawski M, Giniainazzi L, Gajda J, Lehmann T, Niewiadomski H, Nyczyk P, others . Graph of thoughts: Solving elaborate problems with large language models. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. 2024, 17682–17690
42. Lanchantin J, Toshniwal S, Weston J, Sukhbaatar S, others . Learning to reason and memorize with self-notes. *Advances in Neural Information Processing Systems*, 2023, 36: 11891–11911
43. Deng Y, Zhang W, Chen Z, Gu Q. Rephrase and respond: Let large language models ask better questions for themselves. *CoRR*, 2023
44. Wang C, Liu X, Yue Y, Tang X, Zhang T, Jiayang C, Yao Y, Gao W, Hu X, Qi Z, others . Survey on factuality in large language models: Knowledge, retrieval and domain-specificity. *arXiv preprint arXiv:2310.07521*, 2023
45. Zhao P, Zhang H, Yu Q, Wang Z, Geng Y, Fu F, Yang L, Zhang W, Cui B. Retrieval-augmented generation for ai-generated content: A survey. *CoRR*, 2024
46. Cheng M, Luo Y, Ouyang J, Liu Q, Liu H, Li L, Yu S, Zhang B, Cao J, Ma J, others . A survey on knowledge-oriented retrieval-augmented generation. *arXiv preprint arXiv:2503.10677*, 2025
47. Järvelin K, Kekäläinen J. Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems (TOIS)*, 2002, 20(4): 422–446
48. Sankoff D, Kruskal J B. Time warps, string edits, and macro-molecules: the theory and practice of sequence comparison. Reading: Addison-Wesley Publication, 1983
49. Yujian L, Bo L. A normalized levenshtein distance metric. *IEEE transactions on pattern analysis and machine intelligence*, 2007, 29(6): 1091–1095
50. Lin C Y. ROUGE: A package for automatic evaluation of summaries. In: *Text Summarization Branches Out*. July 2004, 74–81
51. Papineni K, Roukos S, Ward T, Zhu W J. Bleu: a method for automatic evaluation of machine translation. In: Isabelle P, Charniak E, Lin D, eds, *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*. July 2002, 311–318
52. Banerjee S, Lavie A. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In: *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*. 2005, 65–72

53. Zhang T, Kishore V, Wu F, Weinberger K Q, Artzi Y. BERTScore: Evaluating Text Generation with BERT. In: 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020. 2020
54. Rengo M, Beadini S, Alfano D, Abbruzzese R. A system for comprehensive assessment of rag frameworks. arXiv preprint arXiv:2504.07803, 2025
55. Zhang J, Zhang Q, Wang B, Ouyang L, Wen Z, Li Y, Chow K H, He C, Zhang W. Ocr hinders rag: Evaluating the cascading impact of ocr on retrieval-augmented generation. arXiv preprint arXiv:2412.02592, 2024
56. Vu T, Iyyer M, Wang X, Constant N, Wei J, Wei J, Tar C, Sung Y H, Zhou D, Le Q, Luong T. FreshLLMs: Refreshing large language models with search engine augmentation. In: Ku L W, Martins A, Srikumar V, eds, Findings of the Association for Computational Linguistics: ACL 2024. August 2024, 13697–13720
57. Sælemyr J, Femdal H T. Chunk smarter, retrieve better: Enhancing llms in finance: An empirical comparison of chunking techniques in retrieval augmented generation for financial reports. Master’s thesis, NORWEGIAN SCHOOL OF ECONOMICS, 2024
58. Finardi P, Avila L, Castaldoni R, Gengo P, Larcher C, Piau M, Costa P, Caridà V. The chronicles of rag: The retriever, the chunk and the generator. arXiv preprint arXiv:2401.07883, 2024
59. Muennighoff N, Tazi N, Magne L, Reimers N. Mteb: Massive text embedding benchmark. arXiv preprint arXiv:2210.07316, 2022
60. Enevoldsen K, Chung I, Kerboua I, Kardos M, Mathur A, Stap D, Gala J, Siblini W, Krzemiński D, Winata G I, others . Mmteb: Massive multilingual text embedding benchmark. arXiv preprint arXiv:2502.13595, 2025
61. Es S, James J, Anke L E, Schockaert S. Ragas: Automated evaluation of retrieval augmented generation. In: Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations. 2024, 150–158
62. Leng Q, Uhlenhuth K, Polyzotis A. Best practices for llm evaluation of rag applications (2023). URL <https://www.databricks.com/blog/LLM-auto-eval-best-practices-RAG>
63. Zhang H, Semnani S, Ghassemi F, Xu J, Liu S, Lam M. Spaghetti: Open-domain question answering from heterogeneous data sources with retrieval and semantic parsing. In: Findings of the Association for Computational Linguistics ACL 2024. 2024, 1663–1678
64. Finsås M, Maksim J. Optimizing rag systems for technical support with llm-based relevance feedback and multi-agent patterns. Master’s thesis, NTNU, 2024
65. Patil S G, Zhang T, Wang X, Gonzalez J E. Gorilla: Large language model connected with massive apis. Advances in Neural Information Processing Systems, 2024, 37: 126544–126565
66. Dai L, Xu Y, Ye J, Liu H, Xiong H. Seper: Measure retrieval utility through the lens of semantic perplexity reduction. arXiv preprint arXiv:2503.01478, 2025
67. Qi Z, Xu R, Guo Z, Wang C, Zhang H, Xu W. Long2rag: Evaluating long-context & long-form retrieval-augmented generation with key point recall. In: Findings of the Association for Computational Linguistics: EMNLP 2024. 2024, 4852–4872
68. Li M, Li X, Chen Y, Xuan W, Zhang W. Unraveling and mitigating retriever inconsistencies in retrieval-augmented large language models. In: Findings of the Association for Computational Linguistics ACL 2024. 2024, 4833–4850
69. Min S, Krishna K, Lyu X, Lewis M, Yih W t, Koh P, Iyyer M, Zettlemoyer L, Hajishirzi H. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation. In: Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. 2023, 12076–12100
70. Song Y, Kim Y, Iyyer M. Veriscore: Evaluating the factuality of verifiable claims in long-form text generation. arXiv preprint arXiv:2406.19276, 2024
71. Chen L, Zhang R, Guo J, Fan Y, Cheng X. Controlling risk of retrieval-augmented generation: A counterfactual prompting framework. In: Findings of the Association for Computational Linguistics: EMNLP 2024. 2024, 2380–2393
72. Fu J, Ng S K, Jiang Z, Liu P. Gptscore: Evaluate as you desire. In: Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers). 2024, 6556–6576
73. Saad-Falcon J, Khattab O, Potts C, Zaharia M. Ares: An automated evaluation framework for retrieval-augmented generation systems. In: Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers). 2024, 338–354
74. Zhao X, Zhang H, Pan X, Yao W, Yu D, Chen J. Thrust: Adaptively propels large language models with external knowledge. Advances in Neural Information Processing Systems, 2023, 36: 69930–69948
75. Zhu K, Feng X, Du X, Gu Y, Yu W, Wang H, Chen Q, Chu Z, Chen J, Qin B. An information bottleneck perspective for effective noise filtering on retrieval-augmented generation. arXiv preprint arXiv:2406.01549, 2024
76. Li D, Yan J, Zhang T, Wang C, He X, Huang L, Xue H, Huang J. On the role of long-tail knowledge in retrieval augmented large language models. arXiv preprint arXiv:2406.16367, 2024
77. Sun Z, Zang X, Zheng K, Song Y, Xu J, Zhang X, Yu W, Li H. Red-deep: Detecting hallucination in retrieval-augmented generation via mechanistic interpretability. arXiv preprint arXiv:2410.11414, 2024
78. Liu Y, Huang L, Li S, Chen S, Zhou H, Meng F, Zhou J, Sun X. Recall: A benchmark for llms robustness against external counterfactual knowledge. arXiv preprint arXiv:2311.08147, 2023
79. Wu S, Xie J, Chen J, Zhu T, Zhang K, Xiao Y. How easily do irrelevant inputs skew the responses of large language models? arXiv preprint arXiv:2404.03302, 2024
80. Liang X, Niu S, Li Z, Zhang S, Wang H, Xiong F, Fan J Z, Tang B, Song S, Wang M, others . Saferag: Benchmarking security in retrieval-augmented generation of large language model. arXiv preprint arXiv:2501.18636, 2025
81. Kang M, Gürel N M, Yu N, Song D, Li B. C-rag: Certified generation risks for retrieval-augmented language models. In: International Conference on Machine Learning. 2024, 22963–23000

82. Cheng X, Wang X, Zhang X, Ge T, Chen S Q, Wei F, Zhang H, Zhao D. xrag: Extreme context compression for retrieval-augmented generation with one token. *arXiv preprint arXiv:2405.13792*, 2024
83. Asai A, Wu Z, Wang Y, Sil A, Hajishirzi H. Self-rag: Learning to retrieve, generate, and critique through self-reflection. In: *The Twelfth International Conference on Learning Representations*. 2023
84. Trivedi H, Balasubramanian N, Khot T, Sabharwal A. Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions. In: *ACL (1)*. 2023, 10014–10037
85. Chen J, Lin H, Han X, Sun L. Benchmarking large language models in retrieval-augmented generation. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. 2024, 17754–17762
86. Zou W, Geng R, Wang B, Jia J. Poisonedrag: Knowledge corruption attacks to retrieval-augmented generation of large language models. *arXiv preprint arXiv:2402.07867*, 2024
87. Zhang Y, Li Q, Du T, Zhang X, Zhao X, Feng Z, Yin J. Hijackrag: Hijacking attacks against retrieval-augmented large language models. *arXiv preprint arXiv:2410.22832*, 2024
88. Chaudhari H, Severi G, Abascal J, Jagielski M, Choquette-Choo C A, Nasr M, Nita-Rotaru C, Oprea A. Phantom: General trigger attacks on retrieval augmented language generation. *arXiv preprint arXiv:2405.20485*, 2024
89. Shafran A, Schuster R, Shmatikov V. Machine against the rag: Jamming retrieval-augmented generation with blocker documents. *arXiv preprint arXiv:2406.05870*, 2024
90. Zeng S, Zhang J, He P, Liu Y, Xing Y, Xu H, Ren J, Chang Y, Wang S, Yin D, others. The good and the bad: Exploring privacy issues in retrieval-augmented generation (rag). In: *Findings of the Association for Computational Linguistics ACL 2024*. 2024, 4505–4524
91. Cheng P, Ding Y, Ju T, Wu Z, Du W, Yi P, Zhang Z, Liu G. Trojanrag: Retrieval-augmented generation can be backdoor driver in large language models. *arXiv preprint arXiv:2405.13401*, 2024
92. Chaudhari H, Severi G, Abascal J, Jagielski M, Choquette-Choo C A, Nasr M, Nita-Rotaru C, Oprea A. Phantom: General trigger attacks on retrieval augmented language generation. *arXiv preprint arXiv:2405.20485*, 2024
93. Perez E, Huang S, Song F, Cai T, Ring R, Aslanides J, Glaese A, McAleese N, Irving G. Red teaming language models with language models, 2022
94. Shrestha R, Zou Y, Chen Q, Li Z, Xie Y, Deng S. Fairrag: Fair human generation via fair retrieval augmentation. *CoRR*, 2024, abs/2403.19964
95. Zhou Y, Liu Z, Jin J, Nie J Y, Dou Z. Metacognitive retrieval-augmented large language models. In: *WWW*. 2024, 1453–1463
96. Sudhi V, Bhat S R, Rudat M, Teucher R. Rag-ex: A generic framework for explaining retrieval augmented generation. In: *SIGIR*. 2024, 2776–2780
97. Ding T, Banerjee A, Mombaerts L, Li Y, Borogovac T, Weinstein J P D I C. Vera: Validation and evaluation of retrieval-augmented systems. *arXiv preprint arXiv:2409.03759*, 2024
98. Anthropic. Reducing latency, January 2025
99. Hofstätter S, Chen J, Raman K, Zamani H. FiD-Light: Efficient and Effective Retrieval-Augmented Text Generation. In: *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23*. July 2023, 1437–1447
100. Kwiatkowski T, Palomaki J, Redfield O, Collins M, Parikh A, Alberti C, Epstein D, Polosukhin I, Devlin J, Lee K, Toutanova K, Jones L, Kelcey M, Chang M W, Dai A M, Uszkoreit J, Le Q, Petrov S. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 2019, 7: 453–466
101. Yang Z, Qi P, Zhang S, Bengio Y, Cohen W W, Salakhutdinov R, Manning C D. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In: *Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2018
102. Thorne J, Vlachos A, Christodoulopoulos C, Mittal A. FEVER: a large-scale dataset for fact extraction and VERification. In: *NAACL-HLT*. 2018
103. Dinan E, Roller S, Shuster K, Fan A, Auli M, Weston J. Wizard of Wikipedia: Knowledge-powered conversational agents. In: *Proceedings of the International Conference on Learning Representations (ICLR)*. 2019
104. DeYoung J, Jain S, Rajani N F, Lehman E, Xiong C, Socher R, Wallace B C. Eraser: A benchmark to evaluate rationalized nlp models. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 2020, 4443–4458
105. Zhang S, Liu X, Liu J, Gao J, Duh K, Van Durme B. Record: Bridging the gap between human and machine commonsense reading comprehension. *arXiv preprint arXiv:1810.12885*, 2018
106. Lyu Y, Li Z, Niu S, Xiong F, Tang B, Wang W, Wu H, Liu H, Xu T, Chen E. Crud-rag: A comprehensive chinese benchmark for retrieval-augmented generation of large language models. *ACM Trans. Inf. Syst.*, 2025, 43(2)
107. Xiong G, Jin Q, Lu Z, Zhang A. Benchmarking retrieval-augmented generation for medicine. In: *Findings of the Association for Computational Linguistics ACL 2024*. 2024, 6233–6251
108. Wang S, Khramtsova E, Zhuang S, Zuccan G. Feb4rag: Evaluating federated search in the context of retrieval augmented generation. In: *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2024, 763–773
109. Kamaloo E, Thakur N, Lassance C, Ma X, Yang J H, Lin J. Resources for brewing beir: reproducible reference models and an official leaderboard, 2023
110. Friel R, Belyi M, Sanyal A. Ragbench: Explainable benchmark for retrieval-augmented generation systems. *arXiv preprint arXiv:2407.11005*, 2024
111. Yu X, Cheng H, Liu X, Roth D, Gao J. ReEval: Automatic hallucination evaluation for retrieval-augmented large language models via transferable adversarial attacks. In: Duh K, Gomez H, Bethard S, eds, *Findings of the Association for Computational Linguistics: NAACL 2024*. June 2024, 1333–1351
112. Wang S, Liu J, Song S, Cheng J, Fu Y, Guo P, Fang K, Zhu Y, Dou Z. Domainrag: A chinese benchmark for evaluating domain-specific



- retrieval-augmented generation. CoRR, 2024
113. Roychowdhury S, Soman S, Ranjani H, Gunda N, Chhabra V, BALA S K. Evaluation of rag metrics for question answering in the telecom domain. ICML 2024 Workshop on Foundation Models in the Wild, 2024
  114. Pipitone N, Alami G H. Legalbench-rag: A benchmark for retrieval-augmented generation in the legal domain. arXiv preprint arXiv:2408.10343, 2024
  115. Zhu K, Luo Y, Xu D, Wang R, Yu S, Wang S, Yan Y, Liu Z, Han X, Liu Z, others . Rageval: Scenario specific rag evaluation dataset generation framework. CoRR, 2024
  116. Galla D, Hoda S, Zhang M, Quan W, Yang T D, Voyles J. Courage: A framework to evaluate rag systems. In: Rapp A, Di Caro L, Meziane F, Sugumaran V, eds, Natural Language Processing and Information Systems. 2024, 392–407
  117. Kasai J, Sakaguchi K, Le Bras R, Asai A, Yu X, Radev D, Smith N A, Choi Y, Inui K, others . Realtime qa: What’s the answer right now? Advances in neural information processing systems, 2023, 36: 49025–49043
  118. Wu X, Li S, Wu H T, Tao Z, Fang Y. Does RAG introduce unfairness in LLMs? evaluating fairness in retrieval-augmented generation systems. In: Rambow O, Wanner L, Apidianaki M, Al-Khalifa H, Eugenio B D, Schockaert S, eds, Proceedings of the 31st International Conference on Computational Linguistics. January 2025, 10021–10036
  119. Liu J, Ding R, Zhang L, Xie P, Huang F. Cofe-rag: A comprehensive full-chain evaluation framework for retrieval-augmented generation with enhanced data diversity. arXiv preprint arXiv:2410.12248, 2024
  120. Wang S, Tan J, Dou Z, Wen J R. Omnieval: An omnidirectional and automatic rag evaluation benchmark in financial domain. arXiv preprint arXiv:2412.13018, 2024
  121. Yang X, Sun K, Xin H, Sun Y, Bhalla N, Chen X, Choudhary S, Gui R D, Jiang Z W, Jiang Z, Kong L, Moran B, Wang J, Xu Y E, Yan A, Yang C, Yuan E, Zha H, Tang N, Chen L, Scheffer N, Liu Y, Shah N, Wanga R, Kumar A, Yih W t, Dong X L. Crag - comprehensive rag benchmark. In: Globerson A, Mackey L, Belgrave D, Fan A, Paquet U, Tomczak J, Zhang C, eds, Advances in Neural Information Processing Systems. 2024, 10470–10490
  122. Papadimitriou I, Gialampoukidis I, Vrochidis S, others . Rag playground: A framework for systematic evaluation of retrieval strategies and prompt engineering in rag systems. arXiv preprint arXiv:2412.12322, 2024
  123. Katsis Y, Rosenthal S, Fadnis K, Gunasekara C, Lee Y S, Popa L, Shah V, Zhu H, Contractor D, Danilevsky M. Mtrag: A multi-turn conversational benchmark for evaluating retrieval-augmented generation systems. arXiv preprint arXiv:2501.03468, 2025
  124. Xu Z, Li Y, Ding R, Wang X, Chen B, Jiang Y, Zheng H, Lu W, Xie P, Huang F. Let llms take on the latest challenges! a chinese dynamic question answering benchmark. In: Proceedings of the 31st International Conference on Computational Linguistics. 2025, 10435–10448
  125. Gao Y, Xiong Y, Wu W, Huang Z, Li B, Wang H. U-niah: Unified rag and llm evaluation for long context needle-in-a-haystack. arXiv preprint arXiv:2503.00353, 2025
  126. Selvaraj T. Calculate the total cost of a retrieval augmented generation (rag) solution, February 2024
  127. Zhang J, Li G, Su J. Sage: A framework of precise retrieval for rag. arXiv preprint arXiv:2503.01713, 2025
  128. Su J, Healey J, Nakov P, Cardie C. Fast or better? balancing accuracy and cost in retrieval-augmented generation with flexible user control. CoRR, 2025
  129. Şakar T, Emekci H. Maximizing rag efficiency: A comparative analysis of rag methods. Natural Language Processing, 2025, 31(1): 1–25
  130. Datta A, Fredrikson M, Leino K, Lu K, Sen S, Shih R, Wang Z. Exploring conceptual soundness with trulens. In: NeurIPS 2021 Competitions and Demonstrations Track. 2022, 302–307
  131. LangChain . Evaluating rag architectures on benchmark tasks, November 2023
  132. Mahboub A, Za’ter M E, Al-Rfooh B, Estaitia Y, Jaljuli A, Hakouz A. Evaluation of semantic search and its role in retrieved-augmented-generation (rag) for arabic language. arXiv preprint arXiv:2403.18350, 2024
  133. Dong G, Song X, Zhu Y, Qiao R, Dou Z, Wen J R. Toward general instruction-following alignment for retrieval-augmented generation. arXiv preprint arXiv:2410.09584, 2024
  134. Salemi A, Zamani H. Evaluating retrieval quality in retrieval-augmented generation. In: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR ’24. 2024, 2395–2400
  135. Jaech A, Kalai A, Lerer A, Richardson A, El-Kishky A, Low A, Hellary A, Madry A, Beutel A, Carney A, others . Openai o1 system card. arXiv preprint arXiv:2412.16720, 2024



Aoran Gan is working toward the PhD degree in the School of Artificial Intelligence and Data Science, University of Science and Technology of China. His research interests include text mining, knowledge graph and large language models.



Hao Yu is pursuing a MS degree at McGill University and is affiliated with Quebec Artificial Intelligence Institute. His research focuses on multilingual and low-resource NLP, as well as RAG systems for misinformation detection.



Kai Zhang is an Associate Researcher at the University of Science and Technology of China. His general area of research is natural language processing and knowledge discovery. He is a member of ACM, SIGIR, AAAI, and CCF.



Shiwei Tong is a senior data scientist at Tencent Company. His research focuses on Game Data Mining and Game Applications driven by Large Language Models.



Qi Liu is a professor in the School of Artificial Intelligence and Data Science at USTC. His area of research is data mining and knowledge discovery. He has published prolifically in refereed journals and conferences. He is an Associate Editor of IEEE TBD and Neurocomputing.



Wenyu Yan is currently pursuing MS degree in University of Science and Technology of China. His research interests focus on conversational search, retrieval-augmented generation, etc.



Enhong Chen is a professor in the School of Computer Science and Technology at USTC. His general area of research includes data mining and machine learning, social network analysis, and recommender systems. He was on program committees of numerous conferences including SIGKDD, ICDM, and SDM.



Zhenya Huang is currently an Associate Professor with USTC. His main research interests include data mining, knowledge reasoning, natural language processing, and intelligent education. He has published more than 50 papers in refereed journals and conference proceedings.



Guoping Hu is senior vice president of iFLYTEK, director of the National Key Laboratory of Cognitive Intelligence. He has been honored with the First Prize of State Science and Technology Advancement Award and garnered over 300 authorized patents.