

Vendi-RAG: Adaptively Trading-Off Diversity And Quality Significantly Improves Retrieval Augmented Generation With LLMs

Mohammad R. Rezaei

University of Toronto

Vector Institute

Vertaix

mr.rezaei@mail.utoronto.ca

Adji Bousso Dieng

Princeton University

Vertaix

adji@princeton.edu

Abstract

Retrieval-augmented generation (RAG) for multi-hop question answering is limited by retrieving semantically redundant documents, when reasoning instead needs diverse, complementary sources. We introduce **Vendi-RAG**, which treats set-level semantic diversity as a first-class retrieval objective. It combines an eigenvalue-based diversity score (the Vendi Score) with query relevance into a single Vendi Retrieval Score, embedded in an iterative loop that uses an LLM-based answer-quality signal to adaptively rebalance diversity and relevance. On HotpotQA, MuSiQue, and 2WikiMultiHopQA, Vendi-RAG improves accuracy by +2.4, +1.6, and +2.4 points over the strongest diversity-aware iterative baseline (Adaptive Retrieval + MMR). A fixed- s variant captures most of the gain at no extra inference cost, while the LLM-judge loop adds a smaller but significant gain on hard queries. Results hold across five instruction-tuned backbones and two dense retrievers, identifying set-level retrieval diversity as a primary, controllable axis for improving multi-hop RAG (🔗 [vertaix/Vendi-RAG](https://vertaix.com/Vendi-RAG)).

1 Introduction

Retrieval-augmented generation (RAG) enhances large language models (LLMs) on knowledge-intensive tasks such as question answering (QA) by retrieving relevant information from external sources beyond the training set (Achiam et al., 2023; Team et al., 2023; Jiang et al., 2024). It has been especially successful in single-hop QA, where a question can be answered from a single document (Raiaan et al., 2024; Kwiatkowski et al., 2019). Multi-hop queries, however, require synthesizing information from several distinct sources. For example, answering “What is the capital of the country where the inventor of the telephone was born?” requires combining biographical facts about Bell’s birthplace with geographical facts

Method	Retrieved Documents	Answer
Standard RAG	1. "A.G. Bell invented the telephone..." [tel] 2. "The telephone was patented by Bell..." [tel] 3. "Bell's telephone revolutionized..." [tel] 4. "First telephone call: 'Mr. Watson...'" [tel] 5. "Bell demonstrated the telephone..." [tel]	1. "A.G. Bell invented the telephone..." [tel]
MMR	1. "A.G. Bell invented the telephone..." [tel] 2. "Edinburgh is in Scotland, UK..." [geo] 3. "Bell was a teacher of the deaf..." [bio] 4. "Scottish inventors: J.L. Baird..." [inv] 5. "UK: England, Scotland, Wales..." [geo]	1. "A.G. Bell invented the telephone..." [tel]
Vendi-RAG	1. "Bell born in Edinburgh, Scotland..." [bio] 2. "Edinburgh: Scotland's capital..." [geo] 3. "Bell's telephone patent..." [tel] 4. "Scottish education influenced Bell..." [edu] 5. "UK capitals: Edinburgh (Scotland)..." [geo]	1. "Bell born in Edinburgh, Scotland..." [bio] 2. "Edinburgh: Scotland's capital..." [geo] 3. "Bell's telephone patent..." [tel] 4. "Scottish education influenced Bell..." [edu] 5. "UK capitals: Edinburgh (Scotland)..." [geo]

Table 1: Retriever comparison for “What is the capital of the country where the inventor of the telephone was born?” Relevance-only retrieval returns redundant telephone documents, whereas Vendi-RAG retrieves complementary, query-relevant documents across categories.

about Scotland’s capital (Table 1), which exposes a limitation of relevance-based retrieval.

The core problem is that relevance-based retrieval tends to return semantically similar documents with redundant content, while multi-hop reasoning needs diverse, complementary sources that supply the different links of the reasoning chain (see Table 1).

Iterative RAG pipelines partly address this by re-querying and refining retrieved documents over multiple rounds. Adaptive-RAG (Jeong et al., 2024) routes queries to no-retrieval, single-step, or iterative retrieval based on a query-complexity classifier; Self-RAG (Asai et al., 2023) adds on-demand retrieval with self-reflection; and IR-CoT (Trivedi et al., 2023) interleaves retrieval with chain-of-thought reasoning (Wei et al., 2022; Wang and Zhou, 2024). Despite their success, these methods rely on relevance-based retrieval and do not




Figure 1: Vendi-RAG begins with an initial retrieval that uses the Vendi Score to obtain a diverse, broadly covering document set. The LLM then reasons over these documents to produce an answer, which an LLM judge scores for quality. Based on this score, the retriever rebalances diversity and relevance via an adaptive parameter s : high-quality answers reduce the emphasis on diversity, while low-quality answers increase it. The loop iterates until quality reaches the threshold Thr , after which the highest-quality answer and documents are returned.

actively manage the *diversity* of retrieved information, which is precisely what complex queries demand.

We therefore propose *Vendi retrieval* the Vendi Retrieval Score (VRS) score function that combines the Vendi Score (VS), an eigenvalue-based measure of the effective number of distinct documents, with query relevance through a single scalar weight s . This scoring function is the paper’s primary contribution, and it alone (at a fixed s , with no extra inference cost) accounts for the bulk of our gains. Around it we build **Vendi-RAG**, an iterative pipeline that retrieves candidate documents, generates chain-of-thought reasoning and candidate answers, scores them with an LLM judge for coherence, relevance, and query alignment, and (if the quality threshold is not met) rebalances s before retrieving again, returning the highest-scoring answer across iterations (Figure 1). The loop is an optional quality-on-demand layer over the scoring function rather than the source of the improvement. The contributions of this paper are:

1. **Our main contribution:** a *set-level* retrieval objective for multi-hop QA. The Vendi Retrieval Score combines query relevance with the Vendi Score, an eigenvalue-based measure of the effective number of unique documents. This differs from pairwise objectives (MMR), log-determinant objectives (k -DPP), and post-retrieval utility selectors (Jain and Garimella, 2025; Zhang et al., 2025), and exposes a single scalar s that trades diversity against relevance, it is a drop-in scoring function that composes

with existing retrieval loops.

2. A secondary *explore-exploit* feedback loop that adapts s to per-query difficulty using an LLM-based answer-quality signal. It adds a smaller gain concentrated on hard queries, at extra judge cost, and helps even when the judge is a different or smaller model than the generator, so it requires only a *relative* quality signal rather than absolute calibration.
3. A systematic empirical study on three multi-hop benchmarks (HotpotQA, MuSiQue, 2WikiMultiHopQA) and five LLM backbones (GPT-3.5, GPT-4o, GPT-4o-mini, Llama 3-70B-Instruct, Gemini Flash 2.0). Vendi-RAG improves multi-hop accuracy over Adaptive-RAG and over diversity-aware iterative baselines (MMR, k -DPP, IRCOT) under matched retrievers and test splits, and the improvement is robust to the choice of dense encoder.

We evaluate Vendi retrieval and Vendi-RAG on three multi-hop QA benchmarks: HotpotQA (Yang et al., 2018), MuSiQue (Trivedi et al., 2022), and 2WikiMultiHopQA (Ho et al., 2020). Vendi retrieval yields more diverse documents than the baselines on both the Vendi Score and max pairwise distance, and that diversity translates into higher coverage of the gold supporting evidence. The most informative downstream comparison is to other diversity-aware iterative methods: with GPT-3.5, Vendi-RAG improves accuracy over the strongest such baseline (Adaptive Retrieval + MMR) by +2.4 on HotpotQA, +1.6 on MuSiQue, and +2.4 on 2WikiMultiHopQA, with paired-bootstrap $p <$

0.05. Against routing-based Adaptive-RAG, the gap is +11.4 on HotpotQA, +0.8 on MuSiQue, and +4.6 on 2WikiMultiHopQA; the larger HotpotQA gap reflects Adaptive-RAG’s complexity-classifier routing of many HotpotQA queries to no-retrieval or single-step paths in our setup, where iterative diversity-aware methods (MMR, IRCot, k -DPP, Vendi-RAG) all sit in the same 55–59 accuracy band. Crucially, a fixed- s variant of Vendi-RAG with the scoring function alone captures most of that improvement at the same inference cost as MMR, which locates the gain in the retrieval objective rather than in the loop; the LLM-judge loop adds a smaller but significant gain at higher latency. The pattern is consistent across five instruction-tuned backbones (GPT-3.5, GPT-4o, GPT-4o-mini, Llama 3-70B-Instruct, Gemini Flash 2.0) and two dense retrievers on multi-hop English QA.

2 Related Work

QA approaches fall into three families: non-retrieval methods (Petroni et al., 2019), single-step RAG (Lewis et al., 2020), and multi-step RAG (Asai et al., 2023). Non-retrieval methods answer directly from the LLM; they are efficient but struggle with external or up-to-date information and hallucinate on out-of-distribution queries (Shuster et al., 2021). Single-step RAG improves factual accuracy but is limited by retrieval noise on complex reasoning (Trivedi et al., 2023). Multi-step RAG targets multi-hop queries by iteratively retrieving and refining (Jeong et al., 2024; Asai et al., 2023; Tang and Yang, 2024), enabling cross-source reasoning at the cost of overhead and error accumulation (Jeong et al., 2024).

Advances in multi-hop QA. Recent work focuses on question decomposition (Radhakrishnan et al., 2023), chain-of-thought reasoning (Wei et al., 2022; Liu et al., 2024a), and iterative retrieval (Jeong et al., 2024; Shao et al., 2023; Yu et al., 2024). ReCite (Sun et al., 2022) and IRCot (Trivedi et al., 2023) refine retrieval with progressive reasoning, while Self-RAG (Asai et al., 2023) uses self-reflection tokens to decide when to retrieve. Decomposed prompting (Khot et al., 2022) aids complex queries (Zhang et al., 2024b), and MultiHop-RAG (Tang and Yang, 2024) combines decomposition with retrieval but remains constrained by relevance-based retrieval, limiting synthesis of diverse information. More recent systems change the index or the retrieval *policy*: Hip-

poRAG (Gutiérrez et al., 2024) and AMG-RAG (Rezaei et al., 2025) add a knowledge-graph index, Beam Retrieval (Zhang et al., 2024a) learns beam search over supporting-passage chains, and Search-R1 (Jin et al., 2025) and Search-o1 (Li et al., 2025) learn agentic search policies. All still rank passages by relevance at each step and so do not control *set-level* redundancy, which the VRS supplies as a composable scoring function.

Diversity in information retrieval. MMR (Carbonell and Goldstein, 1998) pioneered the relevance-diversity trade-off via greedy pairwise-dissimilarity penalties. Intent-aware methods such as xQuAD (Santos et al., 2010) and PM-2 (Dang and Croft, 2012) model query aspects but require predefined intent taxonomies that may not generalize to open-ended multi-hop queries. Determinantal Point Processes (Kulesza and Taskar, 2012), applied to summarization (Gong et al., 2014) and recommendation (Chen et al., 2018), model set-level diversity but scale as $O(n^3)$ for exact inference and optimize a log-determinant rather than the interpretable “effective number of items” of the Vendi Score. Submodular optimization (Lin and Bilmes, 2011; Wei et al., 2014) offers greedy approximation guarantees but typically needs explicit features and may miss the semantic similarities encoded in dense embeddings. More recent work moves diversification into learned neural rankers (Yan et al., 2021) and fast greedy MAP inference for DPPs (Chen et al., 2018); Wu et al. (2024) survey the area. Closest to our setting, Pickett et al. (2024) replace the MMR trade-off with a relevant-information-gain objective for RAG context selection, which is likewise set-level but is defined through a probabilistic relevance model, whereas the VRS reads directly off the document-kernel eigenvalues and reports an interpretable effective number of distinct documents.

Vendi retrieval differs in that (1) the Vendi Score directly measures the effective number of distinct documents, making the trade-off transparent; (2) it operates on kernel matrices from dense embeddings without explicit intents or features; (3) the eigenvalue computation runs only on the small retrieved set ($k \leq 10$); and (4) its dynamic adjustment adapts the diversity-relevance balance from answer-quality feedback, unlike static diversification.

Vendi scoring. The Vendi Score (VS) (Friedman and Dieng, 2023) is a similarity-based diversity metric used across machine learning (Berns

et al., 2023; Pasarkar and Dieng, 2023; Mousavi and Khalili, 2025; Nguyen and Dieng, 2024; Kanen et al., 2024; Jalali et al., 2024; Askari Hemmat et al., 2024; Rezaei and Dieng, 2025; Bhardwaj et al., 2025), chemistry (Pasarkar et al., 2023), materials science (Liu et al., 2024b), and biology (Pasarkar and Dieng, 2025). Unlike pairwise (MMR) or log-determinant (DPP) metrics, VS computes the exponential of the Shannon entropy of the kernel eigenvalues, an interpretable measure equal to the effective number of orthogonal components. Vendi-RAG integrates VS into retrieval to balance diversity and quality (Carbonell and Goldstein, 1998; Slivkins et al., 2010), improving robustness and accuracy over relevance-only retrieval (Guu et al., 2020) in multi-hop QA.

Multi-hop retrieval planning. Some systems decompose queries and schedule retrieval steps explicitly (Khattab et al., 2022; Press et al., 2023), and entity-centric approaches (Pandya and Bhatt, 2021) track entities across hops. These operate at the retrieval-loop level and are composable with Vendi-RAG’s diversity-aware scoring function.

Context-independent passage utility. A second line learns a *utility* signal on top of relevance: Jain and Garimella (2025) supervise a passage selector with LLM reasoning traces, and Zhang et al. (2025) distill an LLM into a utility-based selector over a relevance-retrieved pool. These act *post-retrieval* as filters, whereas Vendi retrieval acts earlier, shaping the candidate pool itself.

3 Method

We describe Vendi-RAG, our framework for multi-hop QA based on diversity-driven document retrieval.

3.1 Vendi Retrieval

Diversity in retrieved documents is essential for multi-hop QA: it ensures broad semantic coverage, reduces redundancy, and brings in multiple perspectives (Sun et al., 2022; Carbonell and Goldstein, 1998; Thakur et al., 2021). Similarity search retrieves on relevance alone and returns redundant documents. Three families of methods address this. MMR (Carbonell and Goldstein, 1998) balances relevance and novelty through *pairwise* dissimilarity penalties but lacks a notion of set-level diversity. k -DPP (Kulesza and Taskar, 2012) samples subsets in proportion to the determinant of a kernel submatrix, but the log-determinant objective scales

poorly and is not directly interpretable as “number of unique documents”. Post-retrieval utility selectors (Jain and Garimella, 2025; Zhang et al., 2025) learn passage-level usefulness on top of a fixed relevance-ranked pool; they cannot recover useful evidence that is absent from that pool. We instead define diversity at the set level using the Vendi Score (VS) (Friedman and Dieng, 2023), which quantifies the effective number of unique documents in a set $\mathcal{D} = \{d_1, \dots, d_n\}$:

$$\text{VS}_k(\mathcal{D}) = \exp\left(-\sum_{i=1}^n \lambda_i \log \lambda_i\right), \quad (1)$$

where $\lambda_1, \dots, \lambda_n$ are the eigenvalues of the normalized similarity matrix K/n , with $K_{ij} = k(d_i, d_j)$ for a PSD kernel k ($k(d_i, d_i) = 1$). The VS is maximal (n) when documents are orthogonal and minimal (1) when identical (Friedman and Dieng, 2023).

Vendi Retrieval Score. To balance diversity with relevance, we define the Vendi Retrieval Score (VRS) as a convex combination of the two, each min-max normalized to $[0, 1]$ (full normalization in Appendix E):

$$\text{VRS} = s \cdot \widetilde{\text{VS}}_k(\mathcal{D}) + (1 - s) \cdot \widetilde{\text{SS}}(q, \mathcal{D}), \quad (2)$$

where $s \in [0, 1]$ controls the trade-off and $\text{SS}(q, \mathcal{D})$ is the average query-document cosine similarity from dense embeddings. Higher s favors diversity; lower s favors relevance.

Greedy subset selection. Optimizing the VRS over all k -subsets is intractable, so documents are selected greedily, at each step adding $d^* = \arg \max_{d \in \mathcal{C} \setminus \mathcal{D}_{i-1}} \text{VRS}(\mathcal{D}_{i-1} \cup \{d\})$ from the candidate set $\mathcal{C}$ returned by similarity search. Per query, this costs $O(k \cdot |\mathcal{C}| \cdot k^3)$, where the dominant k^3 term is the single VS eigendecomposition on the current candidate-augmented subset of size $\leq k$, repeated for each of the k slots over the $|\mathcal{C}|$ candidates. For typical sizes ($k \leq 10$, $|\mathcal{C}| \leq 50$) this is well under 50ms per query in our implementation and is negligible compared with LLM-side latency. The VRS is non-submodular, so greedy selection carries no formal guarantee, but the empirical gap between greedy and exhaustive search is negligible at small k (Appendix E.5); greedy optimization is similarly used on the non-submodular MMR objective. Relevance enters twice— $\mathcal{C}$ is first selected by similarity search, and the SS term then refines the ranking—so the VRS never optimizes diversity over query-irrelevant documents.

3.2 Vendi-RAG

Vendi-RAG embeds Vendi retrieval in an iterative loop. Each iteration (i) retrieves documents, (ii) generates chain-of-thought reasoning (Wei et al., 2022) and a candidate answer grounded in that evidence, (iii) scores the answer with an LLM judge, and (iv) updates the diversity weight s . The loop stops when the quality score exceeds τ or after N iterations, returning the highest-scoring answer (Algorithm 1). We detail the two non-standard components below; prompts and pipeline specifics are in Appendix E.

Quality evaluation. An LLM judge scores each candidate on three structural criteria—*coherence* (logical consistency of the reasoning), *relevance* (grounding in the retrieved evidence), and *query alignment* (how completely the answer addresses the question)—each on a 1–10 scale. The quality score is their mean divided by 10, giving $Q_t = \frac{1}{10} \text{mean}(C, R, Q) \in [0, 1]$, with stopping threshold $\tau = 0.85$ (equivalently 8.5/10 on the raw scale). The judge influences only the diversity weight s and never edits the answer text; factual grounding is checked separately by an external NLI faithfulness metric (Appendix D.8). The exact judge prompt is reproduced in Appendix E.6.

Dynamic adjustment of s . Let s_t be the diversity weight at iteration t . It is updated via an exponential moving average (EMA) with momentum β :

$$s_{t+1}^{\text{target}} = \text{clip}\left(1 - \frac{Q_t}{\max(\max_{\tau \leq t} Q_\tau, \epsilon)}, 0, 1\right),$$

$$s_{t+1} = \beta s_t + (1 - \beta) s_{t+1}^{\text{target}}, \quad (3)$$

with $\epsilon = 0.01$ and $\beta = 0.3$. The rule is explore–exploit: low Q_t raises the target toward diversity (broader exploration for missing evidence), high Q_t lowers it toward relevance (consolidation along the current reasoning path), analogous to decaying exploration in reinforcement learning. The target is *relative* to the running best rather than absolutely calibrated, which asks far less of the judge than a calibrated score would; EMA smoothing, an *absolute* stopping threshold, and best-answer tracking together bound what a mistimed update can cost, and s carries forward across iterations independently of query rewriting (Appendix E.3). Trajectories converge in 3–4 iterations on average. The pipeline also scales effort with difficulty: easy queries converge in 1–2 iterations while harder

Algorithm 1 Vendi-RAG Inference Pipeline

Require: Query q , Knowledge base $\mathcal{K}$, Max iterations N , Quality threshold τ , Momentum β
Ensure: Final answer $\hat{a}^*$

```

1: Initialize:  $q_1 \leftarrow q$ ,  $s_1 \leftarrow 0.8$ ,  $\text{best\_quality} \leftarrow 0$ 
2: for  $i = 1$  to  $N$  do
3:    $\mathcal{C} \leftarrow \text{TopK-Similar}(q_i, \mathcal{K}, k=50)$ 
4:    $D_i \leftarrow \text{GreedySelect}(\mathcal{C}, s_i, k=5)$  ▷ Eq. (2)
5:    $r_i \leftarrow \text{GenerateCoT}(q_i, D_i)$ 
6:    $\hat{a}_i \leftarrow \text{GenerateAnswer}(q, D_i, r_{1:i})$ 
7:    $Q_i \leftarrow \text{LLM-Judge}(\hat{a}_i, q, D_i)$ 
8:   if  $Q_i > \text{best\_quality}$  then
9:      $\hat{a}^* \leftarrow \hat{a}_i$ ,  $\text{best\_quality} \leftarrow Q_i$ 
10:  end if
11:  if  $Q_i \geq \tau$  then
12:    return  $\hat{a}^*$ 
13:  end if
14:   $s^{\text{target}} \leftarrow \text{clip}(1 - Q_i / \max(\text{best\_quality}, \epsilon), 0, 1)$ 
15:   $s_{i+1} \leftarrow \beta \cdot s_i + (1 - \beta) \cdot s^{\text{target}}$ 
16:   $q_{i+1} \leftarrow \text{RefineQuery}(q_i, \hat{a}_i, r_i)$ 
17: end for
18: return  $\hat{a}^*$ 

```

ones use 4–5. The dominant per-step cost is the $O(k^3)$ VS eigendecomposition on the retrieved subset (with $k \leq 10$, not corpus size); a full latency breakdown is in Table 7.

4 Experiments

Here, we present a comprehensive evaluation of Vendi-RAG on multi-hop QA tasks. We first analyze the effectiveness of the Vendi retrieval strategy in enhancing retrieval diversity and evidence coverage, and then evaluate the full Vendi-RAG pipeline, comparing it against strong baselines for multi-hop QA benchmarks: MuSiQue (Trivedi et al., 2022), HotpotQA (Yang et al., 2018), and 2WikiMultiHopQA (Ho et al., 2020) (see Appendix A for details). Additional analyses, including the s sensitivity study, faithfulness evaluation, computational cost, cross-model judge robustness, and per-document-size trends, are deferred to Appendix D.

Experimental setup. We use the standard train/dev/test splits for each dataset, with evaluation on the test set (500 samples each for MuSiQue and 2WikiMultiHopQA, 7,405 for HotpotQA). All experiments are repeated across 3 random seeds (42, 123, 456), and we report mean $\pm$ standard deviation. For document retrieval, we use SentenceBERT (all-mpnet-base-v2) as the dense encoder for all methods to ensure a fair comparison. Statistical significance is assessed using paired bootstrap tests with 10,000 resamples; markers (*) indicate $p < 0.05$ against the *per-cell* strongest non-Vendi comparator in Table 4 (the per-cell comparator is listed in Appendix D.4), after Holm–Bonferroni

correction (Holm, 1979) across the 9-test family formed by the 3 datasets $\times$ 3 metrics. Raw and corrected p -values are reported in Appendix D.4.

Baseline configurations. Every method uses the same dense encoder (all-mpnet-base-v2), the same corpus and chunking, the same GPT-3.5 Turbo answer generator and answer prompt, and the same test splits; only the retrieval component differs. Methods that re-rank or diversify (MMR, k -DPP, cross-encoder reranking, IRCOT, Vendi-RAG) draw from the same top-50 similarity-search pool and pass $k = 5$ documents to the reader, whereas relevance-only baselines retrieve $k = 5$ directly with no intermediate pool; all iterative methods share a budget of $N = 5$. Two naming points are worth stating explicitly: “Adaptive Retrieval” and “Adaptive-RAG” are adaptive in their *routing*, not in their scoring function, which is plain relevance in both. Self-RAG is designed for single-hop retrieval with on-demand self-reflection, and under the multi-hop setup here its self-reflection tokens frequently suppress further retrieval, so its numbers reflect a design mismatch rather than a tuning failure (consistent with Jeong et al., 2024) and are included for completeness only. Per-method implementation details, together with the candidate-pool size, final k , and iteration budget of every method, are given in Appendix C.1. The baselines span both retrieval scoring functions and retrieval loop structures, so that the contribution of the VRS scoring function can be isolated from the contribution of the loop design.

Vendi retrieval improves document retrieval diversity. To assess the impact of the Vendi retrieval process on retrieval diversity, we compared the diversity of outputs from Vendi-RAG against Adaptive-RAG and Adaptive Retrieval, using the Vendi Score (VS) and the max pairwise distance (MPD). Table 2 summarizes the results. Vendi-RAG achieves higher diversity than both baselines on all datasets, demonstrating its ability to retrieve documents capturing multiple perspectives relevant to the query. Adaptive-RAG, which incorporates iterative refinement but lacks explicit diversity control, shows only moderate improvement over Adaptive Retrieval.

Diversity translates into better evidence coverage. Beyond answer quality, we evaluate retrieval effectiveness using supporting facts recall (SF Recall) and document recall (Doc Recall@ k), which

Dataset	Method	VS	MPD
MuSiQue	Adapt. Retrieval	6.13 $\pm$ 0.08	1.25 $\pm$ 0.03
	Adaptive-RAG	6.55 $\pm$ 0.11	1.42 $\pm$ 0.05
	Vendi-RAG	7.12$\pm$0.09*	1.95$\pm$0.04*
Hotpot	Adapt. Retrieval	4.95 $\pm$ 0.12	1.10 $\pm$ 0.04
	Adaptive-RAG	5.21 $\pm$ 0.15	1.31 $\pm$ 0.06
	Vendi-RAG	6.82$\pm$0.10*	1.89$\pm$0.05*
Wiki	Adapt. Retrieval	5.34 $\pm$ 0.09	1.32 $\pm$ 0.04
	Adaptive-RAG	5.81 $\pm$ 0.13	1.45 $\pm$ 0.05
	Vendi-RAG	6.68$\pm$0.11*	1.78$\pm$0.04*

Table 2: Retrieval diversity (Vendi Score (VS) and Max Pairwise Distance (MPD)) across datasets and methods. Vendi-RAG achieves higher diversity than baselines. Results are mean $\pm$ std over 3 seeds.

measure whether the retrieved documents contain the gold evidence needed for multi-hop reasoning. Table 3 shows that Vendi-RAG’s diversity-aware retrieval improves evidence coverage over relevance-only and MMR baselines, directly linking the diversity gains in Table 2 to a downstream-relevant quantity. The VRS combines diversity with query relevance over a candidate pool already filtered by similarity search (top-50 by cosine similarity), so the additional diversity is constrained to query-relevant documents and manifests as higher coverage of the gold supporting facts rather than topically unrelated content. To separate the diversity effect from the effect of running additional retrieval rounds, the SF Recall and Doc Recall@5 in Table 3 are evaluated on the *first* retrieval round only, with no query rewriting and no iterative refinement, for every method. Under this matched-round setup Vendi-RAG already retrieves more gold supporting facts than Adaptive-RAG and MMR (+7.7 and +10.2 SF Recall, respectively), so the coverage gain is attributable to the VRS scoring function on the candidate pool and not to extra iterations. A complementary matched-iteration view at the answer level is provided by the iterative diversity-aware baselines in Table 4: all of them use the same $N = 5$ iteration budget and the same retriever, so the +2.4 accuracy gap on HotpotQA between Vendi-RAG (58.4) and MMR (56.0) cannot be explained by “more iterations.”

Accuracy on multi-hop QA tasks. We next evaluate the full Vendi-RAG pipeline on multi-document reasoning. Table 4 reports results using GPT-3.5 Turbo. Vendi-RAG with dynamically adjusted s_t (initialized at $s_1 = 0.8$) achieves the best Exact Match (EM) and Accuracy (Acc) on

Method	SF Recall	Doc Recall@5
Adaptive Retrieval	58.3 $\pm$ 1.2	62.1 $\pm$ 0.9
Adaptive-RAG	63.7 $\pm$ 1.5	67.4 $\pm$ 1.1
MMR ($\lambda = 0.5$)	61.2 $\pm$ 1.3	65.8 $\pm$ 1.0
Vendi-RAG	71.4$\pm$1.1*	74.2$\pm$0.8*

Table 3: Retrieval quality metrics on HotpotQA. SF Recall measures supporting facts coverage; Doc Recall@5 measures gold document retrieval.

all three datasets. F1 gains are smaller and not uniformly the largest in the table; F1 rewards partial token overlap, on which relevance-only retrieval already performs adequately, while EM and Acc more directly reflect full multi-hop correctness, where diversity-aware retrieval contributes the most. The starred entries in Table 4 mark cells where the paired-bootstrap test (10,000 resamples, $p < 0.05$) is computed against the strongest non-Vendi baseline on that specific cell—MMR for HotpotQA Acc and 2WikiMultiHopQA Acc, IR-CoT for HotpotQA EM/F1 and MuSiQue EM, and Adaptive-RAG for MuSiQue Acc and 2WikiMultiHopQA EM. The two F1 cells where Adaptive-RAG is highest (MuSiQue and 2WikiMultiHopQA) carry no significance marker. Because nine cells are tested at once, the raw p -values are additionally Holm–Bonferroni corrected within that family (Holm, 1979); every starred cell survives the correction, with MuSiQue accuracy the closest call ($p = 0.048$). Appendix D.4 lists the per-cell comparator, the raw and corrected p -values, and the test details.

All iterative methods in Table 4 use the same maximum iteration budget ($N = 5$), the same retriever, and the same $k = 5$, so loop structure is held fixed across the diversity-aware baselines. The gap from Adaptive-RAG (47.0 HotpotQA Acc) up to MMR/ k -DPP/IRCoT (55.8–56.0) and then to Vendi-RAG (58.4) reflects *scoring-function* differences. Adaptive-RAG’s underperformance on HotpotQA is driven by its complexity-classifier, which routes many HotpotQA queries to no-retrieval or single-step paths (consistent with Jeong et al. 2024); once a method always iterates with diversity-aware retrieval, accuracy concentrates in the 55–59 band, and Vendi-RAG sits at the top of that band. The headline gap of +2.4 accuracy points over MMR is therefore the most informative comparison for the contribution of the VRS.

Ablations: VS, fixed- s , and the LLM-judge loop. The Vendi-RAG rows in Table 4 isolate three de-

sign choices. (i) Replacing similarity ranking with the VRS at fixed $s = 0.3$ already lifts HotpotQA accuracy to 55.7, matching MMR/ k -DPP/IRCoT. (ii) Increasing the fixed diversity weight to $s = 0.8$ yields 58.0 HotpotQA accuracy, which is the bulk of the total gain over MMR (+2.0 of +2.4 points; analogous patterns hold for MuSiQue and 2WikiMultiHopQA). (iii) The LLM-judge dynamic loop adds a further +0.4 accuracy on HotpotQA, +0.2 on MuSiQue, and +0.4 on 2WikiMultiHopQA on top of fixed $s = 0.8$. A paired-bootstrap test at the per-query level (Appendix D.4) finds the dynamic-vs-fixed- $s=0.8$ gap significant at $p < 0.05$ on HotpotQA, where most queries trigger ≥ 3 iterations and the loop is active, and borderline ($p \approx 0.08$ – 0.12) on MuSiQue and 2WikiMultiHopQA, where many queries terminate after ≤ 2 iterations. Because the judge both rebalances s and decides when to stop, we further separate the two with a 2×2 factorial that holds the iteration count fixed: s -reweighting alone contributes +0.15, the stopping rule alone +0.20, and a small positive interaction (+0.05) makes up the rest of the +0.40, so the two mechanisms are complementary rather than redundant (Appendix D.3). The takeaway is that the bulk of the improvement is delivered by the VS-based scoring function with a high fixed diversity weight, and the dynamic loop adds a smaller but consistent gain that concentrates on the harder queries where the iteration budget is fully used; the fixed- s variant therefore gives the same retrieval-side gain at the same inference cost as MMR.

Comparison with diversity-aware retrievers (MMR, k -DPP). Two natural set-level diversification baselines are MMR (Carbonell and Goldstein, 1998), which uses greedy pairwise dissimilarity, and k -Determinantal Point Processes (Kulesza and Taskar, 2012), which sample subsets in proportion to the determinant of a kernel submatrix. Both use the same dense encoder as Vendi-RAG to keep the comparison apples-to-apples. Table 4 shows that MMR and k -DPP land in the same range and are themselves competitive iterative-loop alternatives to Adaptive-RAG, yet Vendi-RAG (dynamic) still improves over the better of the two by +1.2 EM and +2.4 Acc on HotpotQA, +1.4 EM and +2.4 Acc on 2WikiMultiHopQA, and +1.8 EM and +1.6 Acc on MuSiQue. The Vendi Score’s set-level eigenvalue formulation differs from the pairwise-distance objective of MMR and from the log-determinant objective of DPPs: it directly mea-

Methods	MuSiQue			HotpotQA			2WikiMultiHopQA		
	EM	F1	Acc	EM	F1	Acc	EM	F1	Acc
No Retrieval	20.40±0.5	31.30±0.6	24.40±0.4	37.40±0.7	51.04±0.8	43.20±0.6	37.00±0.5	48.50±0.7	43.40±0.5
Single-step Approach	16.40±0.6	26.70±0.7	23.60±0.5	39.60±0.8	50.44±0.9	45.60±0.7	46.80±0.6	57.69±0.8	52.60±0.6
Adaptive Retrieval	18.80±0.5	30.30±0.6	24.80±0.4	38.60±0.7	50.70±0.8	43.20±0.6	44.20±0.5	55.11±0.7	50.60±0.5
IRCoT	22.60±0.6	31.50±0.7	28.40±0.5	41.20±0.8	55.30±0.9	55.80±0.7	46.00±0.6	57.40±0.8	58.20±0.6
LightRAG	17.40±0.6	27.51±0.7	23.80±0.5	37.40±0.8	47.30±0.9	42.30±0.7	42.21±0.6	52.24±0.8	47.50±0.6
Adaptive-RAG	21.80±0.6	32.60±0.7	29.60±0.5	40.40±0.8	52.56±0.9	47.00±0.7	46.60±0.6	60.09 ±0.8	56.80±0.6
Self-RAG [‡]	1.20±0.3	8.20±0.5	11.80±0.4	5.60±0.4	17.86±0.6	30.60±0.5	3.00±0.3	19.14±0.5	39.00±0.5
Adaptive Retrieval + MMR	22.6±0.5	31.0±0.6	28.8±0.4	41.0±0.7	55.0±0.8	56.0±0.6	45.8±0.5	57.0±0.7	59.0±0.5
Adaptive Retrieval + k -DPP	22.4±0.5	31.2±0.6	28.6±0.4	40.8±0.7	54.8±0.8	55.6±0.6	45.6±0.5	56.8±0.7	58.6±0.5
Adaptive Retrieval + Cross-Enc. Rerank [†]	—	—	27.1±0.4	—	—	54.2±0.6	—	—	57.6±0.5
Vendi-RAG(fixed $s = 0.0$)	22.4±0.5	30.4±0.6	28.2±0.4	40.6±0.7	54.1±0.8	55.4±0.6	45.4±0.5	56.8±0.7	58.8±0.5
Vendi-RAG*(fixed $s = 0.3$)	22.4±0.5	31.1±0.6	28.6±0.4	40.7±0.7	55.0±0.8	55.7±0.6	45.3±0.5	56.8±0.7	59.1±0.5
Vendi-RAG(fixed $s = 0.8$)	23.0±0.5	31.2±0.6	30.2±0.4	42.0±0.7	56.9±0.8	58.0±0.6	47.0±0.5	58.7±0.7	61.0±0.5
Vendi-RAG(fixed $s = 1.0$)	22.9±0.5	31.0±0.6	29.4±0.4	41.6±0.7	55.8±0.8	57.0±0.6	46.4±0.5	57.8±0.7	60.2±0.5
Vendi-RAG(dynamic, $s_1 = 0.8$)	24.4±0.5*	32.52±0.6	30.4±0.4*	42.2±0.7*	57.04±0.8*	58.4±0.6*	47.2±0.5*	58.9±0.7	61.4±0.5*

Table 4: Performance on multi-hop QA datasets using GPT-3.5 Turbo. Results are mean $\pm$ std over 3 seeds. * indicates paired-bootstrap significance (10,000 resamples, $p < 0.05$ after Holm–Bonferroni correction across the 9-cell family) against the strongest non-Vendi comparator on that specific cell (Appendix D.4). [‡] Self-RAG is designed for single-hop retrieval with self-reflection; its multi-hop numbers reflect this design mismatch (see baseline description) and are included for completeness only. [†] The cross-encoder reranking baseline (BAAI/bge-reranker-base over the same top-50 pool) was run as a single-step controlled comparison and is evaluated on accuracy only. Vendi-RAG ($s_1 = 0.8$, dynamic) achieves the best EM and Acc on all three datasets; F1 is competitive, with Vendi-RAG highest on HotpotQA. IRCoT, MMR, and k -DPP are obtained under the same retriever, k , and test splits as Vendi-RAG and occupy the same performance band as the other diversity-aware baselines.

Backbone	Method	HotpotQA Acc
Llama 3-70B	Adaptive-RAG	46.4±0.6
	Adapt. Retr. + MMR	48.3±0.6
	Vendi-RAG	50.8±0.5
Gemini Flash-2	Adaptive-RAG	48.1±0.6
	Adapt. Retr. + MMR	49.9±0.6
	Vendi-RAG	52.5±0.5

Table 5: Cross-architecture accuracy on HotpotQA. Vendi-RAG (dynamic, $s_1 = 0.8$) outperforms Adaptive-RAG and the stronger diversity-aware MMR baseline. Per-dataset EM and F1 are in Appendix D.6.

sure the effective number of distinct documents, and the diversity-relevance trade-off is controlled through a single scalar s . A full MMR λ sweep (Appendix D.11) shows this gap is not an artifact of suboptimal λ . Finally, since Vendi-RAG selects from a top-50 pool, part of its advantage could come from re-ranking that pool at all rather than from diversity; a strong relevance-only cross-encoder reranker over the identical pool lands well above cosine-similarity selection but still below both MMR and Vendi-RAG on all three datasets (Appendix D.2), so better selection alone recovers only part of the gain.

Robustness across LLM backbones. To check that the gains are not specific to a single model family, we evaluate Vendi-RAG (with $s_1 = 0.8$) against Adaptive-RAG across five backbones spanning three model families: GPT-3.5, GPT-4o, GPT-4o-mini (Figure 3), the open-weight Llama 3-70B-

Instruct, and Gemini Flash 2.0. Table 5 summarizes HotpotQA accuracy for the two new architectures; full per-dataset numbers are in Appendix D.6. Vendi-RAG yields a +4.4 accuracy gain over Adaptive-RAG on both Llama 3-70B and Gemini Flash 2.0, comparable to the gains on the GPT family. Because Adaptive-RAG is not the strongest baseline on HotpotQA (Table 4), we also extend the MMR comparison to these two backbones so that the cross-backbone claim is tested against the strongest competitor rather than a weak one. MMR sits above Adaptive-RAG on both (48.3 and 49.9) but below Vendi-RAG, leaving a +2.5 and +2.6 margin that closely matches the +2.4 MMR gap measured at GPT-3.5. The headline result therefore holds against the stronger baseline across model families, not only against the routing-based one. Larger backbones (GPT-4o) achieve the strongest absolute performance, and the relative advantage of Vendi-RAG holds with smaller models like GPT-4o-mini and with open-weight architectures.

Effect of the retrieved-set size. Figure 2 sweeps the number of retrieved documents k on HotpotQA. Vendi-RAG outperforms Adaptive-RAG on every metric for $k \geq 3$ and the gap widens with k , with $s_1 = 0.8$ the best configuration throughout; the retrieved-set Vendi Score rises with k as well, and because the VS is invariant under duplication (Friedman and Dieng, 2023) this is direct evidence that the added documents are non-redundant




Figure 2: Performance of Vendi-RAG and Adaptive-RAG across document set sizes $k \in \{2, 3, 5, 8, 10, 15\}$, measured on HotpotQA with the GPT-4o-mini backbone (mean over 3 seeds). Vendi-RAG variants are labeled by their fixed initialization value $s_1 \in \{0.3, 0.8, 1.0\}$. Vendi-RAG outperforms Adaptive-RAG across all metrics for $k \geq 3$ and the gap widens with k ; the $s_1 = 0.8$ variant is the best Vendi-RAG configuration. Beyond $k = 10$ the context window starts to require document truncation, which introduces a confound for the largest- k points.

rather than repeated. Performance plateaus around $k = 8$ – 10 , beyond which context-window limits force truncation. Full discussion is in Appendix D.9.

Retriever robustness. Downstream performance can depend strongly on the first-stage retriever, so we repeat the HotpotQA experiment with a different dense encoder (BGE-large, Xiao et al., 2024) in place of all-mpnet-base-v2. With BGE-large, Adaptive-RAG reaches 48.7 accuracy and Vendi-RAG (dynamic, $s_1 = 0.8$) reaches 59.2; the +10.5 gain is in the same range as the +11.4 obtained with all-mpnet-base-v2 (Table 4), and the relative ordering of all baselines is preserved (full numbers in Appendix D.10). For the MMR baseline, sweeping $\lambda \in \{0.3, 0.5, 0.7\}$ on HotpotQA yields 54.8, 56.0, and 55.2 accuracy respectively (Appendix D.11); the best MMR setting is still 2.4 accuracy points below Vendi-RAG, so the gap is not an artifact of MMR hyperparameter tuning.

5 Conclusion

Traditional RAG pipelines optimize for relevance alone and tend to retrieve semantically redundant documents, which is a poor match for multi-hop reasoning that depends on combining complementary evidence. We introduced Vendi-RAG, a framework that elevates retrieval diversity to a

first-class objective via the Vendi Retrieval Score and adapts the diversity-relevance balance with an explore–exploit feedback loop driven by an LLM-based answer-quality signal. Across HotpotQA, MuSiQue, and 2WikiMultiHopQA with GPT-3.5, Vendi-RAG improves accuracy by +2.4, +1.6, and +2.4 points over the strongest diversity-aware iterative baseline (Adaptive Retrieval + MMR); the corresponding gaps against routing-based Adaptive-RAG are +11.4, +0.8, and +4.6, with the largest gap explained by the routing behavior described in Section 4. A fixed- s variant of Vendi-RAG captures most of the improvement at no extra inference cost over MMR; the LLM-judge loop adds a smaller, statistically significant gain on hard queries at higher latency, giving practitioners a tunable accuracy–cost trade-off. The pattern is stable across five instruction-tuned LLM backbones (GPT-3.5, GPT-4o, GPT-4o-mini, Llama 3-70B-Instruct, Gemini Flash 2.0) and two dense encoders on multi-hop English QA, indicating that the gain is driven by the retrieval objective rather than by any individual model in our evaluation suite. The framework is composable: Vendi retrieval is a scoring function that can be combined with retrieval-loop methods (IRCoT, Adaptive-RAG) and with post-retrieval utility selectors.

Limitations

Our study is restricted to English multi-hop QA over three benchmarks (HotpotQA, MuSiQue, 2WikiMultiHopQA), five instruction-tuned backbones (GPT-3.5, GPT-4o, GPT-4o-mini, Llama 3-70B-Instruct, Gemini Flash 2.0), and two dense retrievers (all-mpnet-base-v2, BGE-large). Transfer to non-English benchmarks, tasks where complementary evidence matters less (single-hop QA, factoid lookup), and base (non-instruction-tuned) models is untested, so our consistency across this suite is not a guarantee of universal model-agnosticism; the suite also predates current frontier reasoning models, which may write better query refinements on their own and shrink the headroom diversity-aware retrieval exploits. Vendi-RAG also inherits the standard RAG dependence on knowledge-base quality and coverage. The dynamic LLM-judge loop adds compute on hard queries: latency rises from 5.8s to 8.4s and tokens from 3.4K to 5.1K per query (Table 7), mostly from judge calls rather than the eigendecomposition, for a dollar overhead of \$2–\$3 per 1,000 queries at current GPT-3.5 pricing (Appendix C.6). The fixed- s variant matches MMR’s iteration count and token budget, drops judge calls, and retains most of the gain (+2.0 of +2.4 HotpotQA points over MMR), making it preferable for latency- or cost-critical settings. The dynamic loop’s additional gain over fixed- $s=0.8$ is significant on HotpotQA ($p < 0.05$) but borderline on MuSiQue and 2WikiMultiHopQA (Appendix D.4); it is most valuable when the iteration budget is fully used. The loop relies on the LLM judge providing a useful *relative* quality signal; the cross-judge control shifts accuracy by under 0.5 points (Appendix D.7), but we do not measure inter-judge agreement or calibration directly, and judges may behave differently on distant domains. The VRS is also non-submodular and optimized greedily, so selection carries no approximation guarantee beyond the small measured gap (Appendix E.5), and iterations do not carry forward previously retrieved documents. Future work includes a direct judge-reliability study, a learned reward model in place of the judge, joint optimization of s with per-hop retrieval planning, and Vendi retrieval as an agentic-framework tool.

Ethical Considerations

Vendi-RAG uses LLMs both for answer generation and as a judge for answer quality. The same ethi-

cal concerns that apply to LLM-based QA systems apply here: outputs can reflect biases present in the underlying model and in the retrieved knowledge base, and grounded retrieval cannot rule out subtly biased or factually outdated content. The diversity-aware retrieval mechanism itself does not introduce new categories of risk: it operates over a candidate pool that is first filtered by relevance, so it can only redistribute attention across documents that have already passed the relevance filter; it does not synthesize new content. Faithfulness is monitored externally via an NLI-based metric (Appendix D.8). We recommend that downstream deployments verify the knowledge base for coverage and bias and pair Vendi-RAG with the safety filters and provenance tracking that are standard for production RAG systems.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. Self-rag: Learning to retrieve, generate, and critique through self-reflection. *arXiv preprint arXiv:2310.11511*.
- Reyhane Askari Hemmat, Melissa Hall, Alicia Sun, Candace Ross, Michal Drozdal, and Adriana Romero-Soriano. 2024. Improving geo-diversity of generated images with contextualized vendi score guidance. In *European Conference on Computer Vision*, pages 213–229. Springer.
- Sebastian Berns, Simon Colton, and Christian Guckelsberger. 2023. Towards Mode Balancing of Generative Models via Diversity Weights. *arXiv preprint. ArXiv:2304.11961 [cs.LG]*.
- Utkarsh Bhardwaj, Vinayak Mishra, Suman Mondal, and Manoj Warriar. 2025. A robust machine learned interatomic potential for nb: Collision cascade simulations with accurate defect configurations. *arXiv preprint arXiv:2502.03126*.
- Jaime Carbonell and Jade Goldstein. 1998. The use of mmr, diversity-based reranking for reordering documents and producing summaries. In *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 335–336.
- Laming Chen, Guoxin Zhang, and Eric Zhou. 2018. Fast greedy map inference for determinantal point process to improve recommendation diversity. *Advances in neural information processing systems*, 31.

- Van Dang and W. Bruce Croft. 2012. Diversity by proportionality: An election-based approach to search result diversification. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 65–74, Portland, Oregon, USA. ACM.
- Dan Friedman and Adji Bousso Dieng. 2023. The Vendi Score: A Diversity Evaluation Metric for Machine Learning. *Transactions on Machine Learning Research*.
- Boqing Gong, Wei-Lun Chao, Kristen Grauman, and Fei Sha. 2014. Diverse sequential subset selection for supervised video summarization. In *Advances in Neural Information Processing Systems 27*, pages 2069–2077.
- Zirui Guo, Lianghao Xia, Yanhua Yu, Tu Ao, and Chao Huang. 2025. **LightRAG: Simple and fast retrieval-augmented generation**. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 10746–10761, Suzhou, China. Association for Computational Linguistics.
- Bernal J Gutiérrez, Yiheng Shu, Yu Gu, Michihiro Yasunaga, and Yu Su. 2024. Hipporag: Neurobiologically inspired long-term memory for large language models. *Advances in neural information processing systems*, 37:59532–59569.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020. Retrieval augmented language model pre-training. In *International conference on machine learning*, pages 3929–3938. PMLR.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. Constructing a multi-hop qa dataset for comprehensive evaluation of reasoning steps. *arXiv preprint arXiv:2011.01060*.
- Sture Holm. 1979. A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, pages 65–70.
- Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansy, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim, and Yossi Matias. 2022. True: Re-evaluating factual consistency evaluation. *arXiv preprint arXiv:2204.04991*.
- Akriti Jain and Aparna Garimella. 2025. Modeling contextual passage utility for multihop question answering. In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 464–471.
- Mohammad Jalali, Azim Ospanov, Amin Gohari, and Farzan Farnia. 2024. Conditional vendi score: An information-theoretic approach to diversity evaluation of prompt-based generative models. *arXiv preprint arXiv:2411.02817*.
- Soyeong Jeong, Jinheon Baek, Sukmin Cho, Sung Ju Hwang, and Jong C Park. 2024. Adaptive-rag: Learning to adapt retrieval-augmented large language models through question complexity. *arXiv preprint arXiv:2403.14403*.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan Arik, Dong Wang, Hamed Zamani, and Jiawei Han. 2025. Search-r1: Training llms to reason and leverage search engines with reinforcement learning. *arXiv preprint arXiv:2503.09516*.
- Nithish Kannen, Arif Ahmad, Marco Andreetto, Vinodkumar Prabhakaran, Utsav Prabhu, Adji Bousso Dieng, Pushpak Bhattacharyya, and Shachi Dave. 2024. Beyond aesthetics: Cultural competence in text-to-image models. *arXiv preprint arXiv:2407.06863*.
- Omar Khattab, Keshav Santhanam, Xiang Lisa Li, David Hall, Percy Liang, Christopher Potts, and Matei Zaharia. 2022. Demonstrate-search-predict: Composing retrieval and language models for knowledge-intensive nlp. *arXiv preprint arXiv:2212.14024*.
- Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. 2022. Decomposed prompting: A modular approach for solving complex tasks. *arXiv preprint arXiv:2210.02406*.
- Alex Kulesza and Ben Taskar. 2012. Determinantal point processes for machine learning. *Foundations and Trends in Machine Learning*, 5(2-3):123–286.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and Zhicheng Dou. 2025. Search-o1: Agentic search-enhanced large reasoning models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 5420–5438.
- Hui Lin and Jeff Bilmes. 2011. A class of submodular functions for document summarization. In *Proceedings of the 49th Annual Meeting of the Association for*

- Computational Linguistics: Human Language Technologies*, pages 510–520, Portland, Oregon, USA. Association for Computational Linguistics.
- Jingyu Liu, Jiaen Lin, and Yong Liu. 2024a. How much can rag help the reasoning of llm? *arXiv preprint arXiv:2410.02338*.
- Tsung-Wei Liu, Quan Nguyen, Adji Bousso Dieng, and Diego A Gómez-Gualdrón. 2024b. Diversity-driven, efficient exploration of a mof design space to optimize mof properties. *Chemical Science*, 15(45):18903–18919.
- Mohsen Mousavi and Nasser Khalili. 2025. Vsi: An interpretable bayesian feature selection method based on vendi score. *Knowledge-Based Systems*, 311:112973.
- Quan Nguyen and Adji Bousso Dieng. 2024. Quality-Weighted Vendi Scores And Their Application To Diverse Experimental Design. In *International Conference on Machine Learning*.
- Hariom A Pandya and Brijesh S Bhatt. 2021. Question answering survey: Directions, challenges, datasets, evaluation matrices. *arXiv preprint arXiv:2112.03572*.
- Amey P Pasarkar, Gianluca M Bencomo, Simon Olsson, and Adji Bousso Dieng. 2023. Vendi sampling for molecular simulations: Diversity as a force for faster convergence and better exploration. *The Journal of chemical physics*, 159(14).
- Amey P Pasarkar and Adji Bousso Dieng. 2023. Cousins of the vendi score: A family of similarity-based diversity metrics for science and machine learning. *arXiv preprint arXiv:2310.12952*.
- Amey P Pasarkar and Adji Bousso Dieng. 2025. The vendiscope: An algorithmic microscope for data collections. *arXiv preprint arXiv:2502.10828*.
- Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. 2019. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*.
- Marc Pickett, Jeremy Hartman, Ayan Kumar Bhowmick, Raquib-ul Alam, and Aditya Vempaty. 2024. Better rag using relevant information gain. *arXiv preprint arXiv:2407.12101*.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah A Smith, and Mike Lewis. 2023. Measuring and narrowing the compositionality gap in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5687–5711.
- Ansh Radhakrishnan, Karina Nguyen, Anna Chen, Carol Chen, Carson Denison, Danny Hernandez, Esin Durmus, Evan Hubinger, Jackson Kernion, Kamile Lukošiuūtė, et al. 2023. Question decomposition improves the faithfulness of model-generated reasoning. *arXiv preprint arXiv:2307.11768*.
- Mohaimenul Azam Khan Raiaan, Md Saddam Hossain Mukta, Kaniz Fatema, Nur Mohammad Fahad, Saddam Sakib, Most Marufatul Jannat Mim, Jubaer Ahmad, Mohammed Eunus Ali, and Sami Azam. 2024. A review on large language models: Architectures, applications, taxonomies, open issues and challenges. *IEEE Access*.
- Mohammad R Rezaei, Reza Saadati Fard, Jayson Parker, Rahul G Krishnan, and Milad Lankarany. 2025. Agentic medical knowledge graphs enhance medical question answering: Bridging the gap between llms and evolving medical knowledge. In *EMNLP (Findings)*, pages 12682–12701.
- Mohammad Reza Rezaei and Adji Bousso Dieng. 2025. The α -alternator: Dynamic adaptation to varying noise levels in sequences using the vendi score for improved robustness and performance. *arXiv preprint arXiv:2502.04593*.
- Rodrygo L. T. Santos, Jie Peng, Craig Macdonald, and Iadh Ounis. 2010. Explicit search result diversification through sub-queries. In *Advances in Information Retrieval, 32nd European Conference on IR Research, ECIR 2010*, pages 87–99, Milton Keynes, UK. Springer.
- Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. 2023. Enhancing retrieval-augmented large language models with iterative retrieval-generation synergy. *arXiv preprint arXiv:2305.15294*.
- Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. Retrieval augmentation reduces hallucination in conversation. *arXiv preprint arXiv:2104.07567*.
- Aleksandrs Slivkins, Filip Radlinski, and Sreenivas Golapudi. 2010. Learning optimally diverse rankings over large document collections. In *Proc. of the 27th International Conference on Machine Learning (ICML 2010)*.
- Zhiqing Sun, Xuezhi Wang, Yi Tay, Yiming Yang, and Denny Zhou. 2022. Recitation-augmented language models. *arXiv preprint arXiv:2210.01296*.
- Yixuan Tang and Yi Yang. 2024. Multihop-rag: Benchmarking retrieval-augmented generation for multihop queries. *arXiv preprint arXiv:2401.15391*.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. Beir: A heterogenous benchmark for zero-shot evaluation of information retrieval models. *arXiv preprint arXiv:2104.08663*.

- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. MuSiQue: Multi-hop questions via single-hop question composition. *Transactions of the Association for Computational Linguistics*, 10:539–554.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2023. Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 10014–10037.
- Xuezhi Wang and Denny Zhou. 2024. Chain-of-thought reasoning without prompting. *arXiv preprint arXiv:2402.10200*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Kai Wei, Rishabh Iyer, and Jeff Bilmes. 2014. Submodular optimization and machine learning: Theoretical advances, connections, and applications (tutorial). In *Tutorial at the 31st International Conference on Machine Learning (ICML)*.
- Haolun Wu, Yansen Zhang, Chen Ma, Fuyuan Lyu, Bowei He, Bhaskar Mitra, and Xue Liu. 2024. Result diversification in search and recommendation: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 36(10):5354–5373.
- Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. 2024. C-pack: Packed resources for general Chinese embeddings. *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- Le Yan, Zhen Qin, Rama Kumar Pasumarthi, Xuanhui Wang, and Michael Bendersky. 2021. Diversification-aware learning to rank using distributed representation. In *Proceedings of the Web Conference 2021*, pages 127–136.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2369–2380.
- Tian Yu, Shaolei Zhang, and Yang Feng. 2024. Auto-rag: Autonomous retrieval-augmented generation for large language models. *arXiv preprint arXiv:2411.19443*.
- Hengran Zhang, Keping Bi, Jiafeng Guo, Jiaming Zhang, Shuaiqiang Wang, Dawei Yin, and Xueqi Cheng. 2025. Distilling a small utility-based passage selector to enhance retrieval-augmented generation. In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 22–30.
- Jiahao Zhang, Haiyang Zhang, Dongmei Zhang, Liu Yong, and Shen Huang. 2024a. End-to-end beam retrieval for multi-hop question answering. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1718–1731.
- Zhihao Zhang, Alan Zhu, Lijie Yang, Yihua Xu, Lanting Li, Phitchaya Mangpo Phothilimthana, and Zhihao Jia. 2024b. Accelerating retrieval-augmented language model serving with speculation. *arXiv preprint arXiv:2401.14021*.

A Datasets

MuSiQue evaluates a model’s ability to synthesize facts spread across multiple document sources. It includes questions spanning diverse domains such as history, science, and culture, requiring logical reasoning and synthesis of interdependent information. Given its emphasis on multi-step comprehension, this dataset challenges models to accurately identify and integrate relevant information to generate correct answers to queries. This is the most challenging dataset among the three.

HotpotQA assesses reasoning and fact verification across various domains, including geography, entertainment, and history. Its questions necessitate reasoning over two or more interconnected documents linked via hyperlinks. Additionally, the dataset includes “comparison” questions that require juxtaposing information from multiple sources, making it a challenging benchmark for evaluating both retrieval quality and reasoning ability.

2WikiMultiHopQA leverages Wikipedia’s complex structure to pose complex reasoning challenges. Questions are derived from real-world knowledge graphs and require navigating reasoning paths across multiple documents. Topics span science, politics, and sports, emphasizing logical relationships such as cause-effect dependencies, making it an essential tool for evaluating structured knowledge reasoning.

B Evaluation Metrics

To compare model performance across different datasets, we employ the following key evaluation metrics:

- **Exact Match (EM):** This metric calculates the percentage of predictions that exactly

match the ground truth answers. It is defined as:

$$\text{EM} = \frac{\text{Number of exact matches}}{\text{Total number of queries}} \times 100 \quad (4)$$

EM is a strict metric that grants credit only when the predicted answer matches the ground truth exactly in both content and format. It is particularly useful for assessing a model’s precision in generating accurate responses.

- **F1 Score (F1):** The F1 score captures the harmonic mean of precision and recall at the token level, providing a balanced measure of correctness. It is defined as:

$$\text{F1} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (5)$$

where precision is the fraction of retrieved tokens that are relevant, and recall is the fraction of relevant tokens that are retrieved. The F1 score is particularly relevant for multi-hop QA tasks, where partial correctness (e.g., retrieving some but not all supporting evidence) is informative.

- **Accuracy (Acc):** Accuracy measures the proportion of correct predictions over all evaluated queries. It is defined as:

$$\text{Acc} = \frac{\text{Number of correct predictions}}{\text{Total number of queries}} \times 100 \quad (6)$$

Unlike EM, which requires exact matches, accuracy provides a broader assessment by capturing overall correctness, including responses that convey the intended meaning.

- **Max Pairwise Distance (MPD):** This metric evaluates the maximum Euclidean distance between pairs of retrieved data points, measuring diversity. It is defined as:

$$\text{MPD} = \max_{i,j} \|x_i - x_j\|_2, \quad i \neq j \quad (7)$$

where x_i and x_j represent document embeddings in the feature space. Higher values indicate greater diversity among retrieved documents.

Each of these metrics offers a unique perspective on model performance. EM is a stringent measure of precision, F1 balances precision and recall, and

accuracy provides an overall correctness measure. Meanwhile, MPD and diversity-based metrics assess the variety and independence of retrieved documents, critical for multi-hop QA tasks requiring integration of diverse information.

C Implementation Details

The Vendi-RAG framework employs dense vector representations derived from transformer-based embeddings to compute the similarity score $\text{SS}(q, \mathcal{D})$ between a query q and a set of documents $\mathcal{D}$. This high-dimensional semantic comparison allows the model to effectively capture contextual relationships and retrieve relevant documents across diverse domains. To evaluate answer quality, we use a consistent LLM backbone with a specifically designed prompt that positions the LLM as an expert judge; the full prompt is reproduced in Appendix E.6. The judge produces three integer scores on a 1–10 scale for coherence, relevance, and query alignment; we then divide their mean by 10 to obtain $Q_t \in [0, 1]$, and use threshold $\tau = 0.85$ (equivalently 8.5/10 on the raw scale) for all experiments. The initial diversity weight is fixed to $s_1 = 0.8$; the ablation studies in Tables 4 and 8 vary this hyperparameter and compare fixed and dynamic schedules.

C.1 Baseline Configurations

Every method uses the same dense encoder (all-mpnet-base-v2), the same corpus and chunking (Appendix C.2), the same GPT-3.5 Turbo answer generator and answer prompt, and the same test splits. Two quantities distinguish the retrieval configurations and are stated explicitly for every method in Table 6: the size of the *candidate pool* returned by first-stage similarity search, and the *final* number of documents k placed in the reader’s context. Methods that re-rank or diversify draw from a top-50 pool and pass $k = 5$ documents downstream; methods that do not re-rank retrieve $k = 5$ directly and have no intermediate pool. Only the retrieval component differs across rows.

- **No Retrieval:** The backbone answers from parametric knowledge alone, with the same answer prompt and no retrieved context.
- **Single-step Approach:** One similarity-search pass over the corpus returning the top $k = 5$ chunks by cosine similarity, concatenated into the answer prompt, with no iteration, no query rewriting, and no re-ranking. It differs from *Adaptive*

Retrieval only in that the latter runs inside the iterative harness (see below).

- **Adaptive Retrieval:** Single retrieval pass using cosine similarity, $k = 5$ documents, run inside the iterative harness with query rewriting. Despite the name (inherited from Jeong et al., 2024), its *scoring function* is plain relevance and is not adaptive; it is the relevance-only control against which the diversity-aware scoring functions are compared.
- **Adaptive-RAG (Jeong et al., 2024):** The routing method of Jeong et al. (2024), using their released query-complexity classifier (a T5-based classifier trained on their silver labels) to route each query to a no-retrieval, single-step, or iterative path. Only the routing is adaptive; within the iterative path, retrieval is relevance-only with $N = 5$ maximum iterations, quality threshold $\tau = 0.8$, the same retriever, and $k = 5$. This is why Adaptive-RAG can fall *below* always-iterating baselines on HotpotQA, where the classifier routes many queries to the shorter paths.
- **Adaptive Retrieval + MMR:** MMR (Carbonell and Goldstein, 1998) re-ranking with $\lambda = 0.5$ over the *same* top-50 candidate pool used by Vendi-RAG, selecting $k = 5$ documents, inside the same iterative harness ($N = 5$). MMR never sees the full corpus; it diversifies within the pool, exactly as the VRS does. A full λ sweep is in Appendix D.11.
- **Adaptive Retrieval + k -DPP:** k -DPP (Kulesza and Taskar, 2012) selection over the same top-50 pool with $k = 5$, using an RBF kernel on the same dense embeddings with relevance-scaled quality terms ($L_{ij} = q_i K_{ij} q_j$, q_i the query-document cosine similarity), MAP selection via greedy log-determinant maximization, inside the same iterative harness ($N = 5$).
- **Adaptive Retrieval + Cross-Encoder Rerank:** The same top-50 pool re-ranked by the BAAI/bge-reranker-base cross-encoder (Xiao et al., 2024), keeping the top $k = 5$ by cross-encoder score. This is a relevance-only but much stronger selector than cosine similarity, and it isolates how much of the Vendi-RAG gain is attributable to better selection from the pool as opposed to diversity-aware selection.
- **IRCoT (Trivedi et al., 2023):** Interleaved retrieval with chain-of-thought reasoning, retrieving per reasoning step using the same dense encoder, the same top-50 pool, and $k = 5$. IRCoT varies the retrieval *loop* rather than the scoring

Method	Pool	Final k	Iters
No Retrieval	—	0	1
Single-step Approach	—	5	1
Adaptive Retrieval	—	5	1
Adaptive-RAG	—	5	up to 5 (routed)
Adapt. Retr. + MMR	50	5	up to 5
Adapt. Retr. + k -DPP	50	5	up to 5
Adapt. Retr. + Rerank	50	5	1
IRCoT	50	5	up to 5
Vendi-RAG	50	5	up to 5

Table 6: Retrieval configuration for every method: the candidate pool returned by first-stage similarity search, the final number of documents passed to the reader, and the iteration budget. “—” means the method retrieves its k documents directly with no intermediate pool. All rows share the same encoder, corpus, chunking, backbone, and test splits.

function and serves as a loop-level iterative baseline.

- **Self-RAG (Asai et al., 2023):** Official implementation with default hyperparameters. Self-RAG is designed for single-hop retrieval with on-demand self-reflection rather than for explicit multi-hop reasoning; under the multi-hop setup here its self-reflection tokens frequently suppress further retrieval and the resulting numbers (e.g., 1.2 EM on MuSiQue) reflect this design mismatch rather than a tuning failure, consistent with the multi-hop drop reported by Jeong et al. (2024). We include Self-RAG for completeness but the more informative comparators are the iterative diversity-aware baselines (MMR, k -DPP, IRCoT).
- **LightRAG:** Graph-based retrieval with default entity extraction settings, maximum 5 retrieved subgraphs (Guo et al., 2025).
- **Vendi-RAG:** Greedy VRS selection of $k = 5$ documents over the same top-50 pool, inside the same iterative harness ($N = 5$), with s either fixed or updated by Eq. (3) from $s_1 = 0.8$.

C.2 Dataset Processing and Chunking

Preparing datasets for question-answering requires transforming data into a searchable vector database to enable efficient retrieval. This workflow includes document chunking and semantic embedding to optimize performance.

The dataset, provided in JSON format with context paragraphs and metadata, is processed by splitting each document into smaller chunks. Each chunk has a maximum size of 512 tokens, with a 50-token overlap to preserve context across chunk

boundaries and facilitate multi-hop reasoning in long documents.

C.3 Embedding Model and Vector Database

We use the **SentenceTransformer** model, specifically `all-mpnet-base-v2`, to generate dense vector representations for documents and queries. These embeddings are stored locally to avoid redundant downloads and improve reusability. The **Chroma** vector database efficiently stores and retrieves these vectorized documents along with metadata, such as document titles and chunk IDs.

C.4 Batch Processing and Database Population

To efficiently populate the vector database, document chunks are processed in batches of up to 10,000. This approach optimizes memory usage while ensuring completeness in the ingestion process. The total number of processed chunks is logged for verification.

C.5 Query Answering Workflow

For queries such as *"Who is the father-in-law of Queen Hyojeong?"*, relevant chunks are retrieved using **Chroma**'s similarity-based search mechanism. The system ranks the top 10 chunks based on their semantic similarity to the query, leveraging embeddings generated by `all-mpnet-base-v2` to ensure precise and relevant results.

C.6 Computational cost analysis.

We provide a per-query cost breakdown on HotpotQA in Table 7. Three cost components are relevant. (i) The VS eigendecomposition is $O(k^3)$ per retrieval step *per query* with $k \leq 10$, which is $<5\text{ms}$ on the retrieved subset and is dwarfed by LLM calls; the greedy-selection budget of $O(k \cdot |\mathcal{C}| \cdot k^3)$ is therefore $<50\text{ms}$ per query for $|\mathcal{C}| = 50$. (ii) The iterative loop adds retrieve/reason calls beyond a single-step pass. (iii) The LLM judge calls add one extra forward pass per iteration. The dynamic- s variant of Vendi-RAG averages 3.2 iterations and 3.2 judge calls per query, against 2.8 iterations and no judge calls for Adaptive-RAG; this drives the latency from 5.8s to 8.4s ($\sim 45\%$ higher) and tokens from 3.4K to 5.1K. At GPT-3.5 API pricing as of writing ($\sim \$0.50$ per 1M input tokens, $\sim \$1.50$ per 1M output tokens, with judge prompts at roughly 80% input / 20% output), the dollar overhead of the judge loop is on the order of $\$2$ – $\$3$ per 1,000 queries on top of the Adaptive-RAG cost, which

is small relative to per-query end-user value in QA settings but non-trivial at very high traffic. The fixed- s variant of Vendi-RAG matches Adaptive-RAG iteration count and removes judge calls entirely, so its cost profile is the same as the iterative diversity-aware baselines (MMR, k -DPP, IRCOT) while retaining the bulk of the accuracy gain (Section 4, ablation discussion).

C.7 Key Configuration Details

The system is configured with the following parameters:

- **Embedding Model:** `all-mpnet-base-v2`, optimized for semantic similarity tasks.
- **Vector Database:** Chroma, persisted to disk for efficient reuse.
- **Chunk Size:** 512 tokens per chunk, with a 50-token overlap.
- **Batch Size:** Up to 10,000 chunks per batch.

D Additional Experiments and Analyses

D.1 Sensitivity Analysis of the VRS Process

To evaluate the robustness of the VRS process and understand its impact on retrieval diversity, we conducted a sensitivity analysis focusing on how varying the parameter s affects document ranking order within a vector database. This analysis elucidates the trade-off between retrieval precision and diversity, which is crucial for enhancing multi-hop reasoning performance. The analysis was performed using 100 randomly sampled queries to cover a diverse range of query types and complexity levels. We evaluated the retrieval pipeline across multiple s values from 0.0 to 1.0, incremented in small steps to capture granular variations. Setting $s = 0.0$ serves as a baseline representing pure similarity search, where retrieval relies exclusively on cosine similarity between embeddings, without any emphasis on diversity. This provides a reference point for measuring the impact of increasing s .

To quantify deviations from the baseline, we employed two complementary ranking comparison metrics:

- **Kendall's τ :** Measures rank order similarity between two lists by evaluating concordant and discordant pairs. Higher τ values indicate stronger similarity to the baseline, while lower values reflect greater diversity introduced by increasing s .

Method	Iterations	Judge Calls	Tokens	Latency (s)
Single-step	1.0	0	1.2K	2.1
Adaptive-RAG	2.8	0	3.4K	5.8
Adaptive Retrieval + MMR	3.0	0	3.6K	6.0
Vendi-RAG (fixed $s=0.8$)	3.0	0	3.6K	6.0
Vendi-RAG (dynamic, $s_1=0.8$)	3.2	3.2	5.1K	8.4

Table 7: Computational cost analysis on HotpotQA (per-query average, GPT-3.5 Turbo). The fixed- s Vendi-RAG variant has the same cost profile as MMR; the dynamic loop adds judge calls and increases latency by roughly 45% in exchange for the additional accuracy gain reported in Section 4.

s	τ	ρ	SF Recall
0.0	1.00	1.00	58.3
0.2	0.797	0.828	61.2
0.4	0.688	0.742	65.7
0.6	0.485	0.528	69.1
0.8	0.265	0.316	71.4
1.0	0.074	0.078	68.9

Table 8: Sensitivity Analysis of the VRS. Higher s values increase diversity (lower τ , ρ) while SF Recall peaks at $s = 0.8$.

- **Spearman’s Rank Correlation ρ :** Assesses the monotonic relationship between two rankings by considering both orderings and positional shifts. Lower ρ values signal substantial deviation from the baseline, indicating increased diversity at higher s values.

To connect rank changes to downstream performance, we additionally measure supporting facts recall at each s value in Table 8. As s increases from 0.0 to 1.0, both Kendall’s τ and Spearman’s ρ decrease progressively, showing that higher s values promote retrieval diversity. SF Recall peaks at $s = 0.8$, indicating that the rank changes induced by the diversity mechanism lead to better coverage of multi-hop evidence.

D.2 Cross-Encoder Reranking Control

Vendi-RAG selects $k = 5$ documents from a top-50 candidate pool, so part of its advantage over relevance-only retrieval could come from re-ranking that pool at all, independently of diversity. To separate the two, we run a strong relevance-only selector over the identical pool: the BAAI/bge-reranker-base cross-encoder (Xiao et al., 2024), keeping its top 5 passages and feeding them to the same GPT-3.5 reader. The cross-encoder improves substantially over cosine-

similarity selection (43.2 $\rightarrow$ 54.2 HotpotQA accuracy, 24.8 $\rightarrow$ 27.1 on MuSiQue, 50.6 $\rightarrow$ 57.6 on 2WikiMultiHopQA), confirming that better selection from the pool is worth several accuracy points on its own. It nonetheless remains below both MMR (56.0/28.8/59.0) and the single-step fixed- $s=0.8$ Vendi-RAG variant (58.0/30.2/61.0) on all three datasets. A stronger relevance signal therefore recovers only part of the gain, and set-level diversity adds a further consistent margin on top of it; the ordering (relevance < cross-encoder < pairwise diversity < set-level diversity) is the same on every dataset.

D.3 Isolating the Loop: Adaptation vs. the Stopping Rule

The LLM judge does two things at once—it rebalances s and it decides when to stop—and the dynamic variant also runs slightly more iterations on average than the fixed- s variant (3.2 vs. 3.0; Table 7), so the +0.4 gain of the dynamic loop over fixed- $s=0.8$ reported in Section 4 is not by itself attributable to s -adaptation. We therefore run the 2×2 factorial in Table 9 on HotpotQA, crossing s -adaptation (on/off) with the stopping rule (judge-controlled vs. fixed at three iterations). Conditions A and C run exactly three iterations for every query, which removes the iteration-count confound; B and D let the judge stop early (1–2 iterations on easy queries, 4–5 on hard ones). Comparing A with C isolates s -reweighting (+0.15) and A with B isolates the stopping rule (+0.20); the full +0.40 of condition D is these two effects plus a small positive interaction (+0.05), so the two mechanisms are complementary rather than redundant. Neither component alone accounts for the loop’s gain, and the effect sizes are small in absolute terms: on HotpotQA, fixed- $s=0.8$ is already close to the ceiling of this scoring function, which compresses the room any single component has in isolation. This

is consistent with, and does not change, the paper’s headline claim: the VRS scoring function at a fixed high diversity weight delivers +2.0 of the +2.4 points over MMR at no extra cost, and the loop is a smaller optional add-on that earns its judge cost only where the iteration budget is actually used.

D.4 Statistical Significance Tests

Significance markers (*) in Table 4 are obtained by a paired bootstrap with 10,000 resamples at the per-query level (test splits: 500 queries for MuSiQue, 500 for 2WikiMultiHopQA, 7,405 for HotpotQA). The strongest non-Vendi comparator differs per cell; Table 10 lists which baseline is used for each marker, identified as the row with the maximum value on that cell among the non-Vendi rows of Table 4. We use the strongest per-cell comparator rather than a single fixed baseline so that significance reflects improvement over the genuinely competitive method in each setting, not over an arbitrary one. Cells in which Vendi-RAG is not bold in Table 4 (MuSiQue F1 and 2WikiMultiHopQA F1, both topped by Adaptive-RAG) carry no * marker, since Vendi-RAG does not outperform the strongest comparator there.

Correcting for multiple comparisons. Testing against the per-cell strongest comparator is deliberately conservative: had we fixed a single baseline (say Adaptive-RAG) for all cells, several markers would have been trivially easy to obtain, because Adaptive-RAG is not the strongest baseline on most cells of Table 4. Testing against whichever non-Vendi method is empirically best on each cell means every * already survives the hardest available comparison. Choosing the comparator per cell does not, however, address *multiplicity*: Table 4 contains 9 cells (3 datasets $\times$ 3 metrics), and reporting an uncorrected p -value for each inflates the family-wise error rate. We therefore apply the Holm–Bonferroni step-down procedure (Holm, 1979) to the raw paired-bootstrap p -values within this 9-test family: sorting the raw p -values ascending, the i -th smallest is multiplied by $9 - i + 1$ and the sequence is enforced to be monotone, with significance declared at the corrected $\alpha = 0.05$. Table 11 reports raw and corrected values for the four cells with the smallest raw p -values, which by construction receive the largest Holm multipliers (9, 8, 7, 6) and are therefore the cells the correction acts on most aggressively. All four remain significant at $\alpha = 0.05$ after correction. MuSiQue

accuracy becomes marginal ($p = 0.048$), and we flag it as such rather than treating it as comfortably significant.

Dynamic- s vs. fixed- $s=0.8$. We also run the paired bootstrap for the dynamic- s row against the fixed- $s=0.8$ row (the strongest fixed variant), per query, on the accuracy metric. We obtain $p < 0.05$ on HotpotQA (where the iteration budget is used most heavily and the dynamic schedule visibly adjusts s_t), and $p \approx 0.08$ on MuSiQue and $p \approx 0.10$ on 2WikiMultiHopQA, where many queries terminate after ≤ 2 iterations and the dynamic schedule has fewer opportunities to act. The dynamic loop therefore delivers a small, consistent, and (on HotpotQA) statistically significant improvement on top of the gain from a strong fixed diversity weight; the fixed- s variant remains a competitive lower-cost option.

D.5 IRCOT under a Matched Setup

IRCoT (Trivedi et al., 2023) varies the retrieval loop (per-step retrieval triggers, no quality-based stopping) rather than the retrieval scoring function. To make the comparison apples-to-apples, we report it under a setup matched to ours: the same dense retriever (all-mpnet-base-v2), the same $k = 5$, and the same test splits (500 samples for MuSiQue and 2WikiMultiHopQA, 7,405 for HotpotQA), with GPT-3.5 Turbo as the backbone. Numbers under this setup appear in the main Table 4. Across all three datasets, IRCoT lands close to Adaptive Retrieval + MMR and Adaptive-RAG on EM and F1, and below the fixed- $s = 0.8$ Vendi-RAG variant on Acc. IRCoT’s per-step scheduling and the VRS scoring function operate on different axes and are composable: a per-step IRCoT trigger can call a VRS-based retriever instead of a similarity-only one, which is a natural direction for future work.

D.6 Robustness Across LLM Backbones

We evaluate Vendi-RAG (dynamic, $s_1 = 0.8$) against Adaptive-RAG using two architecturally distinct backbones beyond the GPT family: Llama 3-70B-Instruct (open-weight) and Gemini Flash 2.0 (proprietary non-GPT). Adaptive-RAG is not the strongest baseline on HotpotQA under our setup, so to avoid overstating the cross-backbone gain we additionally run Adaptive Retrieval + MMR—the strongest comparator at GPT-3.5—on both backbones, holding retriever, candidate pool, k , and

Condition	Adapt. s	Stop	Iters	Acc
A: fixed- s , fixed-iter	✗	fixed	3.0	58.00
B: fixed- s , judge-stop	✗	judge	3.2	58.20
C: adaptive- s , fixed-iter	✓	fixed	3.0	58.15
D: full system	✓	judge	3.2	58.40

Table 9: 2×2 factorial ablation of the Vendi-RAG loop on HotpotQA (GPT-3.5 Turbo), crossing s -adaptation with the stopping rule. “fixed” stopping runs exactly three iterations per query; “judge” stops when $Q_t \geq \tau$. “Iters” is the test-set mean. s -reweighting alone contributes +0.15 (A→C) and the stopping rule alone +0.20 (A→B), with a +0.05 interaction accounting for the remainder of the full +0.40 (A→D). All conditions start from $s_1 = 0.8$.

Dataset	EM	F1	Acc
MuSiQue	IRCoT/MMR (22.6)	—	Adaptive-RAG (29.6)
HotpotQA	IRCoT (41.2)	IRCoT (55.3)	MMR (56.0)
2WikiMultiHopQA	Adaptive-RAG (46.6)	—	MMR (59.0)

Table 10: Per-cell strongest non-Vendi comparator used to compute the significance markers in Table 4. “—” indicates a cell where Vendi-RAG does not exceed the strongest comparator (Adaptive-RAG on both F1 columns) and therefore receives no *.

Cell	Raw p	Holm p	Sig.
HotpotQA Acc	0.0013	0.012	✓
HotpotQA EM	0.0024	0.019	✓
2Wiki Acc	0.0049	0.034	✓
MuSiQue Acc	0.0080	0.048	✓ (marginal)

Table 11: Raw and Holm–Bonferroni-corrected paired-bootstrap p -values (10,000 resamples, per-query, against the per-cell strongest non-Vendi comparator of Table 10), corrected within the 9-test family of Table 4. Listed are the four cells with the smallest raw p -values, which receive the largest Holm multipliers. The two F1 cells topped by Adaptive-RAG are not tested but count toward the family size.

iteration budget fixed. Table 12 reports HotpotQA results.

The accuracy gains on both Llama 3-70B and Gemini Flash 2.0 are comparable to those on the GPT family, indicating that the diversity-aware retrieval mechanism transfers across model architectures. Crucially, the ordering Adaptive-RAG < MMR < Vendi-RAG reproduces on both backbones, and the Vendi-RAG-over-MMR margin (+2.5, +2.6) is essentially the same as at GPT-3.5 (+2.4), so the reported advantage is not an artifact of comparing against the routing-based baseline.

D.7 Judge Robustness and Self-Reinforcement

What the judge can and cannot affect. Same-LLM generation and judging could in principle introduce self-reinforcement bias, and an LLM judge is in general imperfectly calibrated. It is therefore worth being precise about the judge’s

scope in this pipeline, because it bounds what a miscalibrated judge can damage. Three properties limit the risk: (i) the judge evaluates only structural properties (coherence, relevance, query alignment) and influences only the diversity parameter s_t —it never edits, rewrites, or filters the answer text; (ii) factual grounding is assessed by an *external* NLI-based faithfulness metric (Appendix D.8), on which Vendi-RAG attains the highest score of all methods (0.86 vs. 0.82–0.84); (iii) bounded iterations and EMA smoothing prevent any single score from dominating the trajectory. A miscalibrated judge can consequently misallocate retrieval effort across iterations—spending an extra round, or stopping one round early—but cannot inject unfaithful content into the final answer. This is a weaker dependence than in pipelines where a judge selects or rewrites the output.

Cross-judge control. A control experiment on HotpotQA with a different-backbone judge changes final accuracy by less than 0.5 points, indicating that same-model judging does not meaningfully bias the results. This motivates the design choice discussed in Section 3: the update rule requires only a *relative* quality signal, not an absolutely calibrated one, which is a substantially weaker requirement and explains why the loop tolerates a different or smaller judge.

Scope of this evidence. The cross-judge control is an end-to-end proxy: it shows that the final accuracy is insensitive to which model judges, which is the property the loop depends on. It is not a

Method	Backbone	EM	F1	Acc
Adaptive-RAG	Llama 3-70B	39.8 $\pm$ 0.7	51.9 $\pm$ 0.8	46.4 $\pm$ 0.6
Adapt. Retr. + MMR	Llama 3-70B	—	—	48.3 $\pm$ 0.6
Vendi-RAG	Llama 3-70B	43.1$\pm$0.6	56.2$\pm$0.7	50.8$\pm$0.5
Adaptive-RAG	Gemini Flash 2.0	41.2 $\pm$ 0.7	53.4 $\pm$ 0.8	48.1 $\pm$ 0.6
Adapt. Retr. + MMR	Gemini Flash 2.0	—	—	49.9 $\pm$ 0.6
Vendi-RAG	Gemini Flash 2.0	44.6$\pm$0.6	57.8$\pm$0.7	52.5$\pm$0.5

Table 12: Cross-architecture results on HotpotQA. Vendi-RAG improves accuracy by +4.4 points over Adaptive-RAG and by +2.5/+2.6 points over Adaptive Retrieval + MMR for Llama 3-70B-Instruct and Gemini Flash 2.0 respectively. The MMR rows were run for the cross-backbone control on accuracy only; “—” marks metrics not computed.

Method	Faithfulness
Adaptive Retrieval	0.82 $\pm$ 0.02
Adaptive-RAG	0.84 $\pm$ 0.02
Vendi-RAG	0.86$\pm$0.01

Table 13: Faithfulness scores (NLI-based) on HotpotQA.

direct measurement of inter-judge agreement or of agreement with human raters, and we do not claim one. Because the update rule consumes Q_t only through comparisons against the running best, agreement on the *ordering* of answers is what matters rather than agreement on absolute scores, and the accuracy-delta control speaks to the former only indirectly. A direct agreement and calibration study of the judge is left to future work (Limitations).

D.8 Faithfulness Evaluation

To assess whether generated answers are grounded in retrieved evidence, we compute faithfulness scores using an NLI-based metric following (Honovich et al., 2022). Table 13 shows that Vendi-RAG maintains high faithfulness while improving answer quality, indicating that the diversity-driven retrieval does not introduce unsupported claims.

D.9 Impact of the Number of Retrieved Documents

To examine the impact of document set size on retrieval effectiveness, we analyze Vendi-RAG and Adaptive-RAG across varying document sizes and initial settings $s_1 = \{0.3, 0.8, 1.0\}$. Figure 2 illustrates the relationship between document size and performance on HotpotQA. Vendi-RAG consistently outperforms Adaptive-RAG in accuracy for document sizes greater than two with any s_1 setting. As document size increases, accuracy improves for both methods, but the gain is notably higher for

Vendi-RAG. EM and F1 scores exhibit a similar increasing trend, with Vendi-RAG showing a more pronounced improvement. The VS also increases with document size: since the VS is invariant under duplication (Friedman and Dieng, 2023), an increasing VS indicates less redundancy among the retrieved documents, evidence that Vendi-RAG alleviates the redundancy that often plagues RAG pipelines. Performance plateaus around $k = 8-10$; beyond that, context-window constraints require document truncation, which introduces confounds. We therefore extend the analysis to $k = 15$ in this appendix figure but caution that the largest- k points are affected by truncation.

D.10 First-Stage Retriever Robustness

To check that the gains are not tied to a particular dense encoder, we replace all-mpnet-base-v2 with BGE-large (Xiao et al., 2024) and rerun HotpotQA with GPT-3.5 Turbo, keeping all other hyperparameters fixed. Table 14 reports the results. The relative ordering of methods is preserved and Vendi-RAG retains a +10.5 accuracy gap over Adaptive-RAG with BGE-large, comparable to the +11.4 gap with all-mpnet-base-v2. The absolute numbers shift upward by roughly one to two accuracy points for every method, reflecting BGE-large’s stronger candidate pool, but the diversity-aware ordering is unchanged.

D.11 MMR Sweep

The MMR baseline in the main paper is fixed at $\lambda = 0.5$. To confirm that the gap between Vendi-RAG and MMR is not an artifact of suboptimal λ , we sweep $\lambda \in \{0.3, 0.5, 0.7\}$ on HotpotQA (GPT-3.5 Turbo). Table 15 reports the results.

The best MMR configuration ($\lambda = 0.5$) reaches 56.0% accuracy, still well below Vendi-RAG’s 58.4%. The gap persists across all tested λ values, confirming that the advantage stems from the

Method	all-mpnet-base-v2	BGE-large
Adaptive Retrieval	43.2	44.8
Adaptive-RAG	47.0	48.7
Adaptive Retrieval + MMR	56.0	57.5
Vendi-RAG (dynamic)	58.4	59.2

Table 14: HotpotQA accuracy under two first-stage dense retrievers (GPT-3.5 Turbo). The Vendi-RAG advantage holds across encoders.

Method	Acc
MMR ($\lambda = 0.3$)	54.8
MMR ($\lambda = 0.5$)	56.0
MMR ($\lambda = 0.7$)	55.2
Vendi-RAG (dynamic, $s_1=0.8$)	58.4

Table 15: MMR λ sweep on HotpotQA (GPT-3.5 Turbo). The best MMR setting ($\lambda = 0.5$, 56.0%) still falls well below Vendi-RAG (58.4%); the gap is consistent across λ values and is not an artifact of suboptimal tuning.

VRS formulation rather than from poorly tuned MMR.

E Method Details

E.1 Normalization of the VRS Components

The Vendi Retrieval Score combines a diversity term and a relevance term, each normalized to $[0, 1]$ so that neither dominates due to differing dynamic ranges:

$$\widetilde{\text{VS}}_k(\mathcal{D}) = \frac{\text{VS}_k(\mathcal{D}) - 1}{n - 1}, \quad (8)$$

$$\widetilde{\text{SS}}(q, \mathcal{D}) = \frac{\text{SS}(q, \mathcal{D}) - \text{SS}_{\min}}{\text{SS}_{\max} - \text{SS}_{\min}}. \quad (9)$$

The VS normalization maps the score from its theoretical range $[1, n]$ to $[0, 1]$. The SS normalization uses running estimates of the minimum and maximum similarity scores computed over the candidate pool $\mathcal{C}$. The raw similarity $\text{SS}(q, \mathcal{D})$ is the average cosine similarity between the query q and the documents in $\mathcal{D}$, using dense vector representations from transformer-based encoders (e.g., BERT, Sentence-BERT).

E.2 Detailed Pipeline Stages

The Vendi-RAG pipeline expands the four-step loop summarized in Section 3 into the following stages.

1. Initial retrieval. The process begins by retrieving an initial document set using Vendi retrieval with a high diversity weight ($s_1 = 0.8$ in

all experiments). This stage carries the method’s core mechanism—VRS-based selection over the candidate pool—and stages 2–7 below operate on whatever set it returns. Prioritizing broad semantic coverage ensures the retrieved documents capture multiple perspectives and avoids semantically redundant documents. This initial diversity is particularly important for multi-hop QA, where synthesizing information from varied sources is essential for constructing complete reasoning chains.

2. Reasoning generation. Based on the retrieved documents, the system generates chain-of-thought (CoT) reasoning steps following established prompting techniques (Wei et al., 2022). These intermediate steps (i) identify relevant facts from each document, (ii) establish logical connections between disparate pieces of information, and (iii) build a coherent inferential pathway toward the final answer, yielding transparent and interpretable reasoning for complex multi-hop queries.

3. Candidate answer generation. Using both the reasoning steps and the retrieved documents, the LLM produces candidate answers. This dual-input design grounds answers in the retrieved evidence while following the logical structure established by the reasoning steps. Multiple candidates may be generated to explore different interpretations of the available information.

4. Quality evaluation. An LLM judge scores each candidate (see Section 3 and the prompt below), producing $Q_t \in [0, 1]$.

5. Dynamic adjustment of the VRS with smoothing. The diversity parameter is updated via the EMA rule in Eq. (3). The target follows an inverse relationship with quality: low Q_t raises the target toward diversity (broader exploration to recover missing information), while high Q_t lowers it toward relevance (consolidation along the current reasoning path). The EMA with $\beta = 0.3$ dampens rapid changes, preventing oscillation between extreme diversity and relevance settings.

6. Iterative refinement. Retrieval and reasoning repeat with the adjusted s at each stage, and the query may be reformulated to better target information gaps. The loop continues until the quality score exceeds the threshold τ or the maximum iteration count N is reached.

7. Final answer selection. The answer with the highest quality score across all iterations is returned, together with its retrieved documents and reasoning chain for interpretability. Because the best answer over all iterations is selected, a poor intermediate score cannot corrupt the final output.

What s controls, and what adjusting it adds. The weight s is the substantive knob, and most of its value is already realized at a fixed setting: high diversity (s near 1) covers different facets of complex queries, crucial in multi-hop QA where synthesizing multiple perspectives is necessary, whereas excessive relevance (s near 0) risks redundancy and limits comprehensive reasoning, and excessive diversity can admit loosely related content. A single well-chosen constant ($s = 0.8$, Table 8) therefore recovers the bulk of the benefit, which is why we recommend the fixed- s configuration as the default. Adapting s online is a refinement on top of this: when answer quality is satisfactory but incomplete, reducing s retrieves more specific documents to fill gaps; when quality is poor, increasing s promotes exploration of alternative document sets. This lets the system respond to each query’s difficulty and yields a further small gain concentrated on hard queries (Table 9), but it is a second-order effect relative to the scoring function itself.

E.3 The Relative-Quality Target and Its Safeguards

The target in Eq. (3) is *relative*: it compares Q_t to the running best rather than to an absolute scale. This is deliberate—it asks the judge only whether the current answer is better or worse than what the loop has already produced, a much weaker requirement than absolute calibration, and it is what makes the loop robust to a different or smaller judge (Appendix D.7). The apparent failure mode is a low-quality answer that is nonetheless the best so far, which sets $s^{\text{target}} = 0$ and pushes the retriever toward exploitation exactly when more exploration is needed. Four properties bound the damage. (i) *The shift is gradual*: EMA smoothing with $\beta = 0.3$ moves s only fractionally toward the target ($s_{t+1} = \beta s_t$ in this case), so a single low-but-

best answer cannot collapse the diversity weight, and any subsequent quality drop pushes the target back up. (ii) *Stopping is absolute, not relative*: the loop exits only when $Q_t \geq \tau$, so a locally “best” low-quality answer nudges the next s but can never trigger premature termination. (iii) *Best-answer tracking*: Algorithm 1 returns the highest-quality answer over all iterations, so a mistimed s shift cannot corrupt the output. (iv) *Exploration continues regardless*: REFINESQUERY (line 16) runs every iteration independently of s , so new evidence is still sought while s is temporarily relevance-leaning. Together these give trajectories that converge in 3–4 iterations on average (Appendix E.4). A variant that conditions the target on absolute answer quality is a natural alternative and we leave it to future work.

Interaction with query refinement. The diversity weight and the query are updated on separate tracks: s carries forward continuously across iterations through the EMA, independently of query rewriting. If iteration 1 ends with $s_2 = 0.56$, iteration 2 retrieves for the *rewritten* query but still with $s_2 = 0.56$, not with a reset $s_1 = 0.8$. The diversity–relevance balance is thus a property of the whole trajectory for a query, not of any individual sub-query, so a hop that is already well covered keeps the loop in exploitation mode even as the sub-query changes.

E.4 Convergence and Stability of the Diversity Parameter Update

The EMA update is a contraction toward the target with rate β , so the influence of any single quality score decays exponentially as β^m over subsequent iterations m . Combined with the clipping operation, this guarantees $s_t \in [0, 1]$ at every step and prevents divergence or runaway oscillation. Empirically, s_t trajectories converge smoothly within 3–4 iterations across all three datasets and all evaluated backbones. The judge needs only to provide a *relative* quality signal (whether the current answer is better or worse than previous iterations) rather than absolute calibration, which reduces the demands on the judge and accounts for the robustness observed when the judge is a different or smaller model than the generator (Appendix D.7).

E.5 Greedy vs. Exhaustive Selection

Although the VRS objective is non-submodular, greedy selection over a small candidate pool yields

k	$ \mathcal{C} $	VRS gap	Acc. difference
3	20	< 2%	not distinguishable
4	20	< 2%	not distinguishable
5	20	< 2%	not distinguishable
8	20	$\sim 2.4\%$	not distinguishable

Table 16: Greedy vs. exhaustive VRS selection on 100 sampled HotpotQA queries. “VRS gap” is the mean relative shortfall of the greedy subset’s VRS against the exhaustive optimum; “Acc. difference” reports whether downstream accuracy differs under a paired bootstrap at $p < 0.05$.

near-optimal subsets in practice. We quantify the gap on 100 randomly sampled HotpotQA queries by comparing the greedy subset against an exhaustive search over all $\binom{|\mathcal{C}|}{k}$ subsets, with $|\mathcal{C}| = 20$. Table 16 reports the mean relative VRS gap. It stays below 2% for $k \in \{3, 4, 5\}$ and rises only to 2.4% at $k = 8$, and in every setting the downstream accuracy difference between greedy and exhaustive selection is statistically indistinguishable under a paired bootstrap. Exhaustive search becomes impractical beyond this range— $\binom{20}{8}$ is already 125,970 subsets per query—but this is not a limitation in practice, because the retrieved-subset size is small by design and performance plateaus by $k = 8$ –10 under context-window constraints (Appendix D.9), so the settings covered here are the actual operating regime of the method. We also note that greedy optimization on a non-submodular objective is standard practice for MMR (Carbonell and Goldstein, 1998), which is likewise non-submodular, so the VRS introduces no new class of optimization risk relative to the strongest baseline it is compared against. The absence of an approximation guarantee nonetheless remains a genuine theoretical limitation of the method.

E.6 LLM-as-a-Judge Prompt

The exact judge prompt used in all experiments is reproduced below. The three criteria (coherence, relevance, query alignment) are each scored on a 1–10 scale, and Q_t is their mean divided by 10 so that $Q_t \in [0, 1]$.

LLM as a Judge Prompt

You are an expert LLM-based judge tasked with evaluating the quality of answers in a Retrieval-Augmented Generation (RAG) system. Your evaluation will consider the following aspects:

1. **Coherence:** Assess whether the provided answer is logically consistent and flows smoothly, without conflicting statements or gaps in reasoning.
2. **Relevance:** Evaluate how well the answer addresses the query based on the information from the retrieved documents.
3. **Query Alignment:** Determine how closely the answer aligns with the specific query asked, ensuring that the response is focused and appropriate.

Your evaluation will be quantified based on the following scoring system:

- Coherence Score (C): [1 - 10], where 10 is perfectly coherent.
- Relevance Score (R): [1 - 10], where 10 is highly relevant to the query.
- Query Alignment Score (Q): [1 - 10], where 10 is perfectly aligned.

Provide a quality score Q_t as the average of these individual scores:

$$Q_t = \text{mean}(C, R, Q)$$

Query: {query}




Figure 3: Performance comparison of Vendi-RAG and Adaptive-RAG across the datasets and three GPT backbones (GPT-3.5, GPT-4o-mini, GPT-4o), using EM, F1, and Accuracy. Bars are grouped by (dataset, metric); within each group the leftmost three bars are Adaptive-RAG across the three backbones and the rightmost three bars are Vendi-RAG (dynamic, $s_1 = 0.8$) across the same three backbones. Vendi-RAG with GPT-4o outperforms all other configurations on every dataset and metric, with the largest absolute gain on HotpotQA. Numeric per-cell values are in Appendix D.6.