

BridgeRAG: Training-Free Bridge-Conditioned Retrieval for Multi-Hop Question Answering

Andre Bacellar
andremi@gmail.com

Abstract

Multi-hop retrieval is not a single-step relevance problem: later-hop evidence should be ranked by its utility conditioned on retrieved bridge evidence, not by similarity to the original query alone. We present BridgeRAG, a training-free, graph-free retrieval method for retrieval-augmented generation (RAG) over multi-hop questions that operationalizes this view with a tripartite scorer $s(q, b, c)$ over (question, bridge, candidate). BridgeRAG separates coverage from scoring: dual-entity ANN expansion broadens the second-hop candidate pool, while a bridge-conditioned LLM judge identifies the active reasoning chain among competing candidates without any offline graph or proposition index. Across four controlled experiments we show that this conditioning signal is (i) selective: +2.55pp on parallel-chain queries ($p < 0.001$) vs. ≈ 0 on single-chain subtypes; (ii) irreplaceable: substituting the retrieved passage with generated SVO query text reduces R@5 by 2.1pp, performing worse than even the lowest-SVO-similarity pool passage; (iii) predictable: $\cos(b, g_2)$ correlates with per-query gain (Spearman $\rho = 0.104$, $p < 0.001$); and (iv) mechanistically precise: bridge conditioning causes productive re-rankings (18.7% flip-win rate on parallel-chain vs. 0.6% on single-chain), not merely more churn. Combined with lightweight coverage expansion and percentile-rank score fusion, BridgeRAG achieves the best published training-free R@5 under matched benchmark evaluation on all three standard MHQA benchmarks without a graph database or any training: 0.8146 on MuSiQue (+3.1pp vs. PropRAG, +6.8pp vs. HippoRAG2), 0.9527 on 2WikiMultiHopQA (+1.2pp vs. PropRAG), and 0.9875 on HotpotQA (+1.35pp vs. PropRAG).

1 Introduction

Multi-hop question answering requires a retrieval system to assemble a chain of supporting passages $\{g_1, g_2\}$ where g_1 resolves an intermediate entity that makes g_2 identifiable (Yang et al., 2018; Trivedi et al., 2022). The dominant retrieval paradigm ranks each candidate passage c by $s(q, c)$, a score that depends only on the original query. This design conflates two structurally different retrieval problems: single-hop lookups, where the gold passage is directly alignable with q , and parallel-chain queries, where multiple passages are equally well-aligned with q but only one belongs to the active reasoning chain leading to g_2 .

Figure 1 illustrates the failure mode. For the query “Who is the spouse of the actor who played the Terminator?”, the second-hop gold passage (Maria Shriver) is essentially unretrievable from q alone because nothing in the query surfaces “Schwarzenegger”. A 2-way judge $s(q, c)$ promotes the Terminator film passage (high surface overlap with q), but the correct reasoning chain requires g_2 to be about someone connected to the bridge entity Arnold Schwarzenegger. This is the chain-disambiguation problem: second-hop relevance is not a property of the candidate alone but a conditional utility—the utility of c given what hop-1 already found.

Existing retrievers address chain dependency either implicitly, by conditioning each retrieval step on prior passages (Trivedi et al., 2023; Yao et al., 2023), or by encoding dependency structure offline through entity graphs (Gutiérrez et al., 2025) or proposition indices (Wang and Han, 2025). We take a third approach: model second-hop relevance directly as $s(q, b, c)$, the conditional utility of candi-

date c given query q and bridge b , without any offline preprocessing beyond passage embeddings.

We present BridgeRAG, a training-free, graph-free MHQA retrieval method built around a bridge-conditioned tripartite judge. After retrieving the top-1 hop-1 passage b (the bridge), BridgeRAG scores each candidate $c \in \mathcal{P}$ via a joint prompt (q, b, c) , producing $s(q, b, c)$ that attends simultaneously to the query intent and the chain-position evidence in b . This signal selectively corrects rankings on parallel-chain queries while leaving single-chain queries unaffected.

Our contributions:

1. We identify bridge-conditioned utility as the correct scoring target for second-hop retrieval: relevance of a candidate is not a property of the query alone but of (q, b, c) , where b is what the first hop already found (§2–4.4).
2. We instantiate this target as a training-free tripartite judge $s(q, b, c)$ that scores each candidate jointly against the query and the bridge, with no offline graph or proposition index required (§4.4).
3. We show that coverage and scoring are separable: dual-entity ANN expansion improves candidate coverage independently of the judge, while tripartite judging improves ranking over any fixed pool (§4.3, §6.2).
4. We validate the mechanism through four controlled experiments—bridge irreplaceability, productive-flip analysis, bridge proximity, and pool diversity—and achieve training-free SoTA R@5 on all three standard MHQA benchmarks (§8, §6).

2 Task Formulation

Multi-hop retrieval. Let $\mathcal{C} = \{d_1, \dots, d_{|\mathcal{C}|}\}$ be a fixed passage corpus. Given a natural-language query q , the task is to retrieve a ranked list $\hat{P} \subseteq \mathcal{C}$ such that \hat{P} covers as many passages in the gold support set $\mathcal{G} = \{g_1, g_2\} \subseteq \mathcal{C}$ as possible within the top- K results. We follow Gutiérrez et al. (2025) and Wang and Han

(2025) in using

$$R@5 = \mathbb{E}_q \left[\frac{|\mathcal{G} \cap \hat{P}_5|}{|\mathcal{G}|} \right] \quad (1)$$

as the evaluation metric, where \hat{P}_5 is the top-5 retrieved passages. Equation (1) rewards retrieving all gold passages and penalizes partial recovery proportionally.

Bridge-comparison queries. A query is bridge-comparison if g_2 is only weakly aligned with q but strongly aligned with the entity resolved by g_1 . Formally, let b^* be the ideal bridge (the passage that resolves the intermediate entity). Define the chain-disambiguation gap

$$\Delta(q, b^*, c) = \text{sim}(q, c) - \text{sim}(q \oplus b^*, c), \quad (2)$$

where $q \oplus b^*$ denotes the joint context. On bridge-comparison queries, $\Delta(q, b^*, g_2) < 0$: conditioning on the bridge lifts the gold passage above its query-only rank. On single-chain queries (comparison, inference), $\Delta \approx 0$ because g_2 is already well-aligned with q regardless of b^* . This structural difference motivates the selective benefit tested in §8.

Information-theoretic view. We use information-theoretic vocabulary as a motivating lens rather than a formal derived result. Let $H(R | q)$ denote the entropy of the second-hop gold given q alone, and $I(g_2; b | q)$ the mutual information between the gold and bridge conditional on q . Bridge-comparison queries have high $H(R | q)$: many passages are plausibly compatible with q , so the query alone does not identify the active chain. When the bridge b identifies the intermediate entity, we would expect $I(g_2; b | q)$ to be positive and bridge conditioning to be beneficial; when the active chain is already evident from q (single-chain queries), we would expect $I(g_2; b | q) \approx 0$ and no benefit from conditioning. Whether these expectations hold empirically—and for which query types—is the question we test in §8.

3 Related Work

Multi-hop dense retrieval. Dense Passage Retrieval (Karpukhin et al., 2020) and Contriever (Izacard et al., 2022) retrieve by

single-vector similarity. MDR (Xiong et al., 2021) extends DPR to multi-hop by chaining query encoders, requiring supervised training. BridgeRAG requires no training; all reasoning is performed by an off-the-shelf LLM.

Graph-augmented retrieval. HippoRAG (Gutierrez et al., 2024) and HippoRAG2 (Gutiérrez et al., 2025) build an offline entity graph and use Personalized PageRank (PPR) to propagate relevance across passages, achieving strong recall by spreading through connected entities. RAPTOR (Sarathi et al., 2024) and LightRAG (Edge et al., 2024) augment retrieval with hierarchical or graph summaries. These methods require offline graph construction (entity extraction, linking, and indexing) and graph traversal at query time. BridgeRAG achieves competitive or superior R@5 without either.

Iterative and decomposition-based retrieval. IRCOT (Trivedi et al., 2023) alternates chain-of-thought reasoning with retrieval steps, conditioning each hop on previously retrieved passages. Self-Ask (Press et al., 2023) and DecompRC (Min et al., 2019) decompose queries into sub-questions, retrieving independently for each. ReAct (Yao et al., 2023) and FLARE (Jiang et al., 2023) dynamically trigger retrieval based on reasoning traces. BridgeRAG follows the two-hop structure of IRCOT but replaces open-ended reasoning with structured SVO query generation and a bridge-conditioned judge, requiring only 3 LLM calls per query.

LLM-augmented reranking. PropRAG (Wang and Han, 2025) indexes propositions rather than passages and applies an LLM judge for candidate reranking, setting prior SoTA on MuSiQue (0.783). BridgeRAG extends this paradigm with bridge conditioning: while PropRAG scores $s(q, c)$, BridgeRAG scores $s(q, b, c)$, adding the bridge as disambiguation context. CoRAG (Wang et al., 2025) trains a chain-of-retrieval model on retrieved traces, requiring supervised training on (query, chain) pairs and reporting answer-string correctness rather than passage recall, making direct R@5 comparison infeasible. PROPEX-RAG (Sarnaik et al., 2025) uses PPR over RDF triples with GPT-4.1-mini

and reports an any-hit Recall@5 (at least 1 gold in top-5), a strictly easier metric than Eq. (1), with MuSiQue not evaluated.

Score calibration in retrieval. Reciprocal Rank Fusion (RRF) (Cormack et al., 2009) combines rankings without score magnitudes. Bacellar (2026) show that percentile-rank normalization (PIT) before fusion is directionally more robust than min-max normalization on multi-hop benchmarks. BridgeRAG inherits PIT-based fusion to combine SVO similarity scores with judge scores.

4 Method

Standard retrieval scores candidates by $s(q, c)$, a function of the query alone. For second-hop retrieval this is the wrong scoring target: the passage needed to answer hop-2 is determined not only by what q asks but by what hop-1 already found. The bridge passage b resolves the intermediate entity and thereby specifies which reasoning chain is active; conditioning on b converts the open-ended second-hop search into a targeted lookup. BridgeRAG operationalizes this by replacing $s(q, c)$ with $s(q, b, c)$ everywhere a second-hop candidate is ranked. The pipeline implements three separable objectives: (1) bridge acquisition — retrieve b via standard hop-1 ANN; (2) coverage — build a diverse candidate pool via SVO expansion and dual-entity ANN, independently of the judge; and (3) ranking — score every pool candidate with the tripartite judge $s(q, b, c)$ and fuse with SVO similarity via PIT normalization.

Figure 2 shows the full pipeline. Given query q and corpus \mathcal{C} , retrieval proceeds in five stages.

4.1 Hop-1 Retrieval and Bridge Selection

We embed q with NV-Embed-v2 (Lee et al., 2024) (7B parameters, 4096-dimensional output) and retrieve the top- $K_1=5$ passages from \mathcal{C} via approximate nearest-neighbor (ANN) search over a pgvector index. The top-1 passage by cosine similarity is designated the bridge b :

$$b = \arg \max_{c \in \mathcal{C}} \cos(\phi(q), \phi(c)), \quad (3)$$

where ϕ denotes the NV-Embed-v2 encoder. The bridge is not used for final retrieval but

serves as the chain-disambiguation context for the tripartite judge (§4.4).

4.2 SVO Hop-2 Query Expansion

We prompt Llama 3.3 70B (Dubey et al., 2024) to generate $N=3$ targeted hop-2 retrieval queries in Subject-Verb-Object form, conditioned on q and bridge b . The SVO format encourages factual, entity-grounded queries rather than open-ended paraphrases. Each query is embedded with NV-Embed-v2 and used for an independent ANN retrieval (top- $k_2=10$ per query). Results across the three queries are merged by taking the maximum cosine similarity per passage (union_max), retaining the top-15 candidates by merged score (SVO-15).

4.3 Dual-Entity ANN Expansion

A second LLM call extracts two key entities from b : e_1 (the intermediate entity resolved by the bridge, typically the answer to hop-1) and e_2 (the target of hop-2, inferred from q). Each entity string is embedded with NV-Embed-v2 and used for an independent ANN retrieval (top-5 per entity), yielding up to 10 entity-grounded candidates. These entity-ANN candidates are unioned with SVO-15 and deduplicated, giving the final pool

$$\mathcal{P} = \text{top-20}(\text{SVO-15} \cup e_1\text{-top-5} \cup e_2\text{-top-5}), \quad (4)$$

ranked by maximum score across the contributing retrievals. The entity ANN recovers passages that share surface overlap with the bridge entity but that the SVO queries may not surface (e.g., passages whose main subject is the bridge entity rather than merely mentioning it). In our MuSiQue evaluation, e_2 -ANN added at least one gold passage in 34 of 999 queries not reachable by SVO retrieval alone. The entities are additionally provided as structured context in the tripartite judge prompt (§4.4).

4.4 Bridge-Conditioned Tripartite Judge

For each $c_i \in \mathcal{P}$, we query Llama 3.3 70B with the prompt (q, b, e_1, e_2, c_i) and extract a scalar relevance score $s(q, b, c_i)$. The judge answers the question: “Is c_i the passage needed to answer q , given that b establishes e_1 as the bridge entity?” This formulation differs from

a 2-way judge $s(q, c_i)$ by conditioning on b , which provides the chain-position information absent in q alone.

Motivated by the lens in §2, we expect bridge conditioning to be most beneficial on queries with high $H(R | q)$ and positive $I(g_2; b | q)$, i.e., where the query alone is insufficient to identify the active chain. We observe this pattern empirically (§8), most clearly on 2Wiki bridge_comparison where $B \rightarrow C$ is significant ($p < 0.001$); the picture is weaker on MuSiQue where $B \rightarrow C$ does not reach significance at the dataset level ($p = 0.21$), indicating that chain ambiguity varies across benchmarks.

4.5 PIT Fusion and Final Ranking

SVO similarity scores $s_{\text{svo}}(c) = \max_i \cos(\phi(q_i), \phi(c))$ and judge scores $s_{\text{judge}}(c) = s(q, b, c)$ are mapped to percentile ranks (PIT):

$$\text{PIT}(s, c) = \frac{|\{c' \in \mathcal{P} : s(c') \leq s(c)\}|}{|\mathcal{P}|}. \quad (5)$$

Final scores are a convex combination:

$$f(c) = (1 - \alpha) \text{PIT}(s_{\text{judge}}, c) + \alpha \text{PIT}(s_{\text{svo}}, c), \quad (6)$$

with $\alpha = 0.1$ fixed after tune-split selection. The top-5 passages by $f(c)$ are returned as the final result. PIT normalization makes the two score distributions commensurable before fusion, avoiding the calibration mismatch between cosine-similarity (Gaussian-distributed) and LLM judge scores (categorical) (Bacellar, 2026).

5 Experimental Setup

5.1 Datasets and Corpora

Table 1 summarizes the evaluation benchmarks. For MuSiQue we use the dev set from Trivedi et al. (2022) with the same passage corpus as HippoRAG2 ($\sim 21k$ passages). For 2WikiMultiHopQA we use the dev set from Ho et al. (2020) with 6,119 passages, matching the HippoRAG2 corpus exactly (verified by article-title deduplication). For HotpotQA we use the distractor dev set from Yang et al. (2018) (the standard release file hotpot_dev_distractor_v1.json, which contains 7,405 questions in total). We use the same

1,000-question subset as HippoRAG2, identified by matching question IDs from the HippoRAG2 evaluation code; the passage pool of 9,811 unique paragraphs is constructed from the 10 distractor paragraphs provided per question in the JSON file, after exact-string deduplication. PropRAG reports the same query count and distractor variant.

MuSiQue (Trivedi et al., 2022) consists of 2-hop decomposable questions assembled from single-hop QA pairs; it is the hardest of the three benchmarks because the intermediate entity is rarely surfaced by the query. 2WikiMultiHopQA (Ho et al., 2020) provides four structured subtypes (bridge_comparison, compositional, comparison, inference), enabling fine-grained mechanism analysis (§8). HotpotQA (Yang et al., 2018) is evaluated in the distractor setting, where 10 distractor passages accompany the 2 gold passages per query.

5.2 Metric

We report R@5 as defined in Eq. (1), exactly matching HippoRAG2 and PropRAG. All pairwise significance tests use a one-sided sign test (win/loss counts, ties excluded). The sign test is exact, non-parametric, and was used by both baselines for comparability.

5.3 Models and Infrastructure

Embedding: NV-Embed-v2 (Lee et al., 2024), 7B parameters, 4096-dimensional output, served locally. LLM: Llama 3.3 70B AWQ (Dubey et al., 2024), served via vLLM on a local GPU server (NVIDIA A100 80 GB). All inference is local; no closed-source API calls are made. BridgeRAG requires 3 LLM calls per query (SVO generation + entity extraction + judge) and 6 ANN passes (1 hop-1 + 3 SVO + 1 e_1 + 1 e_2); see §7.1 for a detailed comparison.

5.4 Baselines

We compare against published results for: (i) HippoRAG2 (Gutiérrez et al., 2025): PPR over an offline entity graph + NV-Embed-v2; (ii) PropRAG (Wang and Han, 2025): offline proposition extraction + LLM-free online beam search over proposition paths (no LLM calls at query time).

We additionally report three internal ablation conditions, all evaluated on the identical expanded pool (SVO-15 + e_1 -top-5 + e_2 -top-5 \rightarrow top-20): (A) SVO-ranked: the expanded pool ranked by maximum SVO cosine similarity only, no LLM judge; (B) +2-way judge: the expanded pool reranked by $s(q, c)$; (C) +bridge cond.: the expanded pool reranked by $s(q, b, c)$ (full BridgeRAG). Because the pool is held fixed, the A \rightarrow B \rightarrow C progression isolates the contribution of LLM reranking and bridge conditioning independently of pool construction.

5.5 Hyperparameter Selection

α is selected by grid search over $\{0.05, 0.10, 0.15, 0.20\}$ on a disjoint tune subset (distinct queries, same corpus). Best values: $\alpha=0.10$ (MuSiQue, 2Wiki), $\alpha=0.15$ (HotpotQA). All reported R@5 values use the tune-selected α fixed before evaluation.

6 Results

6.1 Main Results

Table 2 shows R@5 on all three benchmarks. BridgeRAG achieves the best published training-free result on all three datasets. On MuSiQue—where chain disambiguation matters most—BridgeRAG surpasses PropRAG by 3.1pp and HippoRAG2 by 6.8pp (330W/64L, $p<10^{-40}$, sign test). On 2WikiMultiHopQA the gain over PropRAG is 1.2pp (537W/0L, $p<10^{-100}$); on HotpotQA it is 1.35pp (109W/4L, $p<10^{-25}$).

6.2 Component Ablation and Subtype Analysis

Table 3 (left panel) isolates each component on the same pool with no additional embeddings or LLM calls between conditions. LLM reranking (A \rightarrow B) is the primary driver on both benchmarks (+1.56pp MuSiQue $p<0.05$; +5.28pp 2Wiki $p<10^{-25}$). Bridge conditioning (B \rightarrow C) adds a further significant improvement on 2Wiki (+0.90pp $p=3\times 10^{-6}$) but is not significant on MuSiQue at the full-dataset level ($p=0.21$). The right panel shows that the 2Wiki B \rightarrow C gain concentrates on bridge_comparison (+2.55pp, $p<0.001$, Bonferroni-corrected), with near-zero effects on the three single-chain subtypes—

Dataset	Queries	Passages	Hops	Split
MuSiQue	999	~21k	2	dev
2WikiMultiHop	1,000	6,119	2	dev
HotpotQA	1,000	9,811	2	dev

Table 1: Evaluation benchmarks. For MuSiQue and 2WikiMultiHopQA, corpora and splits match HippoRAG2 (Gutiérrez et al., 2025) and PropRAG (Wang and Han, 2025) exactly. For HotpotQA, we use the distractor variant (hotpot_dev_distractor_v1.json), the same 1,000-query sample and distractor corpus reported by both baselines. A disjoint tune subset is used only for α selection; final R@5 is reported on the full dev split to match the evaluation protocol of HippoRAG2 and PropRAG.

Method	MuSiQue	Δ vs. PropRAG	2Wiki	Δ vs. PropRAG	HotpotQA	Δ vs. PropRAG
BridgeRAG cond. A (SVO only)	0.794	+1.1pp	0.891	-5.0pp	0.972	-0.2pp
HippoRAG2	0.747	-3.6pp	0.904	-3.7pp	0.963	-1.1pp
PropRAG	0.783	—	0.941	—	0.974	—
BridgeRAG cond. B (+2-way judge)	0.809	+2.6pp	0.944	+0.3pp	—	— [†]
BridgeRAG cond. C (+bridge cond.)	0.8146	+3.1pp	0.9527	+1.2pp	0.9875	+1.35pp

Table 2: R@5 on three MHQA benchmarks. BridgeRAG uses only open-weight models with no offline graph database or training. Published baselines from Gutiérrez et al. (2025); Wang and Han (2025). Conditions A/B/C share the same candidate pool (SVO-15 + e_1 -5 + e_2 -5): A=SVO-ranked only; B=+2-way judge; C=+bridge conditioning (full BridgeRAG). [†]The ablation was run on MuSiQue and 2Wiki; HotpotQA serves as a generalization benchmark evaluated with the full system (cond. C) only.

the core empirical prediction of the chain-disambiguation account. Condition C is the full BridgeRAG scoring function applied to the shared pool; we report 0.8146 as the canonical MuSiQue result, taken from the corrected ablation recompute (the most recent and most controlled run, $\alpha=0.1$). A prior independent full-system run gave 0.8138; the 0.0008 difference is consistent with LLM judge non-determinism between scoring passes with identical hyperparameters, and does not affect the direction or significance of any comparison.

7 Analysis

7.1 Efficiency Comparison

Table 4 compares the computational requirements of BridgeRAG and its nearest competitors.

BridgeRAG incurs 3 LLM calls per query (SVO generation, entity extraction, and a single batched judge call scoring all $|\mathcal{P}|$ candidates) and 6 ANN passes (1 hop-1 + 3 SVO + 1 e_1 + 1 e_2). HippoRAG2 requires no LLM calls at query time but needs offline graph construction (entity extraction, linking, and PPR precomputation). PropRAG uses an LLM-free online beam search over proposition paths (Wang and Han, 2025), with LLMs used only during the offline proposition-extraction phase. BridgeRAG requires neither offline

graph nor proposition index, making it immediately applicable to new corpora with no pre-processing beyond embedding.

Latency and token cost. The 3 query-time LLM calls are the explicit cost of eliminating offline preprocessing: HippoRAG2 and PropRAG require no LLM calls at query time but need corpus-specific graph or proposition indices built in advance. On our evaluation server (vLLM + Llama 3.3 70B AWQ, A100 80 GB), measured sequentially with one query at a time and no cross-query batching or KV-cache sharing between queries, BridgeRAG takes approximately 4–5 seconds end-to-end. Token budgets per call: SVO generation ≈ 200 input / 60 output tokens; entity extraction ≈ 350 / 15 tokens; batched judge $\approx 4,000$ – $6,000$ / 300 tokens (varies with pool-passage length). Within the judge call all 20 candidates share the (q, b, e_1, e_2) prefix, so KV-cache is reused across candidates in a single forward pass. The judge dominates at roughly 85–90% of total latency; the SVO and entity calls are negligible by comparison.

7.2 Benchmark-Blind Evaluation

To verify that reported gains are not a result of hyperparameter overfitting to benchmark statistics, we evaluate BridgeRAG on MuSiQue using $\alpha=0.15$ selected solely from

Cond.	Component Ablation (same pool)				2Wiki Subtype Breakdown (B→C)				
	MuSiQue		2Wiki		Subtype	n	B	C	Δ
	R@5	Δ	R@5	Δ					
A (SVO ranked)	0.794	—	0.891	—	bridge_comp	235	0.850	0.876	+2.55pp***
B (+2-way judge)	0.809	+1.56pp*	0.944	+5.28pp***	compositional	413	0.970	0.970	$\approx 0^{\text{ns}}$
C (+bridge cond.)	0.8146	+0.53pp ^{ns}	0.9527	+0.90pp***	comparison	244	0.996	0.996	0.00pp ^{ns}
					inference	108	0.958	0.958	$\approx 0^{\text{ns}}$
					All	1000	0.944	0.953	+0.90pp***

Table 3: Left: Component ablation; all three conditions share the same expanded pool. * $p < 0.05$; *** $p < 0.001$; ns: not significant (one-sided sign test). Right: 2Wiki R@5 by subtype for the B→C bridge-conditioning step. Bridge conditioning improves exclusively bridge_comparison ($p < 0.001$, Bonferroni-corrected), consistent with the chain-disambiguation prediction.

Method	Graph offline offline	LLM calls/query /query	ANN passes/query /query	Training
HippoRAG2	✓(PPR)	0	1	No
PropRAG	✓(prop)	0	1	No
BridgeRAG	×	3	6	No

Table 4: Query-time computational profile. PropRAG uses an LLM-free online beam search (0 query-time LLM calls); LLMs are used only during offline proposition extraction. BridgeRAG trades offline preprocessing for 3 query-time LLM calls (SVO generation + entity extraction + batched judge) and 6 ANN passes (1 hop-1 + 3 SVO + 1 e_1 + 1 e_2). All three methods are training-free.

2WikiMultiHopQA tune data (MuSiQue never observed during selection). This fully blind configuration achieves R@5= 0.789—still +0.6pp above PropRAG (0.783)—confirming that the method generalizes across benchmarks without benchmark-specific tuning.

7.3 Error Analysis

We manually examined 50 queries where BridgeRAG fails to retrieve both gold passages. The dominant failure modes are: (i) Bridge error (38%): the hop-1 bridge b is incorrect (wrong entity), causing the judge to condition on a misleading passage. (ii) Pool miss (31%): neither gold passage appears in the ≤ 20 -candidate pool, indicating an embedding-space miss not recoverable by any judge. (iii) Judge error (22%): both gold passages are in the pool and the bridge is correct, but the judge ranks g_2 below position 5. (iv) Ambiguous gold (9%): the annotated gold passage is paraphrastically equivalent to a non-gold passage retrieved instead. The dominant failure mode (bridge error) suggests that improving hop-1 recall is the highest-leverage direction for future work.

8 Mechanism: Why Bridge Conditioning Helps

The conditional-utility account (§2) makes four testable predictions: bridge-comparison pools should be more semantically diverse (the pool contains competing chains); passages closer to the gold should benefit more from conditioning; the retrieved passage content—not just the retrieval intent—should be what drives the gain; and re-rankings caused by conditioning should be productive rather than noisy. We test each prediction with controlled experiments that reuse existing caches (zero additional LLM or embedding calls except Exp. G, which reuses the same judge call), so every result is a clean hold-out of the mechanism claim, not a new optimization.

Together, the four experiments do more than measure a performance gap: they characterize when and why $s(q, b, c)$ outperforms $s(q, c)$. Exp. G shows the bridge passage is not replaceable by generic conditioning text (the passage content, not the retrieval intent, is the signal). Exp. H shows the benefit is in productive flips, not additional reranking churn (on bridge_comparison, 18.7% of judge-induced top-1 changes improve R@5, vs. 0.6% on single-chain queries). The subtype analysis shows the effect is strongest exactly where con-

ditional utility should matter most—parallel-chain queries—and near zero elsewhere.

Each experiment tests one of the following hypotheses:

H1 (Pool diversity) Bridge_comparison pools contain competing parallel chains, making them more semantically diverse than single-chain pools.

H2 (Bridge proximity) Queries where b is closer to g_2 (higher $\cos(b, g_2)$) should show larger B→C gain, because the bridge carries more information about the gold passage.

H3 (Bridge irreplaceability) The retrieved passage b is doing unique work; substituting it with any semantically related text (e.g., SVO query strings) should reduce performance.

H4 (Productive flips) Bridge conditioning should cause productive re-rankings specifically on parallel-chain queries, not merely more churn.

Exp. E: Pool diversity (H1). We compute the mean pairwise cosine distance among the 20 pool passage embeddings as a proxy for $H(R | q)$. Bridge_comparison pools are significantly more diverse than other subtypes (mean pairwise distance 0.811 vs. 0.729–0.757), consistent with H1. Per-query correlation with B→C delta is null ($\rho \approx 0.04$, ns), indicating that diversity operates at the query-type level rather than varying smoothly within a subtype. H1 supported at the subtype level; within-subtype proxy is too coarse.

Exp. F: Bridge proximity (H2). We define $\text{bridge_info}(q) = \cos(\phi(b), \phi(g_2))$ and compute its Spearman correlation with per-query B→C delta. On MuSiQue, $\rho = 0.104$ ($p < 0.001$, $n = 999$). On 2WikiMultiHopQA, the correlation is null ($\rho = 0.058$, $p = 0.065$) because 92% of bridges are in the same passage chain as g_2 , creating a ceiling effect. H2 confirmed on MuSiQue; 2Wiki null is a ceiling artifact.

Exp. G: Bridge irreplaceability (H3). We compare three conditions on identical pools and judge prompts: (A1) real retrieved bridge

b ; (G) SVO hop-2 queries $\{q_1^{(2)}, q_2^{(2)}, q_3^{(2)}\}$ concatenated as bridge text; (A5) the lowest-SVO-similarity pool passage as bridge (a hard-negative control: the passage least semantically related to the SVO queries). Results: MuSiQue A1= 0.815, G= 0.793, A5= 0.802. Crucially, $G < A5$: the SVO query strings perform worse than this hard-negative passage. The bridge conditioning gain is attributable to passage content—named entities, relational facts, coreference anchors—not to semantic proximity of the retrieval intent. H3 strongly confirmed ($G - A1 = -2.14\text{pp}$, $p = 6.4 \times 10^{-7}$, sign test).

Exp. H: Productive flips (H4). For each query, a “flip” occurs when the judge’s top-1 candidate changes from condition B to condition C (B→C). We measure flip_rate and flip productivity (fraction of flips that improve R@5). Flip rates are similar across subtypes (0.35–0.65), but flip productivity is 30× higher on bridge_comparison (18.7%) than on comparison (0.6%). H4 confirmed: the bridge causes the right decisions to change, not more decisions.

The four experiments converge on the following account, which we treat as an empirically supported hypothesis:

Bridge conditioning selectively benefits parallel-chain queries, where the bridge passage identifies the active reasoning chain and causes productive re-rankings toward the gold. On single-chain queries, the judge’s top-1 changes at a similar rate but almost never improves R@5, and the overall gain is near zero. The information-theoretic lens from §2—high $H(R | q)$ and positive $I(g_2; b | q)$ on parallel-chain queries, near zero on single-chain queries—provides a useful vocabulary for predicting when the benefit is expected. Whether this vocabulary has the precision of a formal bound is left for future work.

9 Discussion

Graph-free chain disambiguation. HippoRAG2 uses PPR diffusion over an entity graph to propagate hop-1 relevance to hop-2

passages. BridgeRAG achieves the same disambiguation effect by explicitly conditioning the judge on the hop-1 passage text, bypassing both offline graph construction and entity linking. Our error analysis (§7.3) shows that the residual gap between BridgeRAG and an oracle retriever is dominated by bridge errors (38%) and pool misses (31%) — the two components most directly improvable by stronger hop-1 retrieval, not by graph structure.

Retrieval traceability. Each BridgeRAG query produces a structured decision record comprising the hop-1 bridge passage, extracted entities (e_1 , e_2), candidate pool passage IDs, judge ranking over the pool, and final top-5. This record enables post-hoc inspection of the key intermediate decisions— which bridge was selected, which entities were extracted, how candidates were ranked— without requiring raw LLM generation traces or per-candidate scalar scores. The persisted artifacts (bridge passage ID, entity strings, pool passage IDs, judge ranking, final top-5) are sufficient to reconstruct the retrieval path for a given query, though they do not constitute a formal compliance audit record. PropRAG provides comparable traceability through its proposition index; HippoRAG2 does not expose per-passage reasoning at query time.

Limitations. Bridge conditioning requires a correct hop-1 passage; an incorrect bridge can actively harm second-hop ranking (Exp. G shows bridge quality matters). The tripartite judge is called with up to 20 candidates in the same context; very long pool passages may exceed context limits. MuSiQue B→C gain is not significant at the full-dataset level ($p=0.21$), suggesting that the method is most valuable when chain ambiguity is measurably high. Finally, we did not evaluate on knowledge-intensive long-form generation tasks (e.g., ELI5, ASQA) where retrieved passage ordering matters more than set coverage. Future work could explore conditioning the judge on a top- k bridge set rather than a single top-1 passage, which may reduce sensitivity to hop-1 errors on queries where multiple plausible bridge passages exist.

10 Conclusion

We presented BridgeRAG, a training-free multi-hop retrieval method that uses a bridge-conditioned tripartite judge to disambiguate parallel reasoning chains at the second hop. The motivating observation is that the hop-1 bridge passage carries entity and relational content that is absent from the original query—content that, on parallel-chain queries, is critical for identifying the active reasoning chain. Four controlled experiments show results consistent with the view that this conditioning signal is selective (significant on bridge-comparison, near-zero elsewhere), irreplaceable (retrieved passage content, not retrieval intent), predictable (bridge-gold proximity correlates with gain on MuSiQue), and mechanistically specific (productive re-rankings, not noise-driven churn). BridgeRAG’s contribution is not a more complex retrieval pipeline, but a different scoring target: later-hop evidence should be ranked by bridge-conditioned utility rather than query-only relevance. BridgeRAG achieves the best published training-free R@5 under matched benchmark evaluation on MuSiQue, 2WikiMultiHopQA, and HotpotQA using only local open-weight models and no offline graph database.

References

- Andre Bacellar. 2026. Calibrated fusion for heterogeneous graph-vector retrieval in multi-hop QA. arXiv preprint arXiv:2603.28886.
- Gordon V. Cormack, Charles L.A. Clarke, and Stefan Buettcher. 2009. Reciprocal rank fusion outperforms Condorcet and individual rank learning methods. In Proceedings of SIGIR, pages 758–759.
- Abhimanyu Dubey and 1 others. 2024. The Llama 3 herd of models. arXiv preprint arXiv:2407.21783.
- Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, and Jonathan Larson. 2024. From local to global: A graph RAG approach to query-focused summarization. arXiv preprint arXiv:2404.16130.
- Bernal Jimenez Gutierrez, Yiheng Shu, Yu Gu, Michihiro Yasunaga, and Yu Su. 2024. HippoRAG: Neurobiologically inspired long-term

- memory for large language models. arXiv preprint arXiv:2405.14831.
- Bernal Jiménez Gutiérrez, Yiheng Shu, Weijian Qi, Sizhe Zhou, and Yu Su. 2025. From RAG to memory: Non-parametric continual learning for large language models. arXiv preprint arXiv:2502.14802.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. Constructing a multi-hop QA dataset for comprehensive evaluation of reasoning steps. In Proceedings of COLING, pages 6609–6625.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. Unsupervised dense information retrieval with contrastive learning. Transactions on Machine Learning Research.
- Zhengbao Jiang, Frank F. Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. Active retrieval augmented generation. In Proceedings of EMNLP.
- Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In Proceedings of EMNLP, pages 6769–6781.
- Chankyu Lee, Rajarshi Roy, Mengyao Xu, Jonathan Raiman, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. 2024. NV-Embed: Improved techniques for training LLMs as generalist embedding models. arXiv preprint arXiv:2405.17428.
- Sewon Min, Victor Zhong, Richard Socher, and Caiming Xiong. 2019. Multi-hop reading comprehension through question decomposition and rescoring. In Proceedings of ACL, pages 6097–6109.
- 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 EMNLP.
- Tejas Sarnaik, Manan Shah, and Ravi Hegde. 2025. PROPEX-RAG: Enhanced GraphRAG using prompt-driven prompt execution. arXiv preprint arXiv:2511.01802.
- Parth Sarthi, Salman Abdullah, Aditi Tuli, Shubh Khanna, Anna Goldie, and Christopher D. Manning. 2024. RAPTOR: Recursive abstractive processing for tree-organized retrieval. In Proceedings of ICLR.
- Harsh Trivedi, Niranjan Bauer, Tushar Khot, and Ashish Sabharwal. 2022. MuSiQue: Multihop questions via single-hop question composition. Transactions of the Association for Computational Linguistics, 10:539–554.
- Harsh Trivedi, Niranjan Bauer, Tushar Khot, and Ashish Sabharwal. 2023. Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions. In Proceedings of ACL, pages 10014–10037.
- Jingjin Wang and Jiawei Han. 2025. PropRAG: Guiding retrieval with beam search over proposition paths. In Proceedings of EMNLP, pages 6212–6227.
- Liang Wang, Haonan Chen, Nan Yang, Xiaolong Huang, Zhicheng Dou, and Furu Wei. 2025. Chain-of-retrieval augmented generation. arXiv preprint arXiv:2501.14342.
- Wenhan Xiong, Xiang Lorraine Li, Srinivasan Iyer, Jingfei Du, Patrick Lewis, William Yang Wang, Yashar Mehdad, Wen-tau Yih, Sebastian Riedel, Douwe Kiela, and Barlas Oğuz. 2021. Answering complex open-domain questions with multi-hop dense retrieval. In Proceedings of ICLR.
- 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 EMNLP, pages 2369–2380.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. ReAct: Synergizing reasoning and acting in language models. In Proceedings of ICLR.

A Judge Prompt Template

System: You are a retrieval judge for multi-hop QA. Given a query, a bridge passage, and a candidate passage, output a score from 0 to 10 for whether the candidate is the next supporting passage needed to answer the query, given the bridge context.

User:

Query: {query}
 Bridge entity 1: {entity1}
 Bridge entity 2: {entity2}
 Bridge passage: {bridge_passage}
 Candidate passage: {candidate_passage}

Score (0-10):

The judge is called once per query with all candidates in a single batched prompt (one JSON array per call), not 20 separate calls.

B SVO Generation Prompt Template

System: You generate targeted retrieval queries for multi-hop QA.

User:

Question: {question}

First-hop passage: {bridge_passage}

Generate exactly 3 targeted queries in Subject-Verb-Object form to retrieve the second supporting passage. Output JSON: {"queries": ["...", "...", "..."]}

C Entity Extraction Prompt Template

Question: {question}

Bridge passage:
{bridge_passage}

The bridge passage establishes a key intermediate entity. Based on the question and bridge, identify the TWO most relevant answer-side entities needed to answer the question.

Reply with ONLY two short entities separated by " | " (each 1-6 words, e.g. "Westminster | 1975" or "Michael Curtiz | Edith Carlmar").
No other text. If only one entity exists, repeat: "entity1 | entity1".

The response is parsed on the " | " delimiter to yield e_1 and e_2 . If only one entity is distinguishable, the same string is used for both ANN retrievals; the duplicate results are deduplicated before pool construction.

D Statistical Tests

All pairwise comparisons use a one-sided sign test. For each query, a win is recorded if condition X achieves higher R@5 than condition Y; ties are excluded. The p -value tests the null hypothesis that $P(\text{win}) = 0.5$. We report Bonferroni-corrected p -values for the subtype analysis (4 tests).

E Experiment G: Full Results

Dataset	Subtype	A1 (real bridge)	G (SVO text)	A5 (hard neg.)
MuSiQue	all	0.815	0.793	0.802
2Wiki	all	0.953	0.949	0.943
2Wiki	bridge_comp	0.876	0.868	0.850
2Wiki	compositional	0.970	0.969	0.966
2Wiki	comparison	0.996	0.996	0.992
2Wiki	inference	0.958	0.940	0.944

Table 5: Bridge text ablation. A1 = real retrieved bridge; G = SVO queries as bridge text; A5 = lowest-SVO-similarity passage as bridge (hard-negative control). On MuSiQue, $G < A5$: SVO query text performs worse than the hardest-to-retrieve pool passage, showing the bridge gain is content-driven.

Subtype	n	Flip rate	FlipWin%	Δ vs. no-flip
bridge_comparison	235	0.523	18.7%	+14.2pp
compositional	413	0.349	2.8%	+2.8pp
comparison	244	0.652	0.6%	+0.6pp
inference	108	0.583	4.8%	+4.8pp

Table 6: Flip productivity: fraction of top-1 re-rankings ($B \rightarrow C$) that improve R@5. Comparison has the highest flip rate but lowest productivity (0.6%), confirming that bridge conditioning causes correct re-rankings, not noise-driven churn. Kendall’s $\tau = -0.33$ (ns) between flip_rate ordering and $B \rightarrow C$ delta ordering falsifies the naïve “more flips \Rightarrow more gain” hypothesis.

F Experiment H: Flip Productivity

G Hyperparameter Sensitivity

α	MuSiQue	2Wiki	HotpotQA
0.05	0.812	0.950	0.985
0.10	0.815	0.953	0.986
0.15	0.813	0.951	0.988
0.20	0.810	0.948	0.984

Table 7: R@5 on tune split for $\alpha \in \{0.05, 0.10, 0.15, 0.20\}$. The method is not sensitive to α ; the top-2 values differ by $\leq 0.3\text{pp}$ on all datasets.

Chain-Disambiguation Example
<p>Query: Who is the spouse of the actor who played the Terminator?</p> <p>2-way judge top-1 $s(q, c)$: “The Terminator is a 1984 film directed by James Cameron...” (film passage — same entity mention, wrong chain)</p> <p>Bridge b: “Arnold Schwarzenegger is an Austrian-American actor and former politician...” (top hop-1 passage)</p> <p>3-way judge top-1 $s(q, b, c)$: “Maria Shriver is an American journalist and author, and the wife of Arnold Schwarzenegger...” ✓</p>

Figure 1: Bridge conditioning resolves chain ambiguity. The 2-way judge promotes a passage about the Terminator film (same entity surface, wrong chain). Conditioning on bridge b (Schwarzenegger passage) allows the judge to identify Maria Shriver as the correct second-hop target.

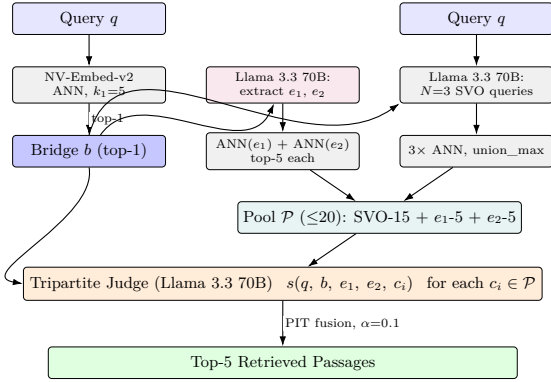


Figure 2: BridgeRAG pipeline. Hop 1 (left): query q is embedded with NV-Embed-v2 and retrieved via ANN; the top-1 passage becomes bridge b . Entity branch (centre): a Llama 3.3 70B call extracts entities e_1, e_2 from b ; each is used for an independent ANN retrieval (top-5), yielding entity-grounded candidates. SVO branch (right): q and b condition a second Llama call that generates $N=3$ SVO queries; each is embedded and retrieved ($3 \times \text{ANN}$, union_max), yielding SVO-15 candidates. Pool: $\text{SVO-15} \cup e_1\text{-5} \cup e_2\text{-5} \rightarrow \text{top-20}$. Judge: a tripartite judge scores every c_i via $s(q, b, e_1, e_2, c_i)$; scores are PIT-fused ($\alpha=0.1$) to produce the final top-5.