

Quantifying the Accuracy and Cost Impact of Design Decisions in Budget-Constrained Agentic LLM Search

Kyle A. McCleary and James M. Ghawaly

Division of Computer Science & Engineering, Louisiana State University
Baton Rouge, Louisiana, USA
jghawaly@lsu.edu, kmcccl24@lsu.edu

Abstract

Agentic Retrieval-Augmented Generation (RAG) systems combine iterative search, planning prompts, and retrieval backends, but deployed settings impose explicit budgets on tool calls and completion tokens. We present a controlled measurement study of how search depth, retrieval strategy, and completion budget affect accuracy and cost under fixed constraints. Using *Budget-Constrained Agentic Search (BCAS)*, a model-agnostic evaluation harness that surfaces remaining budget and gates tool use, we run comparisons across six LLMs and three question-answering benchmarks. Across models and datasets, accuracy improves with additional searches up to a small cap, hybrid lexical and dense retrieval with lightweight re-ranking produces the largest average gains in our ablation grid, and larger completion budgets are most helpful on HotpotQA-style synthesis. These results provide practical guidance for configuring budgeted agentic retrieval pipelines and are accompanied by reproducible prompts and evaluation settings.

Keywords: Retrieval-Augmented Generation, Agentic Search, Budget-Aware Evaluation

1. Introduction

The landscape of Retrieval-Augmented Generation (RAG) has undergone a significant shift in just one year. What began as static retrieve-and-generate pipelines has evolved into sophisticated agentic systems where Large Language Models (LLMs) autonomously plan and execute multi-step retrieval strategies (Singh et al., 2024). Agentic retrieval capabilities are now integrated into mainstream AI tools. OpenAI’s ChatGPT employs function calling for iterative search (OpenAI, 2023), Anthropic’s Claude features a dedicated research mode (Anthropic, 2024), and Google’s Gemini 2.0 includes native tool-use APIs for dynamic information gathering (Gemini Team, 2025).

However, while these advances have primarily focused on maximizing retrieval quality and answer accuracy, a critical dimension remains underexplored: *computational budget constraints*. In real-world deployments, the cost of multiple API calls, the latency of sequential searches, and the computational overhead of processing numerous retrieved documents present significant challenges.

In this setting, we frame the problem as a measurement gap. RAG is now the default strategy for grounding LLMs, and many systems combine iterative querying, intermediate reasoning, hybrid lexical + dense search, and re-ranking under “agentic search.” Existing studies typically report accuracy improvements for individual methods, often with permissive tool usage, but fewer quantify how these design knobs jointly affect both quality and cost under fixed search and token budgets across multiple models and datasets.

In this work, we quantify how common agentic design decisions, including hybrid search, lightweight

re-ranking, pre-planning, reflection, and completion limits, shift accuracy and cost when searches and tokens are capped. We conduct controlled comparisons across six LLMs and three multi-hop QA benchmarks, holding budgets constant while toggling one decision at a time.

To run these experiments we use **Budget-Constrained Agentic Search (BCAS)**, a budget-aware evaluation harness that surfaces remaining search and token allowances as explicit signals and gates tool calls accordingly. BCAS keeps the instrumentation simple: it reuses commodity prompts, avoids bespoke APIs, and records per-question search and token consumption so practitioners can reinterpret the results under their own pricing models.

1.1. Research Questions

To quantify the practical value of budget-aware planning, we subject BCAS to a large-scale study involving six LLMs and three multi-hop benchmarks, guided by three questions:

RQ1: Performance Across Model Sizes

How does model size (parameter count) impact retrieval performance?

RQ2: Budget-Aware Component Tuning

How do individual agentic-RAG components and hyper-parameters influence performance under a fixed search budget, and can those effects be exploited to optimize performance within that budget?

RQ3: Accuracy-Budget Trade-off

How does accuracy vary as we simultaneously tighten the budgets on iterative search steps and generated output tokens in an

agentic-RAG pipeline?

Recent reinforcement-learning-driven systems such as Search-R1 (Jin et al., 2025) and DeepRetrieval (Jiang et al., 2025) demonstrate large gains but incur significant training and inference overheads, so cost-aware behaviour in agentic RAG warrants closer study. The remainder of the paper reviews related work in Section 2, documents the evaluation protocol and budget controls in Section 3, presents results in Section 4, and concludes with discussion and takeaways in Sections 5 and 6.

2. Related Work

Agentic retrieval has progressed from static pipelines to autonomous research agents within a few years, providing the backdrop for BCAS.

Static pipelines. Early systems such as RAG (Lewis et al., 2020) and FiD (Izacard and Grave, 2021) established the retrieve-then-generate paradigm and demonstrated how parametric and non-parametric memories could be combined effectively.

Prompted agency. Lightweight prompting soon enabled iterative reasoning without new training. ReAct (Yao et al., 2023) introduced the “Thought, Action, Observation” loop for search planning, Self-Ask (Press et al., 2023) decomposed questions into sub-queries, and FLARE (Jiang et al., 2023) triggered retrieval on demand. Reflection-driven methods such as Self-RAG (Asai et al., 2024) and CRAG (Yan et al., 2024) further improved robustness by critiquing intermediate results.

Autonomous research agents. Recent work deepens external planning through structured reasoning and reinforcement learning. Self-Reasoning (Xia et al., 2025) trains models to analyse their own trajectories, while Search-R1 (Jin et al., 2025) and DeepRetrieval (Jiang et al., 2025) learn to formulate and schedule queries directly against search engines.

Budget-aware gap. These approaches often assume generous tool usage and many depend on bespoke training or fine-tuning. BCAS contributes an orthogonal evaluation layer: it makes compute constraints explicit at inference time, standardizes the retrieval interface, and logs cost telemetry so design choices can be compared under shared budgets. We quantify how search depth, component choices, and token limits interact across diverse LLMs in this controlled setting.

This distinction matters for how the paper should be read. We do not introduce a new trained search agent or claim a new state of the art on any benchmark. Instead, we use a common scaffold to isolate which design choices matter once searches and completion length are capped. That makes the contribution closer to controlled systems measurement

than to agent leaderboard optimization.

3. Methodology

This section provides implementation details for the BCAS framework, describes the BCAS component ablation studies, and also the iterative search count and output context size scaling studies. We use this loop only as a controlled evaluation scaffold to quantify how search depth, retrieval components, and token limits affect accuracy and cost under explicit budgets. An example of BCAS in action is shown in Figure 1, where DeepSeek V3 (0324) solves a complex HotpotQA question within a strict search budget. Our implementation and evaluation code is available at https://github.com/kmccleary3301/BCAS_RAG.

3.1. Architectural Principles

BCAS is built on three design principles that align with our research questions:

Explicit budget management (RQ3). Models receive their remaining search and token allowances at each step, nudging them toward deliberate query sequencing. The primary levers are `max_searches` and `max_total_tokens`.

Composable component pipeline (RQ2). Retrieval tools, planning hooks, and reflection routines are toggled through configuration, enabling controlled ablations of each capability.

Model-agnostic design (RQ1). Prompts describe tools in plain language rather than bespoke APIs, so any instruction-following LLM can adopt the framework without fine-tuning.

3.2. Prompt Policy

For each evaluated method, we use one prompt template shared across all model families, and we do not tune prompts per model. This choice keeps the interaction protocol fixed so component and budget effects can be compared directly. As a result, cross-model differences in this study should be interpreted as sensitivity to a shared scaffold rather than each model’s best-case tuned performance.

3.3. Budget Accounting

BCAS tracks two resources throughout each trajectory: remaining search calls and cumulative completion budget. The search cap is enforced as a hard tool constraint. Once the allowance is exhausted, the search tool is removed from the available action list and the model must answer with the information it has already gathered. Token accounting is updated after every turn from the provider-reported usage fields, and the loop stops once the configured `max_total_tokens`

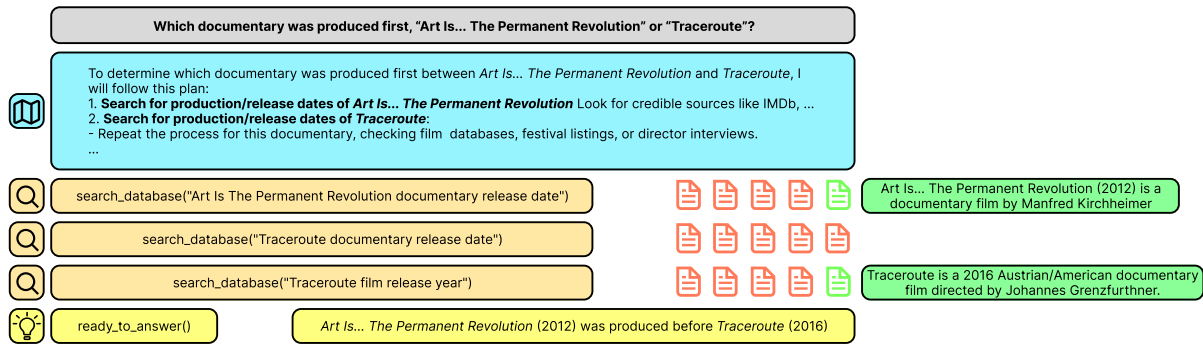


Figure 1: An example of DeepSeek V3 (0324) solving a HotpotQA ‘hard’ problem using BCAS with 4 searches and planning. This demonstrates the core principles of BCAS: the model operates within a strict 4-search budget (RQ1, RQ3) while leveraging system components like planning and iterative search to efficiently locate the answer (RQ2) and terminate early.

threshold is reached. This accounting scheme is deliberately simple. It does not model queueing delays, batching effects, or provider-specific pricing tiers, but it does give a consistent view of how each configuration converts a fixed interaction budget into answer accuracy.

3.4. The BCAS Execution Loop

The BCAS algorithm is implemented as a stateful loop that guides the LLM through a budget-aware reasoning and retrieval process. In each turn, the model:

- 1: **Reasons:** Generates a “thought” about its current progress and what information it needs next.
- 2: **Selects a Tool:** Chooses an action from a dynamically generated list of available tools. The `search_database` tool is only available if the agent is still within its `max_searches` budget.
- 3: **Executes and Observes:** The chosen tool is executed (e.g., a database search is performed), and the result (an “observation”) is returned to the model.
- 4: **Updates Budget:** Token usage is updated, and the remaining budget is recalculated for the next turn.

This loop continues until the model calls the `ready_to_answer` tool or a budget is exhausted.

Optional Pre-planning When enabled, the agent decomposes the question into a step-by-step research plan, considering available search and token budgets for more structured approaches to multi-hop questions.

Optional Reflection This periodically prompts the model to review progress and make strategic adjustments, particularly effective for smaller models in recognizing unproductive search paths and

optimizing budget usage.

These planning and reflection strategies are implemented following patterns from recent state-of-the-art works on agentic search (Xia et al., 2025; Yao et al., 2023; Asai et al., 2024), adapted to operate within explicit budget constraints.

3.5. Information Retrieval Components

To analyze component tuning (RQ2), we equip the agent with three distinct Information Retrieval (IR) strategies that can be configured for each experiment.

- **BM25 with Phrase Boosting:** Our baseline retriever uses a standard BM25 algorithm enhanced with a query parser that boosts two- and three-term phrases. This prioritizes documents with dense coverage of the query’s concepts, inspired by cover density (Clarke et al., 2000).
- **Hybrid Search (HS):** This method combines the BM25 retriever with dense vector search, offering a blend of lexical and semantic matching.
- **Re-ranking (RR):** We use a context re-ranking model to re-score and select the top five most relevant results. This can be applied on top of either BM25 or Hybrid Search.

Our IR backend is built on ParadeDB (Noël, 2024), a PostgreSQL extension. This allows us to construct a hybrid BM25 and vector search engine over our documents.

Retrieval and re-ranking pipeline. Each search request returns chunk-level candidates from the dataset-specific index. In BM25 mode, the system returns the top five BM25 chunks. In hybrid mode (HS), lexical and dense retrieval are combined and the top five chunks are returned to the agent context. In hybrid + re-ranking mode (HS+RR), we first retrieve 100 hybrid candidates, then re-score them

with a cross-encoder re-ranker and keep the top five chunks for the model context.

We keep the final context window deliberately narrow at five returned chunks so that search policy, rather than indiscriminate context expansion, remains the main source of improvement. The 100-candidate pool in the re-ranking condition is large enough to expose meaningful ordering differences while remaining cheap enough to run across the full ablation grid.

Vector Embedding

We used the `BGE-M3` (Chen et al., 2024) model for embedding document chunks and search queries during hybrid retrieval. `BGE-M3` was chosen for its broad effectiveness in multilingual benchmarks, as outlined in (Chen et al., 2024). Vector similarity was measured using the standard cosine similarity metric. In our results, we denote hybrid retriever search as **HS**.

Re-ranking

We tested the effects of re-ranking results at retrieval. We used `bge-reranker-v2-m3` for this task due to competitive performance on common benchmarks (Li et al., 2023). During testing cases involving the re-ranker, we retrieved 100 document chunks using hybrid retrieval, calculated scores for these chunks using the re-rank model, and returned the top five chunks. In our results, we denote re-ranking as **RR**.

3.6. Evaluation

We evaluate BCAS across six LLMs and three benchmarks to measure how accuracy scales with search depth and completion token count. We also evaluate the effects of various retrieval components on the accuracy of unconstrained BCAS search through ablation.

LLM Selection

We evaluated six models across different size categories to analyze performance across model capacities (RQ1): **o4-mini** (OpenAI, 2025), **DeepSeek V3 (0324 version)** (DeepSeek, 2024), **GPT-4.1-mini** (OpenAI, 2024), **Gemma 3 27B** (Team et al., 2024), **Qwen 3 14B** (Cloud, 2024), and **LLaMA 3.1 8B** (Dubey et al., 2024). These models represent diverse capabilities, costs, and architectures, all proficient in tool usage for BCAS tasks.

Grading

We use GPT-4o-mini as a binary correctness judge, following prior evidence that LLM judges can track

human answer grading reasonably well in QA settings (Zheng et al., 2023). The judge receives the question, reference answer, and model answer, and outputs a correct/incorrect label under a constrained rubric. GPT-4o-mini does not, to our knowledge, share a model family with the six evaluated answering models. In addition to literature support, we conducted an independent manual audit of 200 TriviaQA samples, 200 “hard” HotpotQA samples, and 200 2WikiMultihopQA samples. We confirmed correct LLM grading on 200/200 TriviaQA, 198/200 HotpotQA, and 197/200 2WikiMultihopQA samples. The five discrepancies were all false negatives from the judge, so reported accuracies are slightly conservative. We still treat binary judge labels as a proxy metric and discuss this limitation in Section 7.

3.7. Datasets

We performed our scaling evaluations on three diverse question-answering benchmarks:

- 1: TriviaQA** (Joshi et al., 2017) is a widely used benchmark for open-domain question-answering that requires finding specific facts, often from a single document. We use the unfiltered test set, comprising over 500K documents.
- 2: HotpotQA** (Yang et al., 2018) focuses on multi-hop reasoning, where questions require synthesizing information from multiple documents to arrive at an answer.
- 3: 2WikiMultihopQA** (Ho et al., 2020) is the most difficult of the three, demanding reasoning over multiple hops with more complex entity relationships.

These datasets were chosen to represent a range of complexities, from single-fact retrieval to complex multi-step reasoning. Each sample contains a question, answer, and relevant documents.

The documents for each dataset were partitioned into separate collections, and retrieval was restricted to the appropriate collection during testing. For all cases, we ensured that each model was evaluated on the same set of samples, verified by hashing the dataset before evaluation. Due to occasional API failures during evaluation, final sample counts ranged from 467 to 537 samples across datasets, with counts held constant within each dataset. Each run logs per-question metadata, including `searches_used`, `tokens_in`, `tokens_out`, `early_stop`, and grader verdicts, along with a cost ledger so practitioners can recompute dollar metrics under alternative price sheets.

3.8. Ablation Study

We perform an ablation study of BCAS to measure the effect of optional features and altered search

tooling on retrieval accuracy. This was performed on all six models using 467 testing samples from HotpotQA (samples released with our code). Our baseline configuration has `max_total_tokens` set to 16000, `max_searches` set to “unlimited”, the search endpoint set to BM25 only, pre-planning and reflection disabled, and context + search hints enabled. All of our feature adjustments are shown by their net effect on the baseline configuration’s accuracy. We measure the effect of the following:

1. Pre-planning
2. Reflection (no Pre-planning)
3. Pre-planning *and* Reflection
4. Hybrid BM25 + Vector Search (n=5).
5. Hybrid BM25 + Vector Search (n=100) with re-ranking of results (choose top 5).

As discussed in the limitations, we report the full factorial ablation only on HotpotQA, where multi-hop synthesis makes component interactions most visible. We do not claim that the exact ordering of ablation gains is universal across datasets.

3.9. Search Scaling Study

To understand how iterative search capacity affects performance (RQ3), we systematically varied the `max_searches` parameter across four configurations: 1, 2, 3, and unlimited searches. All experiments used hybrid search returning 5 results per query with `max_total_tokens` = 16,000, and both preplanning and reflection disabled. This standardized baseline ensures that observed performance differences are attributable to search budget variations rather than confounding factors. Additionally, we evaluated two enhanced configurations using unlimited search: with preplanning enabled, and with both preplanning and reflection enabled. This design yields six total configurations per model-dataset combination, allowing us to trace both the core search scaling effect and its interaction with strategic reasoning components.

3.10. Context Scaling Study

To examine how token budget constraints influence agent behavior (RQ3), we evaluated performance across five context limits: 500, 1K, 2K, 4K, and 16K tokens using hybrid search with unlimited searches, ensuring that performance variations reflect context budget effects rather than search limitations. We excluded o4-mini due to its reasoning model architecture, where completion tokens include the internal reasoning trace, making artificially low token limits likely to truncate responses mid-reasoning rather than demonstrating budget-aware behavior. Enhanced reasoning configurations are included as appropriate and correspond to identical conditions from the search scaling study. Thus comparisons

across context limits exclude o4-mini to avoid conflating budget policy with model-internal reasoning tokenization.

4. Results

We find three consistent patterns across models and datasets: (i) accuracy improves reliably up to about three searches; (ii) hybrid retrieval with re-ranking yields the most consistent gains among components; and (iii) larger token generation budgets primarily help multi-hop questions with challenging information synthesis requirements. Figure 2 displays the scaling trends; Figure 3 displays the ablation study results; Table 1 reports scaling metrics and Table 2 details the component ablations, addressing RQ1 through RQ3.

4.1. Performance Across Model Capacities (RQ1)

Iterative search narrows capacity gaps: smaller models approach or exceed larger models’ single-search scores when allowed additional searches. For example, on HotpotQA, Qwen 3 14B with unlimited searches and planning attains 75.33%, exceeding o4-mini’s single-search 70.17% but not surpassing o4-mini at two searches (86.05%; Table 1). Similar patterns appear on TriviaQA and 2WikiMultihopQA where multi-step search raises smaller models’ accuracy relative to single-step baselines (Figure 2).

4.2. Budget-Aware Component Tuning (RQ2)

On HotpotQA, re-ranking on top of hybrid search yields the largest average improvement across models (+9.29 points), with hybrid search alone also beneficial (+6.36), while planning and reflection improve smaller models by 4 to 12 points and have limited effect on o4-mini (Table 2). We treat these effect sizes as task-specific to HotpotQA and do not claim that the exact ordering transfers unchanged to TriviaQA or 2WikiMultihopQA.

4.3. Accuracy-Budget Trade-off (RQ3)

Increasing search steps yields consistent gains up to roughly three searches, after which returns plateau; context scaling shows limited change for TriviaQA, a clear lift for HotpotQA when going from 4k to 16k tokens, and minimal gains for 2WikiMultihopQA (Figure 2). Full per-configuration results appear in Table 1.

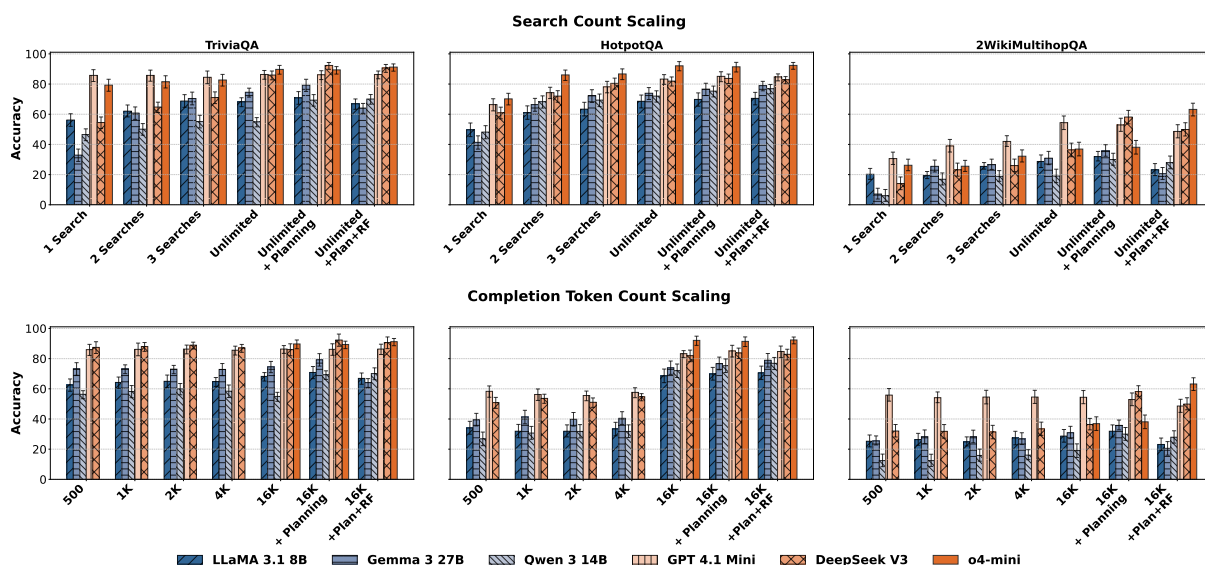


Figure 2: Search scaling and context scaling performance across TriviaQA, HotpotQA, and 2WikiMulti-hopQA. The "Planning" columns show the impact of adding pre-planning and reflection (RF) strategies. This figure provides the primary data for analyzing the accuracy-budget trade-off (RQ3). Confidence intervals computed using 95% Wilson score interval (Wilson, 1927).

Model	Completion Tokens				Searches (16K Token Limit)				Plan	Plan+RF
	500	1K	2K	4K	1	2	3	Unlimited		
TriviaQA										
o4-mini	-	-	-	-	79.33	81.56	82.68	89.72	89.39	91.25
DeepSeek V3	87.45	88.01	88.95	87.08	54.51	64.78	71.07	85.98	92.34	90.65
GPT-4.1-mini	85.96	86.14	86.33	85.58	85.74	85.74	84.49	86.36	86.17	86.36
Gemma 3 27B	73.22	73.22	73.03	72.85	32.91	60.80	70.44	74.58	79.25	64.11
Qwen 3 14B	56.37	58.24	59.93	58.43	46.54	50.10	55.14	54.95	69.35	70.09
LLaMA 3.1 8B	62.55	64.23	64.98	64.79	56.18	62.05	68.76	68.22	70.84	66.92
HotpotQA										
o4-mini	-	-	-	-	70.17	86.05	86.70	92.92	91.43	92.36
DeepSeek V3	50.86	53.65	51.07	54.72	60.95	72.02	80.48	77.94	83.74	82.87
GPT-4.1-mini	58.33	56.25	55.56	57.64	66.38	74.30	78.16	80.47	85.10	84.73
Gemma 3 27B	39.40	41.33	39.83	40.47	41.38	66.38	72.20	68.52	76.66	78.96
Qwen 3 14B	26.82	30.69	31.76	31.76	48.18	68.31	69.38	64.45	75.33	76.91
LLaMA 3.1 8B	34.05	31.91	31.91	33.40	49.68	61.03	63.38	65.31	69.83	70.52
2WikiMulti-hopQA										
o4-mini	-	-	-	-	26.18	25.54	32.19	36.91	38.12	63.20
DeepSeek V3	27.68	26.18	26.39	29.18	14.07	23.16	25.97	36.36	58.15	49.89
GPT-4.1-mini	55.84	54.11	54.55	54.55	30.62	38.97	41.97	54.39	52.89	48.61
Gemma 3 27B	25.48	28.27	28.27	26.77	7.13	25.49	26.57	30.89	35.76	20.56
Qwen 3 14B	12.42	12.42	11.56	13.28	6.00	16.92	18.84	19.06	29.98	28.66
LLaMA 3.1 8B	25.27	26.34	24.84	27.41	20.22	19.35	25.43	28.70	31.91	23.13

Table 1: Full accuracy data (%) for all models across all datasets and BCAS budget configurations. Highest values for each row are bolded. Note that for TriviaQA, the performance was nearly flat across the board.

Model	Plan	RF	Plan+RF	HS	HS+RR
LLaMA 3.1 8B	+4.52	+6.33	+5.21	+5.73	+8.80
Gemma 3 27B	+8.14	+7.37	+10.43	+8.93	+13.01
Qwen 3 14B	+10.87	+5.19	+12.45	+12.36	+17.86
GPT-4.1-mini	+4.63	+3.89	+4.26	+4.07	+5.93
DeepSeek V3	+5.79	+4.55	+4.92	+6.23	+9.02
o4-mini	+0.17	+0.92	+1.10	+0.83	+1.10
Average	+5.69	+4.71	+6.40	+6.36	+9.29

Table 2: Impact of retrieval method components on HotpotQA accuracy (%). This analysis informs component tuning under budget constraints (RQ2).

5. Discussion

The results suggest a practical budget policy in our grid: allocate additional searches first (typically

up to three), improve evidence quality via hybrid retrieval with re-ranking, and expand token generation limits primarily for multi-hop reasoning tasks. This ordering captures most of the observed gains while bounding compute.

One consequence of this framing is that the paper supports configuration decisions more directly than it supports claims about a universally best agent. Every result is conditioned on the same prompt scaffold, retriever stack, and budget rules. That makes the comparisons useful for deciding how to spend a fixed allowance, but it also means the paper should be read as a study of design leverage under control, not as a leaderboard over heterogeneous systems.




Figure 3: Ablation study of BCAS features on HotpotQA: We measure the net effect on retrieval accuracy of different features compared to BCAS baseline with a BM25 retriever. 95% confidence intervals computed using Newcombe Method 9 (Newcombe, 1998).

5.1. Dataset-dependent Budget Leverage

The most unusual effect of scaling the completion token budget appears on HotpotQA. While accuracy remains largely flat for TriviaQA and 2WikiMultihopQA, HotpotQA exhibits a sharp increase when the allowance rises from 4k to 16k tokens. We hypothesize that the datasets differ in the *type* of difficulty they pose: 2WikiMultihopQA seems primarily constrained by retrieval, because locating the correct evidence is hard, whereas HotpotQA appears less retrieval-bound and more limited by the reasoning required to synthesise multiple pieces of evidence into a final response. This would explain why accuracy improves more distinctly with higher iterative search count on 2WikiMultihopQA compared to HotpotQA, whereas HotpotQA sees major gains for higher token count, unlike 2WikiMultihopQA.

These patterns also reveal an asymmetric relationship between search depth and context allocation. Accuracy increases monotonically with search steps across all models and datasets, with diminishing returns above three searches. However, context scaling shows that *constrained completion tokens can outperform unlimited contexts when the model manages multiple searches*. This phenomenon is most evident on 2WikiMultihopQA, where GPT-4.1-mini achieves 55.84% accuracy with 500 completion tokens but only 30.62% with a single search at 16K tokens. Similarly, DeepSeek V3 reaches 27.68% with 500 tokens versus 14.07% with one search. Under tight token constraints, models adopt succinct behaviour, making efficient tool calls that enable multiple searches within budget. Conversely, unlimited single searches encourage verbose responses that consume the entire token budget without additional information gathering.

In our experiments, allocating budget toward search depth often outperforms allocating the same

budget toward longer single-pass generations. For practitioners, moderately restrictive token budgets (500 to 2K) combined with multiple search opportunities can yield better accuracy than generous single-search configurations.

5.2. The o4-mini Anomaly

o4-mini presents the most puzzling behavioural pattern in our evaluation across different model capacities. Despite being among the top performers overall, it shows minimal response to most system enhancements. On HotpotQA, all enhancement strategies yield less than 1.10 percentage points improvement, a stark contrast to other models that see 4 to 18 point gains.

This resistance to enhancement may stem from o4-mini being a reasoning model with built-in chain-of-thought processing. As such, external planning components are often redundant. An exception appears on 2WikiMultihopQA, where combining planning with reflection produces a 25-point increase (38.12% to 63.20%). This suggests that while basic pre-planning may be redundant, reflection can create opportunities for mid-search strategy revision that complement o4-mini’s internal reasoning process.

5.3. Component Hierarchy and Model Variety

The ablation results reveal that component effectiveness varies dramatically by model. Re-ranking with hybrid search provides the most consistent gains (6 to 18 percentage points for most models), making it the most reliable enhancement across model classes. However, the magnitude of improvement varies significantly by model, with smaller ones tending to benefit more from retrieval enhancements than larger ones.

Pre-planning strategies show a similar pattern: smaller models (LLaMA 3.1 8B, Qwen 3 14B) gain 4 to 12 percentage points from planning components, while larger models show more modest improvements. This suggests that external scaffolding becomes less valuable for models with stronger inherent capabilities. Combining pre-planning with reflection slightly degrades performance for some models, indicating that adaptive re-planning can be counter-productive; following an initial plan appears more effective than continuously revising strategy mid-search for conventional models.

5.4. Deployment Recommendations

Based on our evaluation trends, we recommend prioritising iterative search depth, followed by component selection, then token generation limits in budget allocation. For cost-sensitive deployments,

LLaMA 3.1 8B (the smallest model evaluated) with three searches and BM25 already surpasses the single-search baselines of larger systems, except for o4-mini and GPT-4.1-mini. Adding pre-planning and reflection allows it to match or exceed every model limited to one search while remaining dramatically cheaper.

For maximum accuracy, o4-mini with both planning strategies reaches 63.20% on 2WikiMultihopQA, the highest performance observed, but this requires targeted configuration for complex tasks. Models without native chain-of-thought capabilities benefit more from hybrid retrieval and re-ranking than from elaborate reasoning scaffolds, so enhancements should match the underlying capability profile.

The practical implication is that the first deployment question should usually be “How many searches can we afford?” rather than “How large a completion window should we buy?” In our grid, the clearest gains come from one additional well-targeted retrieval step or from better ranking of the retrieved evidence. Longer generations help chiefly when the retrieved material already contains the right pieces and the model still needs room to synthesize them.

6. Conclusion

We present BCAS, a budget-aware evaluation harness for studying agentic retrieval under explicit search and token limits. Using six LLMs across TriviaQA, HotpotQA, and 2WikiMultihopQA, we quantify how search depth, retrieval configuration, and completion budgets trade off accuracy and cost.

Across our experiments, three patterns are consistent. First, increasing search depth improves accuracy up to a small cap in most settings, with diminishing returns after roughly three searches. Second, hybrid lexical+dense retrieval with lightweight re-ranking provides the largest average accuracy gains in our HotpotQA ablation grid. Third, larger completion budgets matter most on synthesis-heavy HotpotQA settings, while retrieval-easier settings show faster saturation.

These results provide a practical configuration order for budgeted deployments: expand search depth first, improve evidence quality through retrieval and re-ranking, and then raise completion budgets when synthesis demands justify the added cost. By making tool and token budgets explicit and logging usage outcomes, BCAS supports reproducible accuracy-cost analysis under realistic constraints.

6.1. Future Directions

Future work should add direct non-agentic baselines, broaden component ablations beyond Hot-

potQA, and expand evaluation to open-web and multimodal environments. A second direction is tighter runtime reporting, including queueing and throughput behavior under production concurrency constraints.

7. Limitations

Our study focuses on factual QA benchmarks over static benchmark corpora rather than open-web retrieval. The reported trends may not transfer directly to domains with rapid knowledge drift, multilingual corpora, multimodal evidence, or long-horizon interactive tasks.

The benchmarks also emphasize short final answers with known references. That is useful for controlled comparison, but it leaves out settings where success depends on citation fidelity, extended synthesis, interactive clarification, or user-specific constraints. A policy that is attractive for benchmark QA may therefore need different stopping rules or retrieval breadth in production workflows.

The full component ablation is reported on HotpotQA only. We do not claim that the exact ordering of component gains is universal across datasets; TriviaQA and 2WikiMultihopQA may shift the relative value of retrieval quality, search depth, and completion budget.

For each method, we use a single prompt template across model families instead of per-model prompt tuning. This improves control for comparative measurement but may under-optimize some models, so cross-model comparisons should be interpreted as sensitivity to a shared scaffold rather than each model’s best-case performance.

Answer correctness is measured with a binary LLM judge at scale. Although our manual audit indicates high agreement, binary grading can miss partial correctness and reasoning quality nuances.

Cost and latency estimates depend on provider pricing, rate limits, and runtime settings. We therefore report token and search usage as stable budget proxies alongside cost estimates, and treat cross-provider cost comparisons as approximate.

We do not include a pure non-agentic single-pass retrieve-then-read baseline in this version. The presented comparisons therefore isolate differences among budgeted agentic policies and do not directly estimate the incremental value of multi-turn agency over one-shot retrieval.

We also do not report head-to-head results against RL-trained open-web search agents. Those systems differ in training regime and retrieval environment, so direct comparison would require aligned tools, corpora, and model training conditions.

8. Data and Code Availability

Code, evaluation harness configurations, prompts, and analysis scripts used in this study are available in the BCAS repository: https://github.com/kmccleary3301/BCAS_RAG. We also provide dataset split hashes and evaluation outputs needed to reproduce the reported plots and tables under the same budget settings.

9. Bibliographical References

- Anthropic. 2024. Introducing claude’s research mode. <https://www.anthropic.com/research>. Accessed: 2025-07-01.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. Self-rag: Self-reflective retrieval-augmented generation. In *International Conference on Learning Representations (ICLR)*.
- Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. [Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation](#).
- Charles L.A. Clarke, Gordon V. Cormack, and Elizabeth A. Tudhope. 2000. [Relevance ranking for one to three term queries](#). *Information Processing & Management*, 36(2):291–311.
- Alibaba Cloud. 2024. [Qwen technical report](#). Accessed: 2025-01-15.
- DeepSeek. 2024. [Deepseek-v3 technical report](#). Accessed: 2025-01-15.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Gemini Team. 2025. [Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities](#). Technical report, Google DeepMind. Accessed: 2025-07-01.
- 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 the 28th International Conference on Computational Linguistics*, pages 6609–6625, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Gautier Izacard and Edouard Grave. 2021. Leveraging passage retrieval with generative models for open domain question answering. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 874–880.
- P. Jiang, J. Lin, L. Cao, R. Tian, S. Kang, Z. Wang, J. Sun, and J. Han. 2025. Deepretrieval: Hacking real search engines and retrievers with large language models via reinforcement learning. *arXiv preprint arXiv:2503.00223*.
- 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 the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- B. Jin, H. Zeng, Z. Yue, J. Yoon, S. Arik, D. Wang, H. Zamani, and J. Han. 2025. Search-r1: Training llms to reason and leverage search engines with reinforcement learning. *arXiv preprint arXiv:2503.09516*.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, Vancouver, Canada. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474.
- Chaofan Li, Zheng Liu, Shitao Xiao, and Yingxia Shao. 2023. [Making large language models a better foundation for dense retrieval](#).
- Robert G. Newcombe. 1998. [Improved confidence intervals for the difference between binomial proportions based on paired data](#). *Statistics in Medicine*, 17(22):2635–2650.
- Philippe Noël. 2024. [Paradedb](#). <https://github.com/paradedb/paradedb>. Accessed: 2024-7-01.
- OpenAI. 2023. Function calling and other api updates. <https://platform.openai.com/docs/guides/function-calling>. Accessed: 2025-07-01.
- OpenAI. 2024. [Introducing gpt-4.1 in the api](#). Accessed: 2025-07-04.

- OpenAI. 2025. [Openai o3 system card](#). Accessed: 2025-06-23.
- 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*.
- A. Singh et al. 2024. Agentic retrieval-augmented generation: A survey on agentic rag. *arXiv preprint arXiv:2501.09136*.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. [Gemma 2: Improving open language models at a practical size](#).
- Edwin B Wilson. 1927. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 22(158):209–212.
- Yuan Xia, Jingbo Zhou, Zhenhui Shi, Jun Chen, and Haifeng Huang. 2025. [Improving retrieval augmented language model with self-reasoning](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(24):25534–25542.
- Shi-Qi Yan, Jia-Chen Gu, Yun Zhu, and Zhen-Hua Ling. 2024. Crag: Corrective retrieval augmented generation. *arXiv preprint arXiv:2401.15884*.
- 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 *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- 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 *International Conference on Learning Representations (ICLR)*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#).