

# 000 001 002 003 004 005 HIERARCHICAL SEQUENCE ITERATION FOR HETERO- 006 GENEous QUESTION ANSWERING 007 008 009

010 **Anonymous authors**  
011

012 Paper under double-blind review  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029

## ABSTRACT

030 Retrieval-augmented generation (RAG) remains brittle on multi-step questions  
031 and heterogeneous evidence sources, trading accuracy against latency and to-  
032 ken/tool budgets. This paper introduces **Hierarchical Sequence (HSeq) Itera-  
033 tion for Heterogeneous Question Answering**, a unified framework that (i) lin-  
034 earize documents, tables, and knowledge graphs into a reversible hierarchical se-  
035 quence with lightweight structural tags, and (ii) perform structure-aware iteration  
036 to collect just-enough evidence before answer synthesis. A Head Agent provides  
037 guidance that leads retrieval, while an Iteration Agent selects and expands HSeq  
038 via structure-respecting actions (e.g., parent/child hops, table row/column neigh-  
039 bors, KG relations); Finally the head agent composes canonicalized evidence to  
040 genearte the final answer, with an optional refinement loop to resolve detected con-  
041 tradictions. Experiments on HotpotQA (text), HybridQA/TAT-QA (table+text),  
042 and MetaQA (KG) show consistent EM/F1 gains over strong single-pass, multi-  
043 hop, and agentic RAG baselines, alongside higher efficiency. Beyond aggregate  
044 metrics, HSeq exhibits three key advantages: (1) a **format-agnostic unification**  
045 that enables a single policy to operate across text, tables, and KGs without per-  
046 dataset specialization; (2) **guided, budget-aware iteration** that reduces unneces-  
047 sary hops, tool calls, and tokens while preserving answer quality; and (3) **evidence**  
048 **canonicalization for reliable QA**, improving consistency and auditability of the  
049 generated answers.  
050

## 051 1 INTRODUCTION

052 Large language models (LLMs), such as ChatGPT (Achiam et al., 2023), LLaMA (Dubey et al.,  
053 2024), Falcon (Zuo et al., 2025), have been increasingly relying on retrieval-augmented generation  
054 (RAG) to ground answers in external evidence. With reliable supplementary knowledge offered  
055 factual errors are reduced, especially in domain-specific questions, leading to higher accuracy and  
056 fewer hallucinations (Zhu et al., 2021b; Gao et al., 2023; Zhao et al., 2024). Yet state-of-the-art  
057 pipelines, remain brittle on multi-step questions and heterogeneous sources, and still struggle to  
058 cope with the following challenges:  
059

060 **C<sub>1</sub> : Coverage in Single-pass Retrievers:** Single-pass pipelines (retrieve-*k* then generate) (Luo  
061 et al., 2023; Glass et al., 2022) focus on isolated retrieval and generation tasks. Although they  
062 can be setup and achieve data retrieval quickly, they struggle to trace complete evidence chains:  
063 dense retrievers, typically trained for pointwise recall and re-ranking, often lack path coverage;  
064 chunking heuristics fragment long documents and break discourse; long-context prompting shifts  
065 budget toward tokens irrelevant to the final answer and provides no explicit *sufficiency* signal.  
066

067 **C<sub>2</sub> : Uncontrolled iteration and latency in multi-agent systems:** With multi-agent collaboration  
068 and reasoning, agentic systems (Liu et al., 2025; Yang et al., 2025; Chen et al., 2025) easily explode  
069 the search space and can achieve multi-step reasoning. However they may fall with branchy plans,  
070 repeated web/file calls, and verbose chain-of-thought prompts, yielding unpredictable token/tool  
071 costs and latency; termination is often heuristic, leading to premature answers or extra wasted loops  
072 with budgets decoupled from the *evidence actually inspected* (Singh et al., 2025).  
073

074 **C<sub>3</sub> : Heterogeneity across formats:** Free text, relational tables, and KGs typically require distinct  
075 indices, retrievers, prompt styles, and controller logic, preventing policy reuse and complicating  
076

054 training and deployment. Although existing heterogeneous RAG systems (Yu, 2022; Christmann  
 055 & Weikum, 2024) are available to deal with multiple formats of data, they may still face issues in  
 056 either weak alignment across representations or lossy and non-reversible serialization that obscures  
 057 provenance and blocks faithful reconstruction.

058 **Hierarchical Sequence Iteration (HSEQ)** for Heterogeneous Question Answering introduces a re-  
 059 versible *hierarchical sequence* interface that linearizes documents, tables, and KGs into a sequence  
 060 of typed segments with lightweight structure (e.g., parent/child locality, offsets or coordinates, min-  
 061 imal schema/time tags). An iteration policy operates on this unified substrate using short, budgeted  
 062 steps: at each step it selects a few promising segments and predicts whether the accumulated set  
 063 is sufficient to answer. A concise *guidance* plan—produced by a lightweight planner or a heuristic  
 064 template—acts as a soft prior over which regions to probe first and when to stop. Once sufficiency is  
 065 predicted, the selected segments are canonicalized into a compact, provenance-preserving package  
 066 consumed by a head module to produce the final answer; an optional verifier can trigger a brief  
 067 refinement if contradictions are detected.

068 To address above issues, this paper introduces **HSEQ**, a **Hierarchical Sequence Iteration Sys-  
 069 tem** that first recasts heterogeneous knowledge source into a *single, LLM-native interface*, then  
 070 turning retrieval into a *guided, budget-aware iterative process*. The reversible HSEQ interface lin-  
 071 earizes documents, tables, and KGs into a sequence of typed segments with lightweight structure  
 072 (e.g., parent/child locality, offsets or coordinates, minimal schema/time tags). An iteration policy  
 073 operates on this unified substrate using short, budgeted steps: at each step it selects a few promis-  
 074 ing segments and predicts whether the accumulated set is sufficient to answer. A concise *guidance*  
 075 plan—produced by a lightweight planner or a heuristic template—acts as a soft prior over which  
 076 regions to probe first and when to stop. Once sufficiency is predicted, the selected segments are  
 077 canonicalized into a compact, provenance-preserving package consumed by a head module to pro-  
 078 duce the final answer; an optional verifier can trigger a brief refinement if contradictions are detected.  
 079 Specifically, our **key contributions** are as followed:

- 080 • **Unified, reversible interface.** A hierarchical sequence representation that standardizes  
 081 text, tables, and KGs with lightweight structure and provenance, enabling a single con-  
 082 troller to operate across formats.
- 083 • **Guided, budget-aware iteration.** A learned selection policy with an explicit sufficiency  
 084 signal that concentrates computation on *evidence actually inspected*, delivering predictable  
 085 latency under token/tool budgets.
- 086 • **Canonicalized evidence for reliable QA.** A compact, provenance-preserving evi-  
 087 dence package that improves answer synthesis and auditability, and supports optional  
 088 contradiction-driven refinement.

## 091 2 RELATED WORK

093 **LLM Finetuning** Large Language Models (LLMs) often adopt finetuning to unlock their capabili-  
 094 ties for downstream applications, like medical (Goyal et al., 2024), economic Guo & Yang (2024),  
 095 or human activity recognition Li et al. (2024). To enhance finetuning efficiency, methods like quan-  
 096 tization (Dettmers et al., 2022) parameter efficient fine tuning (Hu et al., 2022; Dettmers et al., 2023;  
 097 Li & Liang, 2021) can be applied.

099 **Retrieval Augmented Generation** RAG systems help LLMs retrieve extra knowledge according  
 100 to queries and thereby improving the accuracy of LLM response (Fan et al., 2024), with no necessity  
 101 to finetune the model. External databases ensure knowledge offered is domain-specific and timely,  
 102 adding reliability and interpretability (Lewis et al., 2020; Jiang et al., 2023). **Accuracy** of knowledge  
 103 retrieval and **quality** of responses are two key factors for RAG systems evaluation (Yu et al., 2024).  
 104 Apart from text, table, or html sources (Guo et al., 2024b; Chan et al., 2024; Jin et al., 2025),  
 105 recent researches have combined graph-structured data into RAG systems(GraphRAG) to improve  
 106 the efficiency of knowledge interpretability by capturing relationships between entities and utilizing  
 107 triplets as the primary data source (Edge et al., 2024; Peng et al., 2024; Hu et al., 2024; Mavromatis  
 & Karypis, 2024).

108 **Multi Agent QA system** LLM-based Multi-Agent Systems (MASSs) enable groups of intelligent  
 109 agents to coordinate and solve complex tasks collectively at scale, transitioning from isolated mod-  
 110 els to collaboration-centric approaches (Tran et al., 2025). Agents can cooperate with each other  
 111 for tasks like code generation (Hong et al., 2024; Islam et al., 2024), decision making (Nascimento  
 112 et al., 2023; Shinn et al., 2023), while competitions among agents are applied on gaming environ-  
 113 ment Wang et al. (2022) or question answering (Puerto et al., 2021). By interacting with each other,  
 114 the system can be used for both problem solving or world simulation (Guo et al., 2024a)

115  
 116  
 117 **Structural and unified RAG interfaces.** Beyond standard text-centric RAG, a line of work in-  
 118 troduces *structural* or *unified* retrieval layers. Graph-based RAG systems construct heterogeneous  
 119 or chunk-level graphs where nodes represent passages, entities, or sections, and edges encode se-  
 120 mantic or hyperlink connectivity; retrieval then propagates over this graph to improve multi-hop  
 121 reasoning and global coverage (Wu et al., 2024; Huang et al., 2025; Luo et al., 2025). Other sys-  
 122 tems build hierarchical or modular indices over mixed document formats, or define unified data  
 123 schemas for training language agents and tools, but still operate over opaque contexts at inference  
 124 time (Reynolds & Corrigan, 2024; Liu et al., 2025; Chen et al., 2024). These approaches share the  
 125 intuition that adding structure on top of unstructured text helps reasoning, but typically (i) collapse  
 126 different modalities into a single graph or index without a *reversible*, modality-preserving represen-  
 127 tation; (ii) use structure primarily for neighborhood expansion or re-ranking rather than as a generic  
 128 segment schema with explicit level, parent, and alignment fields; and (iii) do not provide formal  
 129 guarantees on faithful reconstruction or budgeted selection cost under an explicit sufficiency-based  
 130 stopping rule.

131 While existing RAG-based methods still suffered from limitation mentioned above, there is a rising  
 132 need for RAG interfaces that (i) preserve modality-specific structure in a *reversible* way rather than  
 133 collapsing all sources into an opaque graph or index; (ii) expose a generic, LLM-native segment  
 134 schema with explicit level, parent, and alignment fields so that a single controller can navigate  
 135 text, tables, and KGs uniformly; and (iii) couple this interface with an explicit sufficiency-aware,  
 136 budget-controlled selection process, so that evidence gathering is both auditable and aligned with  
 137 resource constraints. HSEQ is designed to meet these requirements by treating all sources as typed,  
 138 provenance-aware segments in a hierarchical sequence and pairing this representation with a learned  
 sufficiency head and budget-aware iteration policy.

### 140 3 HSEQ: A MULTI-AGENT HETEROGENEOUS QUESTION ANSWERING 141 FRAMEWORK

#### 142 3.1 BACKGROUND AND SETUP

143 **Heterogeneous QA with budgets.** Given a natural-language query  $q$  and a heterogeneous corpus  
 144  $D = \{(x_j, m_j)\}_{j=1}^N$  with modality  $m_j \in \{\text{text}, \text{table}, \text{kg}\}$ , the goal is to produce an answer  
 145  $y \in \mathcal{Y}$  and optional supporting evidence  $E \subseteq D$  while satisfying resource budgets  $B$  (tokens, tool  
 146 calls, latency, steps). Let  $E^*$  denote a *minimally sufficient* evidence set for  $q$  in  $D$  under a fixed  
 147 answerer.

148 **From retrieval to guided iteration.** We recast retrieval as a short sequence of structure-aware se-  
 149 lections under an explicit sufficiency criterion. A modality-aware adapter  $\tau$  converts  $D$  into a single  
 150 hierarchical sequence  $S_h = \tau(D)$ . A learned iteration policy  $\pi_\theta$  interacts with  $(q, S_h)$  to accumu-  
 151 late a compact evidence set  $M^*$  under budgets  $B$ , guided by a concise plan  $g$ . A canonicalizer  $\kappa$   
 152 packages  $M^*$  for a head module  $\mathcal{H}$ , which produces the final answer. This preserves the familiar  
 153 RAG workflow while adding a principled stopping signal and a unified interface across modalities.

#### 154 3.2 HSEQ ARCHITECTURE

155 The proposed system couples a unified *hierarchical sequence*(HSEQ) representation with an iter-  
 156 ation policy and a head module  $\mathcal{H}$  for answer synthesis. Let  $q$  denote a user query and  $D$  a het-  
 157 erogeneous corpus. A modality-aware adapter  $\tau$  converts  $D$  into a single hierarchical sequence

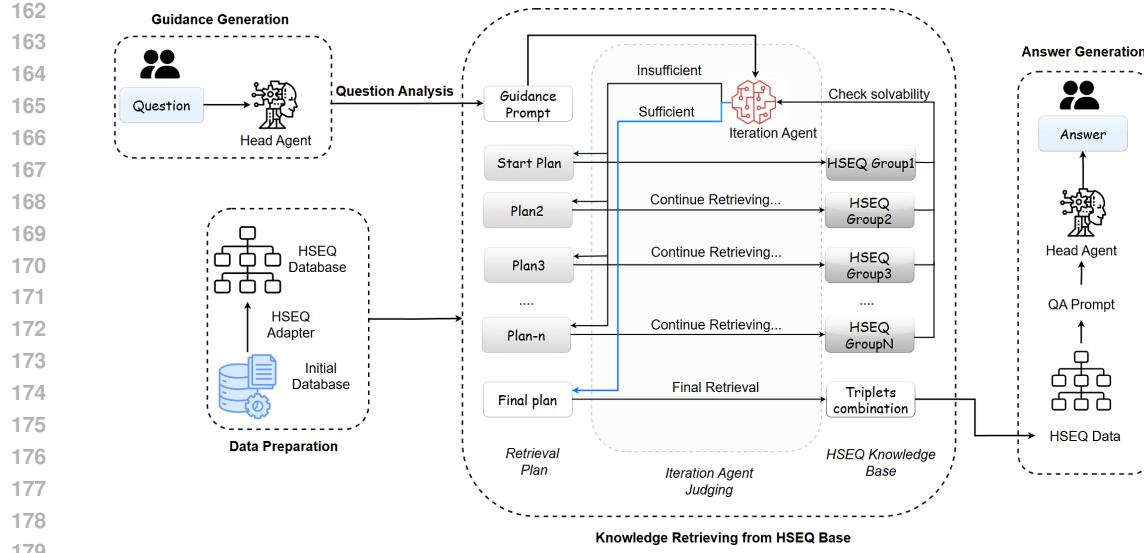


Figure 1: HSEQ overview. (i) HSEQ-A linearizes heterogeneous sources into  $S_h$  with level tags, parent pointers, and standardized metadata; (ii) HSEQ-I iterates over a windowed stream of segments under budgets, guided by  $\Phi$ , and queries  $\Phi$  for sufficiency; (iii)  $\kappa$  compacts  $M_t$  into provenance-preserving evidence; (iv) HSEQ-H produces the final answer and optionally triggers a brief refinement if inconsistencies are detected.

$S_h = \tau(D)$ . An iteration module  $\pi_\theta$  operates on  $(q, S_h)$  and maintains an evolving evidence set  $M_t$ . The final evidence  $M^*$  is then canonicalized by  $\kappa$  and passed to  $\mathcal{H}$  for final answer generation. The end-to-end mapping is summarized as

$$F = (\tau, \pi_\theta, \Phi, \kappa, \mathcal{H}), \quad F(q, D) = \mathcal{H}(q, \kappa(M^*)), \quad (1)$$

Specifically, during iteration module  $\pi_\theta$  selecting and expanding segments on  $S_h$ , the budget-aware sufficiency criterion  $\Phi$  and the budget state  $B_t$  (tokens, tool calls, steps) functioned inside the module to decide when the accumulated evidence is adequate for answering as well as triggering potential early stopping.

$$S_h = \tau(D), \quad M^* = \pi_\theta(q, S_h; \Phi, B). \quad (2)$$

After the iteration,  $\kappa$  maps raw segments  $M^*$  to a normalized evidence package consumable by  $\mathcal{H}$ . The same policy  $\pi_\theta$  is shared across modalities due to the common interface of  $S_h$ .

Generally, to achieve iteration through an unified data structure building from heterogeneous data sources, the HSEQ framework consists of three key modules: HSEQ-Adapter (HSEQ-A), HSEQ-Iterator (HSEQ-I), and HSEQ-Head (HSEQ-H).

### 3.3 HSEQ-ADAPTER(HSEQ-A)

The HSEQ-Adapter is build to produce unified structure(HSEQ  $S_h$ ) that exposes locality (parent/child), alignment (span or coordinates), and lightweight semantics (time, schema, language) in a modality-agnostic format, while remaining *reconstructable*. Formally, each item  $x_j$  is mapped by a modality-specific adapter  $\tau_{m_j}$  to a finite set of segments  $\tau_{m_j}(x_j) \subset \mathcal{S}$  and then concatenated:

$$S_h = \bigsqcup_{j=1}^N \tau_{m_j}(x_j) \in \mathcal{S}^*, \quad \mathcal{S} \ni s = (\text{id}(s), \ell(s), p(s), c(s), \mu(s)). \quad (3)$$

Here  $\ell(s)$  is a level tag matching the raw content, including sentence, paragraph, table, triplet, etc., while  $p(s)$  is a parent pointer recording the roots.  $c(s)$  is compact human-readable content, and  $\mu(s)$  is metadata with fixed keys to record content attributes.

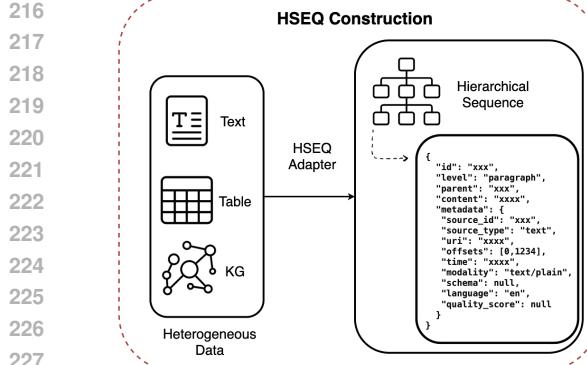


Figure 2: HSEQ  $S_h$  construction: Different modalities of data are transformed into unified sequence by HSEQ-A files.

The single, modality-aware adapter converts heterogeneous sources into a common sequence of *hierarchical segments*. After the construction, each segment is a lightweight record  $s = (id, level, parent, content, metadata)$ , where *level* marks granularity (e.g., *document/paragraph/sentence*, *table\_row/table\_cell*, *triplet/subgraph*), *parent* keeps locality via a unique container pointer, *content* is a compact human-readable summary (text span, serialized row, or compact triple), and *metadata* standardizes provenance and alignment with fixed keys (e.g., *source\_id*, *uri*, *offsets*, *schema*, *time*). Segments are concatenated into a final usable  $S_h$  in parent-before-child order. This minimal contract enables structure-aware neighborhoods and budget-aware iteration without inspecting raw files.

**Concrete fragment examples.** Below we instantiate  $(\ell, p, c, \mu)$  for three modalities (all values are illustrative):

```

238 stext = (id = s1,  $\ell$  = sentence,  $p$  = p1, c = "...capital is Paris...",  

239  $\mu$  = {source_id = doc_12, offsets = (128, 172), time = 1992});  

240 srow = (id = r3,  $\ell$  = table_row,  $p$  = t1, c = "France — 67.4M — EU",  

241  $\mu$  = {source_id = tbl_7, schema = [country.pop,bloc], row = 3});  

242 scell = (id = u3,2,  $\ell$  = table_cell,  $p$  = r3, c = "67.4M",  

243  $\mu$  = {source_id = tbl_7, row = 3, col = 2});  

244 skg = (id = k9,  $\ell$  = triplet,  $p$  = g1, c = "(Paris, capital_of, France)",  

245  $\mu$  = {source_id = kg_2}).
```

Segments are concatenated in parent-before-child order. This minimal contract enables structure-aware neighborhoods and budget-aware iteration without inspecting raw files.

### 3.4 HSEQ-ITERATOR(HSEQ-I)

After HSEQ  $S_h$  are build, the HSEQ-Iterator  $\pi_\theta$  can be used on  $(q, S_h)$  and maintains an evolving evidence set  $M_t$  regarding question  $q$ .

**Guidance prior.** A short guidance  $g = g(q, \text{type})$  is treated as a *prior* over iterator actions.  $g$  is generated before each episode to shape exploration on  $S_h$ . This guidance can come from directly from head agent  $\mathcal{H}$ , or from heuristic templates keyed by type.

**Iteration control.** Let  $M_t \subseteq S_h$  denote the selected evidence at step  $t$ ,  $C_t \subseteq S_h$  a candidate window obeying a budget state  $B_t$ , and  $\mathcal{N}(\cdot)$  the structure-aware neighborhood operators induced by levels, parents, and coordinates. The HSEQ-I module  $\pi_\theta$  functions each step following the policy

$$\pi_\theta(a_t, s_t \mid q, S_h, M_t, C_t, g, B_t),$$

which emits an action  $a_t$  (e.g., selecting up to  $k$  segments from  $C_t$  and/or expanding via  $\mathcal{N}$ ) and a sufficiency prediction  $s_t \in \{0, 1\}$ . A deterministic ordering  $\rho$  over  $S_h$  (e.g., paragraph  $\prec$  row  $\prec$  sentence  $\prec$  triplet) defines the stream exposed to the policy. State evolves via a deterministic update  $M_{t+1} = u(M_t, a_t)$  and  $C_{t+1} = \text{window}(S_h, M_{t+1}, B_{t+1}, \rho)$ . Termination occurs at  $\tau = \min\{t : s_t = 1\}$  or when the budget is exhausted.

With set window size  $W$  and step cap  $T_{\max}$ , the algorithm can be described as Alg. 1, where the *Refresh* operator advances the window while removing already selected segments, keeping the per-step context bounded independent of corpus size.

270 **One-step worked example.** Suppose  $\rho$  is paragraph  $\prec$  row  $\prec$  sentence  $\prec$  triplet and  
 271  $W = 5$ . At  $t = 0$ ,  $C_0$  is the first 5 segments under  $\rho$ . The policy selects  $K_1 \subseteq C_0$  ( $|K_1| \leq k$ ) and  
 272 optionally expands with  $\mathcal{N}_{\text{children}}$  (to get sentences within a paragraph) or  $\mathcal{N}_{\text{row}}$  (to fetch a full table  
 273 row when a cell is promising). Then  $M_1 = M_0 \cup K_1$ . Refresh advances the window to the next  
 274 5 *unseen* segments in  $S_h$ . If  $\Phi$  deems evidence sufficient ( $s_1 = 1$ ) and  $t \geq T_{\min}$ , iteration halts;  
 275 otherwise proceed to  $t = 2$  with updated  $B_t$ .

276  
 277 **3.5 HSEQ-HEAD (HSEQ-H).**

278 The HSEQ-Head module  $\mathcal{H}$  can be used in two parts: 1) Guiding the retrieval for HSEQ-I; and 2)  
 279 Generating final conclusion regarding the question.

280 **Guidance Generation.** Although heuristic templates can be used, regarding an incoming ques-  
 281 tion,  $\mathcal{H}$  is available to be called first to analysis the content, generating guidance including: 1) Initial  
 282 Retrieval Plan; 2) What information may be needed; 3) Potential conditions to stop.

283 **Answer synthesis and optional refinement.** Upon termination at step  $\tau$ , the canonicalizer  $\kappa$  con-  
 284 verts  $M_\tau$  into a compact, provenance-preserving evidence package(ids, levels, offsets/coordinates,  
 285 short snippets). The head module  $\mathcal{H}$  then produces the final prediction:

$$\hat{y} = \mathcal{H}(q, \kappa(M_\tau)).$$

286 An optional verifier  $\xi$  inspects  $\kappa(M_\tau)$  for contradictions; if detected, a brief refinement pass (at most  
 287  $\Delta$  additional steps) resumes iteration in Alg. 1 with tightened guidance  $g'$  and reduced budget  $B'$ .

---

288 **Algorithm 1** Guided Iterative Selection under HSEQ-I

289 **Require:** question  $q$ , HSEQ  $S_h$ , guidance  $g$ , budget  $B$ , window size  $W$ , step cap  $T_{\max}$ , minimum  
 290 steps  $T_{\min}$ , top- $k$   $k$ , ordering  $\rho$   
 291 1:  $M_0 \leftarrow \emptyset$ ;  $C_0 \leftarrow \text{Window}(S_h; W, B_0, \rho)$   
 2: **for**  $t = 1$  to  $T_{\max}$  **do**  
 3:   Update  $a_t$   
 4:    $(K_t, s_t) \xleftarrow{a_t} \pi_\theta(q, g, M_{t-1}, C_{t-1}; B_t)$   $\triangleright K_t \subseteq C_{t-1}, |K_t| \leq k$   
 5:    $M_t \leftarrow M_{t-1} \cup K_t$   
 6:    $C_t \leftarrow \text{Refresh}(S_h, M_t; W, \rho)$   
 7:   **if**  $s_t = 1$  **and**  $t \geq T_{\min}$  **then**  
 8:     **break**  
 9:   **end if**  
 10:   Update  $B_t$   
 11: **end for**  
 12:  $\tau \leftarrow t$ ; **return**  $\kappa(M_\tau)$

---

309 **4 LEARNING TO USE HSEQ WITH OPEN-SOURCE LLMs**

310 This section details how we instantiate HSEQ with open-source LLMs, with Section 3 as the theo-  
 311 retical interface and reuse all symbols without redefining them.

312 **4.1 FINE-TUNING HSEQ-I**

313 **Training tuples and supervision.** We organize supervision as tuples  $(q, \text{type}, S_h, A^*)$ . Besides  
 314  $q$  and  $S_h$ , an optional label type is added. The trajectory  $A^* = \{(a_t^*, s_t^*)\}_{t=1}^{\tau^*}$  contains an action  
 315 and a binary sufficiency signal with  $\tau^* = \min\{t : s_t^* = 1\}$ . When explicit trajectories are unavail-  
 316 able, *weak positives*  $P^* \subseteq S_h$  are induced by high-precision matching between gold answers (or  
 317 oracle spans) and segment content, optionally augmented by lexical overlap with  $q$ . A target action  
 318 sequence is synthesized by greedily selecting from  $P^*$  under the budget (details in App. A.2).

324 **Policy learning.** Let the step state be  $(q, S_h, M_t, C_t, g, B_t)$ . We train  $\pi_\theta$  by supervised risk minimization with parameter-efficient adaptation of a base LLM. With teacher forcing for  $t < \tau^*$ ,

$$327 \min_{\theta} \mathbb{E} \left[ \sum_{t=1}^{\tau^*} \underbrace{\ell_{\text{act}}(\pi_\theta(\cdot | \text{state}_t), a_t^*)}_{\text{action loss}} + \lambda \underbrace{\ell_{\text{stop}}(\pi_\theta(\cdot | \text{state}_t), s_t^*)}_{\text{sufficiency loss}} \right],$$

330 where  $\text{state}_t = (q, S_h, M_t, C_t, g, B_t)$  and  $\lambda > 0$  balances early stopping. When  $A^*$  is synthesized from  $P^*$ , per-step weights attenuate low-confidence choices to reduce label noise (App. A.2).  
331 During experiments, Low-Rank Adaptation is used for finetuning (Hu et al., 2022) (App. A.3.4).  
332

## 334 4.2 HSEQ-H: GUIDANCE AND ANSWER GENERATION

335  
336  
337  
338 **Guidance generation (HSEQ-H).** Given a question  $q$  (and optional type), HSEQ-H produces a short  
339 guidance  $g$  that steers the iterator and specifies a  
340 stop rule. We use two modes: (i) a lightweight  
341 planner that drafts  $g$  in 2–4 sentences; (ii) a heuristic  
342 template keyed by coarse question patterns (e.g.,  
343 number/factoid/yes–no).  $g$  follows a fixed structure:  
344 *first-look targets* (entities/rows/1–2-hop neighbors),  
345 *expansion rule* (parent/child, row/column, or relation  
346 hops), and *stop rule* (e.g., “answer span/number  
347 is explicit and corroborated by  $\geq 1$  snippet”).  $g$  is  
348 cached and reused; on a cache miss, we generate or  
349 fall back to the template.  $g$  is a *soft prior*—the iterator  
350 may override it when stronger signals appear.

### 353 Evidence-conditioned answering (HSEQ-H).

354 After  $\kappa$  compacts  $M_\tau$  to  $Z = \kappa(M_\tau)$  (snippets plus  
355 ids/levels/source and minimal offsets/coordinates),  
356 HSEQ-H performs evidence-conditioned answering:  $\hat{y} = \mathcal{H}(q, Z; g)$  using a minimal prompt:  
357 *answer only* (span/number/yes–no) grounded in  $Z$ , no chain-of-thought. When useful, we also  
358 return supporting ids from  $Z$  for auditability. A lightweight entailment-style check over  $Z$  may  
359 trigger a one-shot *refinement*—the iterator resumes for a few steps under a tightened  $g'$ —otherwise  
360  $\hat{y}$  is emitted.

## 361 5 EXPERIMENT

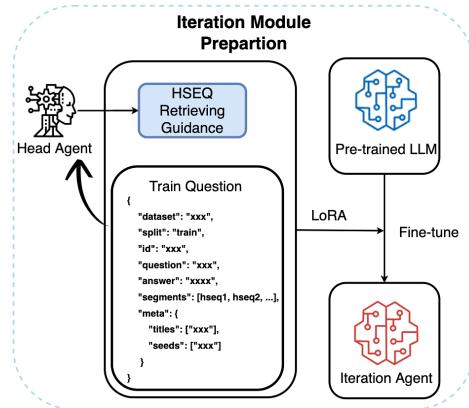
362 HSEQ are evaluated on multiple QA datasets with a focus on both answer quality and efficiency.  
363 Metrics include Accuracy, F1, alongside efficiency indicators that reflect the *evidence actually in-*  
364 *pected*: average iteration steps and end-to-end latency (ms).

### 365 5.1 EXPERIMENT SETUP: BENCHMARKS AND BASELINES

366 **Benchmarks.** To evaluate HSEQ usage from  
367 different data modalities, four benchmarks are used  
368 for experiments, stressing text-only, table–text hybrid,  
369 and KG reasoning: *HotpotQA* (Yang et al.,  
370 2018) (multi-hop reasoning over Wikipedia text),

371 Table 1: Datasets used in our study (modality abbreviations: T=Text,  
372 Tb=Table, KG=Knowledge Graph).

Dataset	Modality	#Train	#Validation	#Test
HotpotQA	T	90447	7405	7405
TAT-QA	Tb + T	13,251	1644	1,663
HybridQA	Tb + T	62,682	3,466	3463
MetaQA-2Hop	KG	119,986	14,872	14,872
MetaQA-3Hop	KG	17,482	14,274	14,274



373 Figure 3: HSEQ-I is trained from multi-  
374 source questions. After guidance sets are  
375 prepared, LoRA is applied for finetuning.

378 *TAT-QA* (Zhu et al., 2021a)  
 379 (table-centric financial  
 380 QA with accompanying paragraphs and numerical operations), *HybridQA* (Chen et al., 2020)  
 381 (Wikipedia tables linked to passages requiring cross-format grounding), and *MetaQA* (Zhang et al.,  
 382 2018) over a Wikidata-style KG (Since 1-hop variants are not emphasized due to triviality, during  
 383 experiment only 2-hop and 3-hop questions are used for experiments).

384  
 385 **Baselines.** Three groups are divided for experiments including:

386

- 387 • **LLM-only QA.** Multiple LLMs are used to directly answers each question from raw inputs  
 388 without HSEQ (no unified adapter, no iterative controller), under same prompt instruction.
- 389 • **RAG-based methods.** Since HSEQ explores different formats of data sources, RAG mod-  
 390 els specializing in separately *Text*, *Table* and *Knowledge Graphs* have been tested.  
 391 Specifically, for HybridQA and TAT-QA, *TAT-LLM* (Zhu et al., 2024), *TableRAG* (Yu  
 392 et al., 2025), *ODYSSEY* (Agarwal et al., 2025), *TTQA-RS* (Bardhan et al., 2024) and  
 393 *HippoRAG* (Jimenez Gutierrez et al., 2024) are chosen for comparison. While for Hot-  
 394 potQA and MetaQA-2hop and 3hop, graph-centric RAG systems *Graph-constrained Rea-*  
 395 *soning(GcR)* (Luo et al., 2024), *Think on Graph (ToG)* (Ma et al., 2024) and *Adap-*  
 396 *tiveRAG* (Jeong et al., 2024) are set as baselines. Each is configured per its recommended  
 397 settings for the corresponding modality.
- 398 • **HSEQ (ours).** (i) The *best* iteration–head pair results and (ii) The *median* pair results over  
 399 a grid of open-source models are provided. Three ablations are also included in experi-  
 400 ments: (i) *LLM-only (no HSEQ)*; (ii) *HSEQ w/o SFT* (iteration agent not fine-tuned) and  
 401 (iii) *heuristic-only guidance under fixed template without HSEQ-H*.

402 **HSEQ variants.** For the *iteration agent* (HSEQ-I) and the *head agent* (HSEQ-H), different LLMs  
 403 are finetuned and used, listed as:

404

405 HSEQ-I: Falcon-H1-3B-Instruct; Qwen3-4B-Instruct-2507; DeepSeek-R1-Distill-Qwen-7B;  
 406 Falcon3-3B-instruct; Falcon3-7B-instruct; Llama-3.2-3B-Instruct.

407 HSEQ-H: Falcon3-10B-instruct; Falcon-H1-7B-Instruct; Llama-3.1-8B-Instruct; DeepSeek-  
 408 R1-Distill-Qwen-7B.

409 Compatible pairs are swept and final “best” and “median” results across benchmarks are counted,  
 410 with hyperparameters settings listed in App. A.3.

## 411 5.2 EXPERIMENT RESULT: HOW COMPETITIVE IS HSEQ WITH OTHER BASELINES?

412 Table 2 summarizes answer quality across all datasets. HSEQ consistently improves over both LLM-  
 413 only and strong RAG baselines, while using controlled iteration and exposing explicit provenance.  
 414 Detailed per-model pairs are reported in Table 3. Efficiency measurements (tokens/latency/steps)  
 415 are in Table 4.

416 Our HSEQ achieves strong and consistent gains on multiple benchmarks. On HotpotQA, MetaQA-  
 417 2hop, and MetaQA-3hop, both the *best* and *median* HSEQ configurations surpass all baselines. On  
 418 TAT-QA, HSEQ’s best run attains the top score overall, while the median run trails slightly behind  
 419 TAT-LLM (Zhu et al., 2024). On the table-and-text HybridQA, HSEQ attains the best accuracy and  
 420 the second-best F1 (just behind HippoRAG (Jimenez Gutierrez et al., 2024)); the median configura-  
 421 tion remains third among baselines.

## 422 5.3 YIELDING BETWEEN EFFICIENCY AND ACCURACY

423 Table 3 lists results using different HSEQ-I and HSEQ-A. The HybridQA results reveal a clear accu-  
 424 racy–efficiency trade-off across HSEQ agent pairs. The highest accuracy/F1 comes from Qwen3-4B  
 425 (HSEQ-I) + Falcon-H1-7B (HSEQ-H) (66.2 / 71.4), with the second-best Qwen3-4B + Llama-3.1-  
 426 8B (65.5 / 71.2). These configurations, however, incur larger iteration depth and latency (about  
 427 3.7–4.1 steps; 16.5–21.5 second). On the efficiency end, Llama-3.2-3B + Llama-3.1-8B delivers  
 428 the lowest steps and latency (2.11; 8.35k ms) with moderate accuracy (55.4 / 57.9), while Falcon3-  
 429 3B + Falcon-H1-7B attains the second-best efficiency (2.25; 11.7k ms) at similar quality. Taken

432 Table 2: Overall QA performance on heterogeneous benchmarks. Shaded cells (N/A) indicate the  
 433 method is not applicable to that benchmark; gray dashes (–) indicate metric not reported. The record  
 434 results use Qwen3-4B-Instruct-2507 for HSEQ-I; and Falcon-H1-7B-Instruct for HSEQ-H

Method	HybridQA		TAT-QA		HotpotQA		MetaQA-2hop		MetaQA-3hop	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
<b>LLM-only (direct QA)</b>										
Falcon3-10B-instruct	22.4	–	35.2	–	16.5	–	43.0	–	39.8	–
Falcon-H1-7B-Instruct	32.9	–	43.7	–	21.1	–	48.3	–	44.6	–
Llama-3.1-8B-Instruct	28.1	–	37.6	–	14.6	–	37.8	–	31.9	–
Qwen3-4B-Instruct-2507	30.3	–	42.1	–	17.8	–	42.2	–	38.5	–
<b>RAG-based methods (single-pass / agentic baselines)</b>										
TAT-LLM	–	–	73.1	81.0	N/A	N/A	N/A	N/A	N/A	N/A
TableRAG	47.9	–	61.9	68.6	N/A	N/A	N/A	N/A	N/A	N/A
ODYSSEY	51.5	66.0	–	–	N/A	N/A	N/A	N/A	N/A	N/A
TTQA-RS	62.3	70.6	–	–	N/A	N/A	N/A	N/A	N/A	N/A
HippoRAG	65.8	72.4	70.1	74.9	53.2	55.7	N/A	N/A	N/A	N/A
Graph-constrained Reasoning (GcR)	N/A	N/A	N/A	N/A	39.2	41.6	86.7	88.1	83.2	80.6
Think on Graph (ToG)	N/A	N/A	N/A	N/A	43.1	44.7	83.2	84.8	81.1	78.5
AdaptiveRAG	N/A	N/A	N/A	N/A	50.3	52.5	88.2	90.1	84.5	85.7
<b>Our method: HSEQ</b>										
HSEQ (best)	<b>66.4</b>	72.1	<b>75.7</b>	<b>83.5</b>	<b>56.3</b>	<b>58.6</b>	<b>95.9</b>	<b>91.1</b>	<b>93.4</b>	<b>88.3</b>
HSEQ (median)	63.9	70.8	73.2	79.6	55.4	57.1	93.2	89.7	90.1	86.6

450  
 451 together, the Pareto frontier spans (i) Qwen-based iterators with larger heads for top accuracy, and  
 452 (ii) lightweight Llama/Falcon pairs for predictable low latency. Different agent pairs can be chosen  
 453 regarding whether accuracy or budget dominates.

454  
 455 Table 3: Overall performance of HSEQ agent pairs on Hybrid-QA: Accuracy/F1 and Efficiency.

Iteration Agent (HSEQ-I)	Head Agent (HSEQ-H)	Accuracy & F1		Efficiency	
		Avg. Acc	Avg. F1	Steps ↓	Latency (ms) ↓
Llama-3.2-3B-Instruct	Falcon3-10B-instruct	60.4	62.8	2.08	12055.5
Qwen3-4B-Instruct-2507	Falcon3-10B-instruct	63.9	64.5	4.1	20577.5
Falcon3-3B-instruct	Falcon3-10B-instruct	59.3	61.1	2.6	10530.1
Llama-3.2-3B-Instruct	Llama-3.1-8B-Instruct	55.4	57.9	<b>2.11</b>	<b>8346.3</b>
Qwen3-4B-Instruct-2507	Llama-3.1-8B-Instruct	65.5	71.2	3.29	16503.2
Falcon3-3B-instruct	Llama-3.1-8B-Instruct	61.2	65.1	2.46	11616.7
Llama-3.2-3B-Instruct	Falcon-H1-7B-Instruct	58.7	63.9	2.41	12080.0
Qwen3-4B-Instruct-2507	Falcon-H1-7B-Instruct	<b>66.2</b>	<b>71.4</b>	3.71	21479.2
Falcon3-3B-instruct	Falcon-H1-7B-Instruct	56.1	58.6	<b>2.25</b>	<b>11714.4</b>
Llama-3.2-3B-Instruct	DeepSeek-R1-Distill-Qwen-7B	62.5	60.2	2.75	15073.7
Qwen3-4B-Instruct-2507	DeepSeek-R1-Distill-Qwen-7B	62.8	66.7	4.07	21094.8
Falcon3-3B-instruct	DeepSeek-R1-Distill-Qwen-7B	61.4	62.0	3.01	13709.7

469  
 470 5.4 EFFICIENCY ANALYSIS  
 471

472 To test HSEQ framework’s latency, *evidence actually inspected* are calculated: iteration steps for  
 473 HSEQ-I and wall-clock latency are calculated. Results are summarized below. “LLM-only” incurs  
 474 a single forward pass (1 step) and thus the lowest raw latency, but this comes at the cost of weaker  
 475 multi-hop accuracy and no explicit provenance in Table 3. In contrast, graph-centric ToG performs  
 476 many expansion steps (11–17 on average), which substantially increases latency (e.g., over 22k ms  
 477 on HotpotQA and 24k ms on MetaQA-3hop), even though it is designed for multi-hop reasoning.

478  
 479 Table 4: Efficiency metrics on HotpotQA, MetaQA-2hop and MetaQA-3hop.

Efficiency		HotpotQA		MetaQA-2hop		MetaQA-3hop	
Method	Steps	Latency (ms) ↓	Steps	Latency (ms) ↓	Steps	Latency (ms) ↓	
LLM-only	1	3266.3	1	2556.4	1	3631.1	
Think on Graph (ToG)	13.28	22708.2	11.73	15707.6	16.58	24307.4	
HSEQ (ours, best)	4.00	6247.0	3.27	5732.2	4.11	10632.8	
HSEQ (ours, median)	4.17	12114.4	3.76	9480.1	4.59	13505.3	

HSEQ occupies a middle ground in this trade-off. Both the best and median HSEQ variants maintain short, budgeted loops of roughly 3–5 steps across datasets, yet reduce latency by more than half relative to ToG on all three benchmarks. This indicates that guided, windowed iteration over HSEQ can retain multi-hop capability while avoiding the long expansion chains and repeated graph traversals of ToG. Compared with LLM-only, HSEQ pays a moderate overhead in latency but gains structured evidence and substantially higher accuracy on multi-step questions. HSEQ provides a more balanced operating point with bounded steps and competitive performance.

## 5.5 ABLATION STUDIES

Ablation studies are set to evaluate each component of HSEQ framework on representative text (HotpotQA) and table-text (HybridQA) tasks. Following tasks are considered: (a) **No SFT** (iteration agent not fine-tuned); (b) **No guidance** (remove  $g$ ); (c) **Heuristic-only guidance** (no planner) ; and (d) **LLM-only** (without multi-agent but use HSEQ as part of prompt for data input).

Table 5: Ablations on benchmarks.

Variant	HybridQA		TAT-QA		HotpotQA		MetaQA-3hop		MetaQA-3hop	
	Acc	F1								
HSEQ (full)	<b>66.4</b>	<b>72.1</b>	<b>75.7</b>	<b>83.5</b>	<b>56.3</b>	<b>58.6</b>	<b>95.9</b>	<b>91.1</b>	<b>93.4</b>	<b>88.3</b>
w/o SFT (base iteration)	57.3	65.7	60.4	66.9	46.5	47.8	78.3	80.1	74.6	72.5
w/o guidance	59.2	62.6	68.8	75.1	50.5	51.2	82.4	83.0	79.2	73.8
heuristic-only guidance	<u>63.8</u>	<u>67.3</u>	<u>70.4</u>	<u>79.9</u>	<u>54.7</u>	<u>56.1</u>	<u>87.3</u>	<u>85.4</u>	<u>83.9</u>	<u>86.1</u>
LLM-only (no HSEQ)	32.9	—	43.7	—	21.1	—	48.3	—	44.6	—

The ablation study demonstrates the necessities of all HSEQ’s components, with differing sensitivity across formats. Using *heuristic-only* guidance yields the smallest degradation from the full system—typically a modest drop in Acc/F1—indicating that a lightweight, template-style prior already guides HSEQ-I effectively when the planner is absent. Removing fine-tuning (*w/o SFT*) causes a larger decline, but with the use of structured HSEQ data, accuracy remains substantially higher than *LLM-only*. Without guidance (*w/o guidance*) influence performance, as in prompt HSEQ-I is only asked to *choose necessary evidence from below to answer the question*. The results underscore the role of guidance as a portable sufficiency prior. Finally, the *LLM-only* setting performs worst across all benchmarks, reflecting the difficulty of recovering minimally sufficient evidence without iterative, structure-aware selection. Overall, the results suggest that (i) HSEQ’s unified data structure is the primary source of robustness, (ii) SFT HSEQ-I provides consistent gains, and (iii) guidance—even a simple heuristic ones from template-would increase overall accuracy strongly.

## 6 CONCLUSION

This paper introduces **HSEQ**, a compact framework for heterogeneous QA that (i) *unifies* text, tables, and knowledge graphs into a reversible hierarchical sequence with lightweight structure and provenance; (ii) performs *guided, budget-aware iteration* that selects small sets of salient segments and predicts *sufficiency* for early stopping; and (iii) feeds a *canonicalized evidence* package to a head module for answer synthesis. By replacing single-shot retrieval and unconstrained agentic loops with short, structure-aware selections equipped with an explicit sufficiency signal, HSEQ concentrates computation on *evidence actually inspected*, delivers predictable latency under token/tool budgets, and preserves auditability through provenance-aware canonicalization.

Across heterogeneous QA benchmarks, HSEQ achieves strong answer quality alongside consistent efficiency, revealing a controllable trade-off between accuracy and cost: larger head with finetuned small iterators achieved both fast and accurate QA. The format-agnostic interface and standardized action schema enable a single learned policy to operate across modalities without per-dataset retrievers, bespoke prompts, or tokenizer changes. *Future work* will extend HSEQ to multi-turn/streaming settings with dynamic corpora, mitigate hallucination on sufficiency judge under noisy evidence.

540 7 ETHICS STATEMENT.  
541542 We affirm adherence to the ICLR Code of Ethics. All experiments use publicly available benchmarks  
543 (HybridQA, TatQA, HotpotQA, MetaQA) under their respective licenses; no new human-subject  
544 data were collected, and no personally identifiable information (PII) is processed. Our HSEQ  
545 construction preserves provenance via identifiers and offsets while avoiding storage of copyrighted text  
546 beyond short snippets necessary for QA. As with any LLM-based system, model outputs may re-  
547 flect societal biases inherited from pretraining corpora; we mitigate this risk by requiring explicit,  
548 auditable evidence and by permitting abstention when sufficiency is not met. We release code and  
549 configuration solely for research use and discourage deployment in high-stakes settings without  
550 domain-specific evaluation and additional safeguards (fairness, privacy, and safety audits).551 8 REPRODUCIBILITY STATEMENT.  
552553 We provide an anonymous GitHub link (<https://anonymous.4open.science/r/HSEQ-anonymous-0DAC>) with code and scripts to (i) construct HSEQ from raw corpora, (ii) fine-tune the iteration  
554 policy with LoRA, and (iii) run guided inference and evaluation. Implement details are shown  
555 in App. A.3, containing models used (App. A.3.1), prompts (App. A.3.2- A.3.3), LoRA adaption  
556 parameters (App. A.3.4) and reproducibility notes (App. A.3.6). Theorems include complete as-  
557 sumptions and proofs (App. A.1). Apart from the code, detailed examples of agents interactions  
558 (example questions, LLM outputs, data retrieved each steps, etc.) are provided in App. A.5 and as a  
559 jsonl file in our anonymous repository.  
560561 REFERENCES  
562563 

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-  
564 man, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical  
565 report. *arXiv preprint arXiv:2303.08774*, 2023.

Ankush Agarwal, Chaitanya Devaguptapu, et al. Hybrid graphs for table-and-text based question  
566 answering using llms. *arXiv preprint arXiv:2501.17767*, 2025.

Jayetri Bardhan, Bushi Xiao, and Daisy Zhe Wang. Ttqa-rs-a break-down prompting approach  
567 for multi-hop table-text question answering with reasoning and summarization. *arXiv preprint  
arXiv:2406.14732*, 2024.

Chi-Min Chan, Chunpu Xu, Ruibin Yuan, Hongyin Luo, Wei Xue, Yike Guo, and Jie Fu. Rq-rag:  
568 Learning to refine queries for retrieval augmented generation. *arXiv preprint arXiv:2404.00610*,  
569 2024.

Wenhu Chen, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, and William Wang. Hy-  
570 bridqa: A dataset of multi-hop question answering over tabular and textual data. *arXiv preprint  
arXiv:2004.07347*, 2020.

Xinyue Chen, Pengyu Gao, Jiangjiang Song, and Xiaoyang Tan. Hiqa: A hierarchical contextual  
571 augmentation rag for multi-documents qa. *arXiv preprint arXiv:2402.01767*, 2024.

Yiqun Chen, Lingyong Yan, Weiwei Sun, Xinyu Ma, Yi Zhang, Shuaiqiang Wang, Dawei Yin,  
572 Yiming Yang, and Jiaxin Mao. Improving retrieval-augmented generation through multi-agent  
573 reinforcement learning. *arXiv preprint arXiv:2501.15228*, 2025.

Philipp Christmann and Gerhard Weikum. Rag-based question answering over heterogeneous data  
574 and text. *arXiv preprint arXiv:2412.07420*, 2024.

Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Gpt3. int8 (): 8-bit matrix  
575 multiplication for transformers at scale. *Advances in neural information processing systems*, 35:  
576 30318–30332, 2022.

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning  
577 of quantized llms. *Advances in neural information processing systems*, 36:10088–10115, 2023.

594 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha  
 595 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.  
 596 *arXiv e-prints*, pp. arXiv–2407, 2024.

597 Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt,  
 598 Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. From local to global: A  
 599 graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*, 2024.

600 Wenqi Fan, Yujuan Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and  
 601 Qing Li. A survey on rag meeting llms: Towards retrieval-augmented large language models. In  
 602 *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*,  
 603 pp. 6491–6501, 2024.

604 Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and  
 605 Haofen Wang. Retrieval-augmented generation for large language models: A survey. *arXiv  
 606 preprint arXiv:2312.10997*, 2023.

607 Michael Glass, Gaetano Rossiello, Md Faisal Mahbub Chowdhury, Ankita Rajaram Naik, Pengshan  
 608 Cai, and Alfio Gliozzo. Re2g: Retrieve, rerank, generate. *arXiv preprint arXiv:2207.06300*, 2022.

609 Sagar Goyal, Eti Rastogi, Sree Prasanna Rajagopal, Dong Yuan, Fen Zhao, Jai Chintagunta, Gautam  
 610 Naik, and Jeff Ward. Healai: A healthcare llm for effective medical documentation. In *Proceed-  
 611 ings of the 17th ACM International Conference on Web Search and Data Mining*, pp. 1167–1168,  
 612 2024.

613 Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V Chawla, Olaf Wiest,  
 614 and Xiangliang Zhang. Large language model based multi-agents: A survey of progress and  
 615 challenges. *arXiv preprint arXiv:2402.01680*, 2024a.

616 Yue Guo and Yi Yang. Econnli: evaluating large language models on economics reasoning. *arXiv  
 617 preprint arXiv:2407.01212*, 2024.

618 Zirui Guo, Lianghao Xia, Yanhua Yu, Tu Ao, and Chao Huang. Lightrag: Simple and fast retrieval-  
 619 augmented generation. *arXiv preprint arXiv:2410.05779*, 2024b.

620 Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Ceyao Zhang, Jinlin  
 621 Wang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, et al. Metagpt: Meta programming for  
 622 a multi-agent collaborative framework. International Conference on Learning Representations,  
 623 ICLR, 2024.

624 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,  
 625 Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

626 Yuntong Hu, Zhihan Lei, Zheng Zhang, Bo Pan, Chen Ling, and Liang Zhao. Grag: Graph retrieval-  
 627 augmented generation. *arXiv preprint arXiv:2405.16506*, 2024.

628 Yiqian Huang, Shiqi Zhang, and Xiaokui Xiao. Ket-rag: A cost-efficient multi-granular indexing  
 629 framework for graph-rag. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge  
 630 Discovery and Data Mining* V. 2, pp. 1003–1012, 2025.

631 Md Ashraful Islam, Mohammed Eunus Ali, and Md Rizwan Parvez. Mapcoder: Multi-agent code  
 632 generation for competitive problem solving. *arXiv preprint arXiv:2405.11403*, 2024.

633 Soyeong Jeong, Jinheon Baek, Sukmin Cho, Sung Ju Hwang, and Jong C Park. Adaptive-rag:  
 634 Learning to adapt retrieval-augmented large language models through question complexity. *arXiv  
 635 preprint arXiv:2403.14403*, 2024.

636 Zhengbao Jiang, Frank F Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang,  
 637 Jamie Callan, and Graham Neubig. Active retrieval augmented generation. In *Proceedings of the  
 638 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 7969–7992, 2023.

639 Bernal Jimenez Gutierrez, Yiheng Shu, Yu Gu, Michihiro Yasunaga, and Yu Su. Hipporag: Neurobi-  
 640 logically inspired long-term memory for large language models. *Advances in Neural Information  
 641 Processing Systems*, 37:59532–59569, 2024.

648 Jiajie Jin, Yutao Zhu, Zhicheng Dou, Guanting Dong, Xinyu Yang, Chenghao Zhang, Tong Zhao,  
 649 Zhao Yang, and Ji-Rong Wen. Flashrag: A modular toolkit for efficient retrieval-augmented  
 650 generation research. In *Companion Proceedings of the ACM on Web Conference 2025*, pp. 737–  
 651 740, 2025.

652 Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal,  
 653 Heinrich Kütller, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented gener-  
 654 ation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:  
 655 9459–9474, 2020.

656 Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv  
 657 preprint arXiv:2101.00190*, 2021.

659 Zechen Li, Shohreh Deldari, Linyao Chen, Hao Xue, and Flora D Salim. Sensorllm: Aligning large  
 660 language models with motion sensors for human activity recognition. 2024.

662 Pei Liu, Xin Liu, Ruoyu Yao, Junming Liu, Siyuan Meng, Ding Wang, and Jun Ma. Hm-rag: Hierar-  
 663 chical multi-agent multimodal retrieval augmented generation. *arXiv preprint arXiv:2504.12330*,  
 664 2025.

665 Haoran Luo, Zichen Tang, Shiyao Peng, Yikai Guo, Wentai Zhang, Chenghao Ma, Guanting  
 666 Dong, Meina Song, Wei Lin, Yifan Zhu, et al. Chatkbqa: A generate-then-retrieve framework  
 667 for knowledge base question answering with fine-tuned large language models. *arXiv preprint  
 668 arXiv:2310.08975*, 2023.

669 Linhao Luo, Zicheng Zhao, Gholamreza Haffari, Yuan-Fang Li, Chen Gong, and Shirui Pan. Graph-  
 670 constrained reasoning: Faithful reasoning on knowledge graphs with large language models.  
 671 *arXiv preprint arXiv:2410.13080*, 2024.

673 Linhao Luo, Zicheng Zhao, Gholamreza Haffari, Dinh Phung, Chen Gong, and Shirui Pan. Gfm-rag:  
 674 graph foundation model for retrieval augmented generation. *arXiv preprint arXiv:2502.01113*,  
 675 2025.

676 Shengjie Ma, Chengjin Xu, Xuhui Jiang, Muzhi Li, Huaren Qu, Cehao Yang, Jiaxin Mao, and Jian  
 677 Guo. Think-on-graph 2.0: Deep and faithful large language model reasoning with knowledge-  
 678 guided retrieval augmented generation. *arXiv preprint arXiv:2407.10805*, 2024.

680 Costas Mavromatis and George Karypis. Gnn-rag: Graph neural retrieval for large language model  
 681 reasoning. *arXiv preprint arXiv:2405.20139*, 2024.

682 Nathalia Nascimento, Paulo Alencar, and Donald Cowan. Self-adaptive large language model (llm)-  
 683 based multiagent systems. In *2023 IEEE International Conference on Autonomic Computing and  
 684 Self-Organizing Systems Companion (ACSOS-C)*, pp. 104–109. IEEE, 2023.

686 Boci Peng, Yun Zhu, Yongchao Liu, Xiaohe Bo, Haizhou Shi, Chuntao Hong, Yan Zhang, and  
 687 Siliang Tang. Graph retrieval-augmented generation: A survey. *arXiv preprint arXiv:2408.08921*,  
 688 2024.

689 Haritz Puerto, Gözde Gülgahin, and Iryna Gurevych. Metaqa: Combining expert agents for multi-  
 690 skill question answering. *arXiv preprint arXiv:2112.01922*, 2021.

691 Andrew Reynolds and Felix Corrigan. Improving real-time knowledge retrieval in large language  
 692 models with a dns-style hierarchical query rag. *Authorea Preprints*, 2024.

694 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:  
 695 Language agents with verbal reinforcement learning. *Advances in Neural Information Processing  
 696 Systems*, 36:8634–8652, 2023.

697 Aditi Singh, Abul Ehtesham, Saket Kumar, and Tala Talaei Khoei. Agentic retrieval-augmented  
 698 generation: A survey on agentic rag. *arXiv preprint arXiv:2501.09136*, 2025.

700 Khanh-Tung Tran, Dung Dao, Minh-Duong Nguyen, Quoc-Viet Pham, Barry O’Sullivan, and  
 701 Hoang D Nguyen. Multi-agent collaboration mechanisms: A survey of llms. *arXiv preprint  
 702 arXiv:2501.06322*, 2025.

702 Jianrui Wang, Yitian Hong, Jiali Wang, Jiapeng Xu, Yang Tang, Qing-Long Han, and Jürgen Kurths.  
 703 Cooperative and competitive multi-agent systems: From optimization to games. *IEEE/CAA Journal  
 704 of Automatica Sinica*, 9(5):763–783, 2022.

705 Junde Wu, Jiayuan Zhu, Yunli Qi, Jingkun Chen, Min Xu, Filippo Menolascina, and Vicente Grau.  
 706 Medical graph rag: Towards safe medical large language model via graph retrieval-augmented  
 707 generation. *arXiv preprint arXiv:2408.04187*, 2024.

708 Ruiyi Yang, Hao Xue, Imran Razzak, Hakim Hacid, and Flora D Salim. Beyond single pass, looping  
 709 through time: Kg-irag with iterative knowledge retrieval. *arXiv preprint arXiv:2503.14234*, 2025.

710 Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov,  
 711 and Christopher D Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question  
 712 answering. *arXiv preprint arXiv:1809.09600*, 2018.

713 Hao Yu, Aoran Gan, Kai Zhang, Shiwei Tong, Qi Liu, and Zhaofeng Liu. Evaluation of retrieval-  
 714 augmented generation: A survey. *arXiv preprint arXiv:2405.07437*, 2024.

715 Wenhao Yu. Retrieval-augmented generation across heterogeneous knowledge. In *Proceedings of  
 716 the 2022 conference of the North American chapter of the association for computational linguistics:  
 717 human language technologies: student research workshop*, pp. 52–58, 2022.

718 Xiaohan Yu, Pu Jian, and Chong Chen. Tablerag: A retrieval augmented generation framework for  
 719 heterogeneous document reasoning. *arXiv preprint arXiv:2506.10380*, 2025.

720 Yuyu Zhang, Hanjun Dai, Zornitsa Kozareva, Alexander Smola, and Le Song. Variational reasoning  
 721 for question answering with knowledge graph. In *Proceedings of the AAAI conference on artificial  
 722 intelligence*, volume 32, 2018.

723 Penghao Zhao, Hailin Zhang, Qinhan Yu, Zhengren Wang, Yunteng Geng, Fangcheng Fu, Ling  
 724 Yang, Wentao Zhang, and Bin Cui. Retrieval-augmented generation for ai-generated content: A  
 725 survey. *arXiv preprint arXiv:2402.19473*, 2024.

726 Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng,  
 727 and Tat-Seng Chua. Tat-qa: A question answering benchmark on a hybrid of tabular and textual  
 728 content in finance. *arXiv preprint arXiv:2105.07624*, 2021a.

729 Fengbin Zhu, Wenqiang Lei, Chao Wang, Jianming Zheng, Soujanya Poria, and Tat-Seng Chua.  
 730 Retrieving and reading: A comprehensive survey on open-domain question answering. *arXiv  
 731 preprint arXiv:2101.00774*, 2021b.

732 Fengbin Zhu, Ziyang Liu, Fuli Feng, Chao Wang, Moxin Li, and Tat Seng Chua. Tat-llm: A  
 733 specialized language model for discrete reasoning over financial tabular and textual data. In  
 734 *Proceedings of the 5th ACM International Conference on AI in Finance*, pp. 310–318, 2024.

735 Jingwei Zuo, Maksim Velikanov, Ilyas Chahed, Younes Belkada, Dhia Eddine Rhayem, Guillaume  
 736 Kunsch, Hakim Hacid, Hamza Yous, Brahim Farhat, Ibrahim Khadraoui, et al. Falcon-h1: A  
 737 family of hybrid-head language models redefining efficiency and performance. *arXiv preprint  
 738 arXiv:2507.22448*, 2025.

746  
 747 **A APPENDIX**

748  
 749 **A.1 THEORETICAL PROPERTIES OF HSEQ**

750  
 751 **A.1.1 PRELIMINARIES AND ASSUMPTIONS**

752  
 753 **Segment schema.** An HSEQ is a finite multiset  $S_h$  of segments  $s = (\text{id}(s), \ell(s), p(s), c(s), \mu(s))$ .  
 754 Here  $\ell$  is a level tag;  $p$  is a parent pointer with  $p(s) = \perp$  if  $s$  is a root;  $c$  is content;  $\mu$  is metadata,  
 755 possibly including offsets (for text), schema and row indices (for tables), and triplet fields (for  
 KGs).

756 **Encoder/decoder.** Let  $\Phi$  map any finite corpus  $X$  (text + tables + KG) to  $S_h = \Phi(X)$ , and let  $\Psi$   
 757 map  $S_h$  back to a corpus  $\Psi(S_h)$ . We assume the following modality-specific invariants are enforced  
 758 by the adapters (they match the implementation but are stated abstractly).  
 759

760 **(T1) Text offsets.** For each text item  $x \in \Sigma^*$ , if  $s$  is a paragraph (resp. sentence) segment for a  
 761 span  $x[a : b]$  (resp.  $x[u : v]$ ) inside a paragraph, then  $\mu(s).\text{offsets} = [a, b]$  (resp.  
 762  $[a + u, a + v]$ ),  $c(s) = x[a : b]$  (resp.  $x[a + u : a + v]$ ), and  $p$  is the unique parent in the  
 763 containment chain (sentence  $\rightarrow$  paragraph  $\rightarrow$  document).

764 **(T2) Table rows.** For a table with header  $H = (h_1, \dots, h_C)$  and  $n$  rows  $(r_i)_{i=1}^n$ , the table-root  
 765 segment stores  $H$  in  $\mu(\cdot).\text{schema}$ ; each row-segment  $s_i$  stores  $c(s_i) = \text{dict}(H \mapsto r_i)$   
 766 and either (a) an explicit row index  $\mu(s_i).\text{offsets} = [i, -1]$ , or (b) a total order on row  
 767 segments consistent with the original row order.

768 **(T3) KG triples.** For a KG edge multiset  $E \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$  (optionally time-stamped), each  
 769 edge  $(h, r, t, \tau)$  corresponds to exactly one triplet segment  $s$  with  $c(s) = (h, r, t)$  and  
 770  $\mu(s).\text{time} = \tau$ ; parent  $p(s)$  is the unique subgraph-root for the neighborhood.  
 771

772 **Benign equivalence.** Define an equivalence relation  $\equiv$  over corpora by (i) ignoring differences in  
 773 text whitespace that do not change the sequence of non-whitespace characters; (ii) allowing a global  
 774 column permutation  $\pi \in S_C$  applied uniformly to the header and all row dictionaries of a table; (iii)  
 775 treating KGs as edge multisets (edge order immaterial).

776 **Ordering and window.** Let  $\rho$  be a total order over  $S_h$  (e.g., paragraph  $\prec$  row  $\prec$  sentence  $\prec$  triplet  
 777 with a deterministic tie-break). The stream induced by  $\rho$  lists  $S_h$  as  $(s_1, \dots, s_N)$ . For a window size  
 778  $W \in \mathbb{N}$ ,  $\text{Window}(S_h; W, \rho)$  returns the first  $W$  items of the stream that are not already selected;  
 779  $\text{Refresh}(S_h, M; W, \rho)$  returns the next  $W$  unseen items after removing  $M$ . Both are *monotone* w.r.t.  
 780  $\rho$ : the sequence of items exposed across refreshes is exactly the  $\rho$ -stream with already selected items  
 781 removed.  
 782

783 **Admissibility.** For a question  $q$ , a supporting set  $E^* \subseteq S_h$  is *answer-supporting* if the head module  
 784  $\mathcal{H}$  yields the correct answer when given only  $E^*$ . An order  $\rho$  is *admissible* for  $(q, S_h)$  if there exists  
 785 a minimal  $L \in \{1, \dots, |S_h|\}$  such that  $E^* \subseteq \{s_1, \dots, s_L\}$  for some answer-supporting  $E^*$ .  
 786

787 **Sufficiency predicate.** Let  $\text{Suff}(M)$  be a predicate that holds iff  $M$  contains some answer-  
 788 supporting subset. We assume a calibrated sufficiency head: whenever  $\text{Suff}(M_t)$  becomes true,  
 789 the policy can set its stop flag  $s_t = 1$  at that step or earlier.<sup>1</sup>  
 790

### 791 A.1.2 FAITHFUL LINEARIZATION

793 **Theorem 1** (Faithful linearization). *For any finite corpus  $X$ , under (T1)–(T3), the encoder  $\Phi$  is  
 794 injective up to  $\equiv$ , i.e.,  $\Psi(\Phi(X)) \equiv X$ .*

796 *Proof.* Write  $X = X_{\text{text}} \uplus X_{\text{tbl}} \uplus X_{\text{kg}}$  and let  $S_h = \Phi(X)$ . We show  $\Psi(\Phi(\cdot))$  acts as identity  
 797 modulo  $\equiv$  on each modality and hence on their disjoint union.

798 *Text.* Consider  $x \in X_{\text{text}}$ . By (T1) each paragraph (resp. sentence) segment  $s$  stores the exact  
 799 substring  $c(s) = x[a : b]$  (resp.  $x[u' : v']$ ) and absolute offsets in  $\mu(s).\text{offsets}$ . Let  $S_x \subseteq S_h$   
 800 be all segments rooted at the document node of  $x$ . The decoder reconstructs  $x'$  by placing every  
 801 paragraph substring at its  $[a, b]$  range and merging overlaps implied by sentence children; uniqueness  
 802 of parents eliminates ambiguity. Because offsets are absolute and children are contained in parents  
 803 by construction, the reconstructed  $x'$  equals  $x$  character-for-character; any whitespace normalization  
 804 is permitted by  $\equiv$ .

805 *Tables.* Let a table have header  $H = (h_1, \dots, h_C)$  and rows  $(r_i)_{i=1}^n$ . By (T2),  $\mu(\cdot).\text{schema}$  stores  
 806  $H$ , and each row segment  $s_i$  stores the dictionary  $c(s_i)$  mapping  $H$  to the row tuple  $r_i$ , together  
 807 with either an explicit row index or a total order consistent with the original order. The decoder

808 <sup>1</sup>This is standard in supervised setups where the stop head is trained to fire at first sufficiency (or with  
 809 tolerance).

reassembles the matrix  $[H; r_1; \dots; r_n]$ . Any global column permutation  $\pi$  yields an equivalent table under  $\equiv$ ; thus the reconstruction is unique modulo schema-order permutations.

*KGs.* Let  $E$  be the multiset of edges. By (T3), each edge  $(h, r, t, \tau)$  corresponds bijectively to one triplet segment with  $c(s) = (h, r, t)$  and  $\mu(s).\text{time} = \tau$ , and parentage is irrelevant to content. The decoder collects the multiset of triplets, which equals  $E$ ; edge order is immaterial and thus fits  $\equiv$ .

Since the three reconstructions are independent and disjointly supported,  $\Psi(\Phi(X)) \equiv X$  follows.  $\square$

### 819 A.1.3 WINDOWED ITERATION: COVERAGE AND COMPLEXITY

821 Let  $E^* \subseteq \{s_1, \dots, s_L\}$  be an answer-supporting set with minimal prefix length  $L$  under an admissible order  $\rho$ . Fix window  $W \geq k \geq 1$  and define the iterative selection with refresh as in the main text.

824 **Lemma 1** (Prefix coverage under  $k$ -selection). *After  $t$  steps, the selected set  $M_t$  contains at least  $\min\{kt, L\}$  items from the  $\rho$ -prefix  $\{s_1, \dots, s_L\}$ . In particular,  $E^* \subseteq M_T$  for  $T = \lceil L/k \rceil$ .*

827 *Proof.* We prove by induction on  $t \geq 0$  that  $|M_t \cap \{s_1, \dots, s_L\}| \geq \min\{kt, L\}$ .

828 Base  $t = 0$ :  $M_0 = \emptyset$  so the bound is 0.

830 Inductive step: assume the claim for  $t - 1$ . At step  $t$ , the window exposes (by monotonicity of 831 Refresh) the earliest  $W$  unseen items under  $\rho$ ; hence at least the next  $k$  unseen items in the 832 prefix  $\{s_1, \dots, s_L\}$  are eligible (because  $W \geq k$ ). Selecting  $k$  new items (or fewer if fewer remain in the 833 prefix) increases the count by at least  $\min\{k, L - (t - 1)k\}$ , giving  $\min\{kt, L\}$ . Once all  $L$  prefix 834 items are selected, the bound saturates at  $L$ .  $\square$

835 **Proposition 1** (Guaranteed halt). *Assume a step cap  $T_{\max}$  and a sufficiency head that can set 836  $s_t = 1$  whenever  $\text{Suff}(M_t)$  holds. Under admissibility, the control loop halts after at most 837  $\min\{T_{\max}, \lceil L/k \rceil\}$  steps.*

839 *Proof.* By Lemma 1, after  $T = \lceil L/k \rceil$  steps,  $E^* \subseteq M_T$ ; hence  $\text{Suff}(M_T)$  holds and the stop 840 head can fire at or before  $T$ . Independently, the hard cap  $T_{\max}$  forces termination by  $T_{\max}$  steps. 841 Therefore  $\tau \leq \min\{T_{\max}, T\}$ .  $\square$

843 **Theorem 2** (Budgeted selection complexity). *Let  $C(W) > 0$  be the (deterministic) per-step context 844 cost determined by window size  $W$ . Under admissibility, the total selection cost is bounded by 845*

$$\text{Cost}_{\text{select}} \leq C(W) \cdot \min\{T_{\max}, \lceil L/k \rceil\},$$

847 independent of  $|S_h|$ . If  $L$  is a nonnegative integer random variable with  $\mathbb{E}[L] = \bar{L} < \infty$ , then 848

$$\mathbb{E}[\text{Cost}_{\text{select}}] \leq C(W) \cdot \mathbb{E}[\min\{T_{\max}, \lceil L/k \rceil\}] \leq C(W) \cdot \min\{T_{\max}, \bar{L}/k + 1\}.$$

851 *Proof.* The first bound follows by multiplying the per-step cost by the halt bound in Proposition 1. 852 For the expectation, use linearity of expectation and the inequality  $\lceil x \rceil \leq x + 1$  for  $x \geq 0$ : 853  $\mathbb{E}[\lceil L/k \rceil] \leq \mathbb{E}[L]/k + 1 = \bar{L}/k + 1$ , and  $\mathbb{E}[\min\{a, X\}] \leq \min\{a, \mathbb{E}[X]\}$  for  $a \geq 0$  and  $X \geq 0$ .  $\square$

### 855 A.2 WEAK-POSITIVE LABELING AND TRAJECTORY SYNTHESIS

857 **Positive identification.** For each instance, segments are sorted by a *level priority* that favors 858 container-like units (e.g., paragraphs, rows). Within a capped candidate set, a positive pool  $P^*$  is 859 constructed by: (i) exact/substring matching of the gold answer in `content`; and (ii) if insufficient, 860 selecting top segments by lexical Jaccard overlap between tokenized  $q$  and segment content.

862 **Sufficiency heuristic.** A sufficiency threshold  $u$  is used to label  $s_t^*$ : if the union of already-selected 863 and newly-picked positives reaches  $\geq u$ , mark `sufficient = 1` and stop; otherwise continue. Small  $u$  encourages minimal-evidence solutions.

864 **Trajectory construction.** Given  $P^*$  and a per-step cap  $k$ , a target sequence is synthesized by  
 865 greedily choosing up to  $k$  unseen positives at each step until sufficiency holds or candidates are  
 866 exhausted. Low-confidence choices (from lexical overlap rather than exact match) can be down-  
 867 weighted in the loss.

868 **Proxy selection metric.** During development, a lightweight proxy evaluates selection quality:  
 869 for a held-out set, the agent’s chosen ids are compared with target ids to compute micro Precision/Recall/F1 over segment identifiers. This tracks selection ability without requiring full QA eval-  
 870 uation.

873 **A.2.1 CANONICALIZATION AND SOUNDNESS**

875 **Definition 1** (Canonicalizer). *A canonicalizer  $\kappa$  maps  $M \subseteq S_h$  to a finite structure  $\kappa(M)$  consisting  
 876 only of tuples  $(\text{id}, \ell, \text{content\_view}, \text{provenance})$  where  $\text{content\_view}$  is a deterministic, lossless  
 877 projection of  $c(s)$  and  $\mu(s)$ , and provenance contains the fields needed to locate  $s$  in  $S_h$  (e.g., offsets  
 878 or coordinates). We say  $\kappa$  is content-preserving if for all  $M$ , the multiset  $\{(c(s), \mu(s)) : s \in M\}$  is  
 879 reconstructible from  $\kappa(M)$ .*

880 **Proposition 2** (Soundness and auditability). *If  $\text{Suff}(M)$  holds and  $\kappa$  is content-preserving, then the  
 881 head  $\mathcal{H}$  applied to  $(q, \kappa(M))$  is supported solely by items in  $M$ , and every atomic support can be  
 882 traced back to a unique segment in  $M$  via id and provenance.*

884 *Proof.* By content preservation,  $\kappa(M)$  contains all information from  $\{(c(s), \mu(s)) : s \in M\}$ ; there-  
 885 fore  $\mathcal{H}$  restricted to  $\kappa(M)$  depends only on evidence in  $M$ . Since  $\kappa$  stores id and provenance per  
 886 item, any atomic support used by  $\mathcal{H}$  can be mapped to a unique  $s \in M$ . Auditability follows.  $\square$

887 **A.2.2 PROBABILISTIC COMPLETENESS UNDER STOCHASTIC SELECTION**

889 We next quantify success probability for a stochastic policy that may fail to pick all supporting items  
 890 even if they appear early in the stream.

891 **Definition 2** (Exposure count). *Fix an admissible  $\rho$  with prefix length  $L$  and selection size  $k \geq 1$ .  
 892 Let  $R = \lceil L/k \rceil$ . An element  $e \in \{s_1, \dots, s_L\}$  is said to be exposed at steps  $1, \dots, R$ , meaning it  
 893 either is in the first window where it lies or remains eligible until selected; monotone refresh ensures  
 894 at most  $R$  exposures before all prefix items are exhausted.*

895 **Assumption 1** (Per-exposure success). *There exists  $p \in (0, 1]$  such that for every  $e \in E^*$  and for  
 896 every step  $t$  at which  $e$  is exposed and not yet selected, the policy includes  $e$  in  $K_t$  with probability  
 897 at least  $p$ , independently across steps for the same  $e$ .*

898 **Theorem 3** (Stochastic completeness). *Under admissibility and Assumption 1, with  $R = \lceil L/k \rceil$   
 899 and  $m = |E^*|$ , the probability that all items in  $E^*$  are selected within  $R$  steps is bounded below by*

$$900 \mathbb{P}[E^* \subseteq M_R] \geq 1 - m(1 - p)^R.$$

901 *Consequently, by Proposition 1, the probability that the loop halts by  $\min\{T_{\max}, R\}$  with a correct  
 902 answer is at least  $1 - m(1 - p)^R$ .*

904 *Proof.* Fix  $e \in E^*$ . By Assumption 1, across its at most  $R$  exposures, the probability that  $e$  is never  
 905 selected is at most  $(1 - p)^R$ . By the union bound over the  $m$  items in  $E^*$ ,

$$907 \mathbb{P}[\exists e \in E^* \text{ not selected by step } R] \leq m(1 - p)^R.$$

908 Taking complements yields the first claim. The second claim follows because once  $E^* \subseteq M_t$ ,  
 909  $\text{Suff}(M_t)$  holds and the stop head can fire; the hard cap can only make halting earlier.  $\square$

910 **A.2.3 DISCUSSION OF ASSUMPTIONS**

912 The injectivity result (Thm. 1) relies on invariants (T1)–(T3), which are satisfied by construction in  
 913 the HSEQ adapters (offsets and row indices/ordering are recorded; triplets are stored verbatim). Ad-  
 914 missibility is a regularity condition stating that an order  $\rho$  exists (often paragraph/row-first) placing  
 915 supporting segments early; in practice this is further improved by guidance. Assumption 1 abstracts  
 916 a calibrated selector that repeatedly assigns nontrivial probability mass to any exposed, still-missing  
 917 support item; the bound in Theorem 3 is conservative (union bound) and can be tightened under  
 additional structure (e.g., adaptive  $k$ , or margin assumptions on scoring).

918 A.3 IMPLEMENTATION DETAILS  
919920 A.3.1 AGENT MODELS USED FOR HSEQ  
921922 Models used for both iteration agent and head agent are shown in Table 6, grouped by size. Most  
923 experiments are done by using small and medium models (as of the result shown in main text).  
924925 Table 6: Iteration-agent and head agent base models grouped by size.  
926

Group	Model (HF id)
SMALL	tiuae/Falcon-H1-0.5B-Instruct
	tiuae/Falcon-H1-1.5B-Instruct
	tiuae/Falcon3-1B-instruct
	meta-llama/Llama-3.2-1B-Instruct
	deepseek-ai/DeepSeek-R1-Distill-Qwen-1.5B
MEDIUM	tiuae/Falcon3-3B-instruct
	tiuae/Falcon-H1-3B-Instruct
	Qwen/Qwen3-4B-Instruct-2507
	tiuae/Falcon3-7B-instruct
	tiuae/Falcon-H1-7B-Instruct
	meta-llama/Llama-3.2-3B-Instruct
	meta-llama/Meta-Llama-3-8B-Instruct
	meta-llama/Llama-3.1-8B-Instruct
	deepseek-ai/DeepSeek-R1-Distill-Qwen-7B
LARGE	deepseek-ai/DeepSeek-R1-Distill-Llama-8B
	tiuae/Falcon3-10B-instruct
	tiuae/Falcon-H1-34B-Instruct
	Qwen/Qwen3-30B-A3B-Instruct-2507
	deepseek-ai/DeepSeek-R1-Distill-Qwen-14B
	deepseek-ai/DeepSeek-R1-Distill-Qwen-32B
	deepseek-ai/DeepSeek-R1-Distill-Llama-70B
	meta-llama/Llama-3.1-70B-Instruct

946 A.3.2 ITERATION-AGENT PROMPTS AND OUTPUT SCHEMA  
947948 **System instruction.** The iteration agent is conditioned with a concise, role-defining system mes-  
949 sage:  
950951 You are an iteration agent working over a hierarchical  
952 sequence (H-Seq).  
953 Given a question and a list of candidate segments (each with  
954 an id and text)  
955 select the top-k segment\_ids that best support answering the  
956 question.  
957 Then decide if the selected evidence is sufficient to stop.  
958 Return ONLY compact JSON with keys: type, args.segment\_ids,  
959 args.strategy, args.top\_k, sufficiency.  
960 WITHOUT ANY EXPLANATION.961 **Prompt template.** Each training step uses a structured multi-section prompt:  
962

```

963     ### Instruction
964     {system-instruction}
965     ### Question
966     {q}
967     ### Guidance
968     {g(q,type)}
969     ### Selected-So-Far
970     - [seg_id] truncated_content
971     ...
972     ### Candidate-Window
973     - [seg_id] truncated_content
974     ...

```

972           `### Output (JSON)`  
 973  
 974       Only identifiers, levels, truncated content, and key metadata of segments are serialized.  
 975

976       **Output schema.** The agent must emit deterministic, machine-checkable JSON:

977  
 978       `{ "type": "select", "args": { "segment_ids": [...],  
 979        "strategy": "guided_topk", "top_k": k }, "sufficiency":  
 980        true/false }`

981       No free-form text is allowed. This constraint simplifies supervision and evaluation.  
 982

983       **Masking for SFT.** During supervised fine-tuning, the loss is applied only to the *output* portion of  
 984       the sequence (prompt tokens are masked), yielding a standard next-token objective over the action  
 985       string while keeping inputs loss-free.

### 986       A.3.3 GUIDANCE GENERATION AND CACHING

988       **Head-generated guidance.** A lightweight planner (“head”) converts  $(q, \text{type})$  into a short plan  $g$   
 989       that specifies: (i) what to retrieve first, (ii) optional branches, and (iii) a sufficiency hint. The planner  
 990       is prompted with:

992       `You are a planning assistant. Given a question, write  
 993       a short retrieval plan for an iteration agent selecting  
 994       evidence snippets. Specify ONLY what to retrieve first,  
 995       possible branches, and when to stop (sufficiency condition).`

996       A short completion is generated and, if too brief or incomplete, a single continuation is requested to  
 997       end with an explicit stop condition.

999       **Heuristic templates.** When a head is unavailable or for ablations, templates keyed by coarse pat-  
 1000       terns produce  $g$ , start with:

1001       `"Plan: retrieve a minimal set of highly relevant snippets; prefer  
 1002       concise facts."`

1004       Then add the following according to  $\mathbb{Q}_{\text{type}}$ :

- **Numeric:** Look for numeric mentions and table rows; stop when final number is explicit or corroborated.
- **Factoid (who/which/where/when):** Focus on short spans that directly contain the answer; stop on a clear statement..
- **Binary:** Retrieve one-two definitive statements; stop when evidence strongly supports yes/no..
- **Default:** Prefer snippets naming key entities/relations; stop when answer is explicitly stated.

1014       **Caching.** Guidance strings are cached per example using a stable key (dataset name and a hash of  
 1015        $q$ ) under a directory organized by head model id. Cache is consulted before running the head planner  
 1016       to reduce overhead.

1018       **Settings.** The planning head is run with short outputs and deterministic decoding. A minimal-  
 1019       length heuristic is applied to avoid truncated guidance.  
 1020

### 1021       A.3.4 LORA ADAPTATION AND OPTIMIZATION

1023       **Parameterization.** The iteration agent is obtained by adding low-rank adapters to a base causal  
 1024       LLM. Adapters are attached to attention projections ( $q_{\text{proj}}$ ,  $k_{\text{proj}}$ ,  $v_{\text{proj}}$ ,  $o_{\text{proj}}$ ) and  
 1025       MLP projections ( $\text{gate}_{\text{proj}}$ ,  $\text{up}_{\text{proj}}$ ,  $\text{down}_{\text{proj}}$ ); vocabulary and positional embeddings  
 1026       are unchanged.

1026 Table 7: Supervised fine-tuning (SFT) hyperparameters for each model-size group. These settings  
 1027 apply to all models within the corresponding group.

Group	Target steps	Batch	GA	LR	ML	MS	Top- $k$	Mi	BF16
SMALL	12000	2	8	$2.0 \times 10^{-5}$	3072	48	2	4	Yes
MEDIUM	9000	2	8	$1.5 \times 10^{-5}$	3072	48	4	4	Yes
LARGE	4500	1	16	$1.0 \times 10^{-5}$	2048	32	5	4	Yes

1033 **Notes.** *Batch* is `--per_device_train_batch_size`. *GradAcc* is `--grad_accum`. *LR* is `--lr`. *ML*,  
 1034 *MS*, *Top- $k$* , *Mi* map to `--max_length`, `--max_segments`, `--top_k`, `--max_iters`. *BF16* indicates  
 1035 `--bf16` enabled.

1037 **Default configuration.** LoRA rank  $r = 16$ , scaling  $\alpha = 32$ , dropout 0.05, no bias; the language  
 1038 head is preserved as a save-module. Mixed-precision and 4-bit weight quantization (NF4 with dou-  
 1039 ble quantization) are used to reduce memory. Gradient checkpointing is enabled.

1040 **Training schedule.** A cosine learning-rate schedule with warmup ratio 0.03 is used; batches are  
 1041 accumulated over several steps to match the target global batch size. Maximum input length is  
 1042 capped to a few thousand tokens; candidate windows and per-step  $k$  are tuned to respect the overall  
 1043 budget.

1045 **Mixture and curriculum.** Examples are sampled across datasets by normalized weights; quotas  
 1046 are computed for a target mixed size and shuffled. A short-to-long curriculum increases the maxi-  
 1047 mum number of steps  $T$  as training progresses.

## 1049 Finetuning Parameters

### 1051 A.3.5 CANONICALIZATION AND SUFFICIENCY

1053 **Canonical evidence package.** At termination, a modality-agnostic canonicalizer  $\kappa$  converts the  
 1054 selected set  $M_\tau$  into a compact, auditable structure

$$1055 \kappa(M_\tau) = \{(\text{id}, \text{level}, \text{uri}, \text{offsets}, \text{source\_type}, \text{snippet}; \text{meta})\}_{s \in M_\tau},$$

1057 with the following contract: (i) **id**: globally unique, deterministically derived (e.g.,  
 1058 `sha1(uri, offsets)`); (ii) **uri**: source identifier with version (e.g., document path or graph  
 1059 name); (iii) **offsets**: zero-based half-open character indices  $[a, b)$  into the *original* source; for  
 1060 tables,  $[i, j]$  denotes row/column coordinates; for KGs, `offsets = (-1, -1)`; (iv) **snippet**: a  
 1061 human-readable content aligned to sentence/field boundaries when possible; (v) **meta**: integrity  
 1062 and alignment helpers (`schema`, `time`, `source_version`, `sha1`). Duplicates are removed by  
 1063 `(uri, offsets)` and the package is *deterministically* ordered by `uri` then `offsets`. Typed  
 1064 views are derived on demand: `text`  $\Rightarrow$  spans with section/paragraph ids; `table`  $\Rightarrow$  `row_id`, `col_ids`,  
 1065 `schema`, `cell_coords`; `KG`  $\Rightarrow$   $(h, r, t)$  plus optional validity time.

1066 **Stopping signal.** The sufficiency head outputs  $s_t \in \{0, 1\}$  at each step. **Training targets** follow  
 1067 a coverage-based heuristic:  $s_t^* = 1$  if and only if the current  $M_t$  satisfies task-specific adequacy  
 1068 (e.g., contains at least one gold-positive segment; achieves full slot coverage for table QA; or yields  
 1069 a unique answer span/number under a fixed head). For weak supervision, per-step weights down-  
 1070 weight low-confidence positives (App. A.2). **Inference** uses a calibrated threshold  $\tau$  on the model’s  
 1071 sufficiency score  $\hat{p}_t$  and enforces a minimum step count  $T_{\min}$ :

$$1073 \text{stop at } \tau = \min\{t \geq T_{\min} : \hat{p}_t \geq \tau\} \text{ or when budget } B \text{ is exhausted.}$$

1074 Optionally, a lightweight contradiction checker triggers a one-shot refinement loop of at most  $\Delta$   
 1075 additional steps with tightened guidance  $g'$  and reduced budget  $B'$ . Thresholds  $(\tau, T_{\min})$  are selected  
 1076 on the development split and may be calibrated via temperature scaling.

### 1078 A.3.6 REPRODUCIBILITY NOTES

- 1079 • **Seed and sampling.** A fixed seed is used for example subsampling and order shuffling.

Table 8: Notation used throughout the paper.

Symbol	Meaning
$q$	Natural-language query (question).
$D = \{(x_j, m_j)\}_{j=1}^N$	Heterogeneous corpus with items $x_j$ and modality tags $m_j \in \{\text{text}, \text{table}, \text{kg}\}$ .
$m_j$	Modality label for the $j$ -th item (text / table / KG).
$\tau, \tau_m$	Modality-aware adapter; $\tau(D)$ produces the unified hierarchical sequence. $\tau_m$ is the adapter for modality $m$ .
$S_h$	The <b>HSEQ</b> (hierarchical sequence): $S_h = \bigsqcup_j \tau_{m_j}(x_j) \in \mathcal{S}^*$ .
$\mathcal{S}$	Segment universe. Each segment $s \in \mathcal{S}$ is a lightweight record.
$s = (\text{id}(s), \ell(s), p(s), c(s), \mu(s))$	Segment fields: unique identifier, level tag (granularity), parent pointer, compact human-readable content, standardized metadata.
$\ell(s)$	Level tag (e.g., document/paragraph/sentence, table_row/table_cell, triplet/subgraph).
$p(s)$	Parent pointer (container linkage) encoding locality in the hierarchy.
$c(s)$	Compact content snippet (text span / serialized table row / triple).
$\mu(s)$	Metadata with fixed keys (e.g., source_id, uri, offsets/coordinates, schema, time).
$\pi_\theta$	<b>HSEQ-I</b> iteration policy (LLM-based) with parameters $\theta$ ; operates over $(q, S_h)$ to select evidence iteratively.
$g = g(q, \text{type})$	Short <i>guidance</i> prior (from planner/head or heuristics) shaping early exploration and stop notion.
$B, B_t$	Budget (global / per-step): token, tool-call, step, and/or latency limits.
$M_t$	Selected-evidence set at step $t$ ; $M^*$ is the final selected set at termination.
$C_t$	Candidate window at step $t$ (bounded by window size and ordering).
$k, W$	Top- $k$ selection cap per step; window size $W$ for the exposed candidate stream.
$T_{\max}, T_{\min}$	Maximal and minimal number of iteration steps (cap and anti-early-stop).
$\rho$	Deterministic ordering over $S_h$ levels (e.g., paragraph $\prec$ row $\prec$ sentence $\prec$ triplet) to form the stream.
$\mathcal{N}(\cdot)$	Structure-aware neighborhood operators (parent/child, row/column, KG relation hops).
$a_t, s_t$	Action at step $t$ (e.g., select up to $k$ segments and/or expand neighborhoods) and sufficiency prediction $s_t \in \{0, 1\}$ .
$\Phi$	Budget-aware sufficiency criterion queried by the iterator to trigger termination.
$\kappa$	Canonicalizer mapping $M_\tau$ to provenance-preserving evidence package (ids, levels, offsets/coordinates, snippets).
$\mathcal{H}$	<b>HSEQ-H</b> head module for answer synthesis from $(q, \kappa(M_\tau))$ ; can also generate guidance $g$ .
$\xi$	Optional verifier; on contradiction detection, triggers a brief refinement loop with tightened $g'$ and reduced $B'$ .
$y, \hat{y}$	Gold answer and system prediction, respectively.
$E^*$	Minimally sufficient evidence set (w.r.t. a fixed answerer) for $q$ in $D$ .
Window( $\cdot$ ), Refresh( $\cdot$ )	Operators to expose a bounded candidate window and to advance it while removing already selected segments.
$\Delta$	Max number of additional refinement steps if the verifier $\xi$ requests a retry.

- **Segment capping.** The number of serialized candidate segments per step is capped to respect the overall token budget; truncation is applied to `content` strings for display.
- **Budget control.** Global limits on steps, tokens, and optional tool calls are enforced; guidance encourages early sufficiency.
- **Hardware.** Experiments are run on maximum 4 NVIDIA H200 Tensor Core GPU. Mixed-precision and 4-bit quantization substantially reduce memory; typical training runs fit on a single GPU.

#### A.4 NOTATIONS

Table 8 lists all symbols used in main context.

1134 A.5 EXAMPLE USING HSEQ  
11351136 A.5.1 CASE STUDY: GUIDED ITERATIVE RETRIEVAL ON HYBRIDQA  
11371138 **Setup.** Query  $q$ : “Who is the author of the novel that inspired the 2004 Russian film directed  
1139 by Timur Bekmambetov?” HSEQ-I (iterator): Qwen3-4B-Instruct-2507; HSEQ-H (head):  
1140 Falcon-H1-7B-Instruct. Guidance mode: head; source: cache (latency  $\approx 0.12$  ms).1141 **Head-generated guidance.** The head planner emits a short plan: (i) identify the 2004 Russian  
1142 film directed by Bekmambetov; (ii) locate the novel that inspired it; (iii) stop once the *author of that*  
1143 *novel* is found. This plan is injected as a prefix and acts as a soft prior on where the iterator should  
1144 probe first.1145 **Guided iteration over  $S_h$ .** The iterator consumes the guidance and operates over the HSEQ  
1146 stream with a fixed window and top- $k$  selection. Table 9 summarizes the six steps (all sufficiency  
1147 flags were `false`; the loop terminates by budget).1148 Table 9: Stepwise selection (abridged). Segment ids prefixed by level:  $p_-$  (paragraph),  $row_-$  (table  
1149 row).  
1150

Step	Key picks (content excerpt)	Sufficient?
1	$p\_6df9c849$ : “Night Watch (...) is a 2004 Russian ... directed by Timur Bekmambetov. It is loosely based on the novel <i>The Night Watch</i> by Serg[ei Lukyanenko]...”	No
2	$p\_c15173df$ , $p\_3bc4a108$ , $p\_54f6ef94$ : contextual paragraphs (“List of Russian films of 2004”, “2004” entries)	No
3	$row\_a44a4a17$ : table row confirming <i>Night Watch</i> with director “Timur Bekmambetov”	No
4–6	additional table rows from the same list ( <i>Arie</i> , <i>Countdown</i> , <i>Dad or Papa</i> , etc.) providing film set context	Yes

1151 **Answer synthesis.** After  $\tau=6$  iterations, the canonicalizer  $\kappa$  compacts the selected set  $M_\tau$  (para-  
1152 graph + corroborating table rows) into a provenance-preserving package (segment ids, levels, offsets,  
1153 snippets). The head  $\mathcal{H}$  is prompted *only* with  $(q, \kappa(M_\tau))$  and outputs:

1154 
$$\hat{y} = \text{Sergei Lukyanenko}.$$

1155 The prediction matches the gold answer (EM/F1 = 1.0). Runtime profile: selection latency  
1156  $\approx 32,185$  ms, head latency  $\approx 1,826$  ms, total  $\approx 34,011$  ms; number of iterations = 6.1157 **Takeaway.** Guidance steers the iterator to a high-yield paragraph in the first step, which already  
1158 contains the sufficient evidence (film identity and source novel). Subsequent steps provide cor-  
1159 roboration from structured rows. The provenance in  $\kappa(M_\tau)$  makes the final answer auditable: the  
1160 paragraph  $p\_6df9c849$  explicitly ties *Night Watch* (2004, Bekmambetov) to the novel *Night Watch*  
1161 by Sergei Lukyanenko, enabling concise and well-grounded answer synthesis by the head.1162 A.5.2 CASE STUDY: GUIDED ITERATIVE RETRIEVAL ON HOTPOTQA  
11631164 **Setup.** Query  $q$ : “Which style is the building located on the East Side of Midtown Manhattan that  
1165 Robert Von Ancken appraised?” HSEQ-I (iterator): Qwen3-4B-Instruct-2507; HSEQ-H  
1166 (head): Falcon-H1-7B-Instruct. Guidance mode: head; source: generated online (latency  
1167  $\approx 8,496$  ms).1168 **Head-generated guidance.** The head planner issues a short plan: (i) identify buildings on the  
1169 East Side of Midtown Manhattan connected to appraiser *Robert Von Ancken*; (ii) once the specific  
1170 building is found, retrieve its architectural style; (iii) stop when the style is clearly linked to the  
1171 appraised building.

1188     **Guided iteration over  $S_h$ .** The iterator follows the guidance with a fixed window and top- $k$  se-  
 1189     lection. Table 10 lists the six steps (all sufficiency flags `false`; termination by budget). Note that  
 1190     Step 1 already surfaces the key paragraph about the Chrysler Building.  
 1191

1192     Table 10: Stepwise selection (abridged). Segment ids prefixed by level:  $p_-$  (paragraph).  
 1193

Step	Key picks (content excerpt)	Sufficient?
1	$p\_a73a8d8f$ : “The <i>Chrysler Building</i> is an <b>Art Deco-style</b> skyscraper located on the East Side of Midtown Manhattan ...”	No
2	$p\_c01522d2$ : “23 Beekman Place ... apartment building ... East Side of Midtown Manhattan ...”	No
3	$p\_7c2aa386$ : “The Helmsley Building ... Midtown Manhattan ...”	No
4	$p\_658d6333$ : “ <i>Robert Von Ancken</i> is a prominent New York City real estate appraiser ...”	No
5	$p\_e97ef7e6$ : “Lenox Hill Neighborhood House ... East Side of Manhattan ...”	Yes

1204     **Answer synthesis.** After  $\tau=5$  iterations, the canonicalizer  $\kappa$  compacts the selected set  $M_\tau$  (in-  
 1205     cluding  $p\_a73a8d8f$  and the Von Ancken paragraph  $p\_658d6333$ ) into a provenance-preserving  
 1206     package. The head answers using only  $(q, \kappa(M_\tau))$ :

1207     
$$\hat{y} = \text{Art Deco-style skyscraper}.$$

1208     The prediction matches the gold answer. Runtime profile: selection latency  $\approx 32,153$  ms, head  
 1209     latency  $\approx 838$  ms, total  $\approx 41,487$  ms; iterations = 5.  
 1210

1212     **Takeaway.** The head’s guidance steers the iterator directly to a paragraph that states both the loca-  
 1213     tion (East Side of Midtown) and the architectural style (Art Deco) of the relevant building (Chrysler  
 1214     Building), while additional picks provide neighborhood and appraiser context. Provenance in  $\kappa(M_\tau)$   
 1215     supports auditable linking from the final answer to its evidence.  
 1216

1217     A.6 STATEMENT FOR THE USE OF LARGE LANGUAGE MODELS (LLMs)  
 1218

1219     We used large language models (LLMs) as general-purpose tools for *writing assistance* and *engi-  
 1220     neering support*. For writing, we employed LLMs to improve clarity and style (e.g., rephrasing  
 1221     sentences, tightening paragraphs, standardizing notation, and proofreading grammar). Drafting,  
 1222     technical claims, algorithms, proofs, experiment design, and all final wording were authored and  
 1223     verified by the authors. For engineering, we consulted LLMs for debugging; all research code, data  
 1224     processing, and experiment scripts were implemented, audited, and executed by the authors. No text  
 1225     or code generated by an LLM was used verbatim without author review; we take full responsibility  
 1226     for the content.  
 1227

1228  
 1229  
 1230  
 1231  
 1232  
 1233  
 1234  
 1235  
 1236  
 1237  
 1238  
 1239  
 1240  
 1241