

Respecting Temporal-Causal Consistency: Entity–Event Knowledge Graphs for Retrieval-Augmented Generation

Anonymous ACL submission

Abstract

Retrieval-augmented generation (RAG) based on large language models often falters on narrative documents with inherent temporal structures. Standard unstructured RAG methods rely solely on embedding-similarity matching and lack any general mechanism to encode or exploit chronological information, while knowledge graph RAG (KG-RAG) frameworks collapse every mention of an entity into a single node, erasing the evolving context that drives many queries. To formalize this challenge and draw the community’s attention, we construct ChronoQA, a robust and discriminative QA benchmark that measures temporal, causal, and character consistency understanding in narrative documents (e.g., novels) under the RAG setting. We then introduce *Entity-Event RAG* (E²RAG), a dual-graph framework that keeps separate *entity* and *event* subgraphs *linked by a bipartite mapping*, thereby preserving the temporal and causal facets needed for fine-grained reasoning. Across ChronoQA, our approach outperforms state-of-the-art unstructured and KG-based RAG baselines, with notable gains on causal and character consistency queries. E²RAG therefore offers a practical path to more context-aware retrieval for tasks that require precise answers grounded in chronological information.

1 Introduction

Large language models (LLMs) have demonstrated remarkable zero-shot and few-shot capabilities across various NLP tasks. Yet, LLMs remain fundamentally constrained by their fixed context window: once the prompt exceeds a certain length, processing becomes slow and the model’s reasoning capability is significantly impaired (Liu et al., 2023; Fiction.live, 2025; Li et al., 2024; Gao et al., 2024; An et al., 2024). To mitigate this limitation, retrieval-augmented generation (RAG) was introduced, coupling the generator with an exter-

nal document retriever so that fresh, grounded evidence can be injected on demand (Lewis et al., 2020). RAG not only extends an LLM’s effective memory but also reduces hallucinations and allows rapid knowledge updates without costly re-training (Shuster et al., 2021; Lewis et al., 2020; Borgeaud et al., 2022; Béchar and Ayala, 2024).

Vanilla unstructured RAG, however, treats each passage in isolation and therefore struggles with reasoning that depends on a document’s temporal or causal structure. For example, the following seemingly straightforward question for *Harry Potter* enthusiasts can be challenging for a RAG system to answer reliably:

Query 1. *Who was jinxing Harry’s broom during his first Quidditch match?*

Because effects in a story typically follow their causes in time, losing chronological ordering also obscures causal links. If a RAG system naively retrieves the seemingly most relevant passage from the earlier part of the story, it will incorrectly conclude that *Snape* was responsible for hexing Harry’s broom. This error arises because it overlooks crucial information revealed later—that *Professor Quirrell* was actually causing the interference, while *Snape*’s suspicious actions were attempts to counteract Quirrell’s jinx. The key reason is that the document’s *chronological context information*, which is crucial for generating answers that are both temporally coherent and causally sound, is not preserved in the vanilla unstructured RAG after being chunked.

To capture richer structure, knowledge-graph (KG) RAG methods such as GraphRAG build an entity graph over the corpus and retrieve communities of related nodes (Edge et al., 2024). Unfortunately, the standard graph-construction pipeline relies on aggressive de-duplication of entity mentions; this collapses distinct temporal or contextual facets of the same character into a single node, discarding

information that is crucial in domains like narrative fiction, where characters evolve continuously. Consider the following query:

Query 2. *How would Hermione Granger react if a friend proposed breaking school rules after the troll incident?*

Hermione is introduced as a rule-obsessed know-it-all who reprimands classmates for the smallest infractions (Ch. 9), yet after the troll incident (Ch. 10), she forges a deep friendship with Harry and Ron and is soon helping them brew an illegal potion and sneak past teachers to protect the Stone. If every “Hermione Granger” mention is collapsed into a single KG node, a query will likely retrieve only her predominant rule-abiding persona—the class monitor who scolds Ron for casting *Lumos* in the corridor—while overlooking later chapters in which she calmly steals ingredients and slips past teachers to protect the Stone (Ch. 16). Thus, the aggressive de-duplication of entity mentions in KG-RAG methods can erase her arc from strict disciplinarian to pragmatic rule-breaker, masking precisely the nuance the question probes and leading the system to an outdated, inaccurate answer.

To rigorously define the challenge of temporal-causal consistency in RAG applications and evaluate how current RAG systems handle the aforementioned narrative-focused challenges, we construct **ChronoQA**, a retrieval-style QA benchmark drawn from nine public-domain narrative documents. Every question targets one of eight fine-grained reasoning facets, ranging from causal and character consistency to symbolism and thematic insight.

To address the aforementioned challenges, we propose an *Entity–Event KG* composed of (i) an entity subgraph; (ii) an event subgraph; and (iii) a *bipartite edge set* that maps each entity mention to the events in which it participates, thereby preserving those nuances. Instead of collapsing duplicates, we first extract both entities *and* their associated events, then link every event to the *specific* mention of each participating entity; because we never merge mentions that arise in different parts of the story, each entity node carries its own *context-specific description*. The resulting graph retains multiple, context-specific representations of entities while still exposing the relational structure needed for retrieval. We term the resulting RAG framework, which retrieves over the proposed entity–event KG, *Entity–Event RAG* (E^2RAG). At

retrieval time, beyond the vanilla pipeline, we introduce an augmentation operation that calls the backbone LLM only once to inject richer context. Empirically, this single-call augmentation lets E^2RAG achieve significant gains on fine-grained narrative-reasoning queries, outperforming state-of-the-art unstructured and KG-based RAG baselines and pointing to a practical path toward more context-aware retrieval for LLMs.

The contributions of this work are as follows:

- We release *ChronoQA*, the first open, passage-grounded benchmark that stresses temporal and causal reasoning over book-length narratives under a RAG setting.
- We propose E^2RAG , a dual-graph retrieval framework that keeps every entity mention distinct and grounded in its associated event, thereby preserving the causal and temporal consistency of evolving characters and plot.
- Extensive experiments demonstrate that variants of E^2RAG outperform state-of-the-art unstructured and KG-based RAG baselines on *ChronoQA*.

2 Related Work

2.1 Retrieval-Augmented Generation (RAG)

Although today’s large language models can store vast amounts of world knowledge, that knowledge is *static*, *unverifiable*, and *expensive* to refresh. Lewis et al. (2020) demonstrated that attaching a live retriever to a generator lets the model pull *up-to-date evidence on demand*, achieving higher accuracy than closed-book BART while returning the very passages that support each claim. Subsequent analyses show that retrieval-augmented language models markedly cut hallucination rates—especially on fact-heavy tasks—because retrieved text provides an external “ground truth” that the decoder can copy or paraphrase instead of guessing (Lv et al., 2024; Sree Mala et al., 2025). Gao et al. (2023) go a step further: their HYDE method lets the LLM hallucinate a “hypothetical” answer, embeds it, and uses that vector to query the index, delivering *precise zero-shot dense retrieval* without relevance labels—an idea we later adapt in Section 4. RAG also sidesteps model-update costs: keeping the index current is far cheaper and faster than re-training or fine-tuning multi-billion-parameter networks, a point underscored by adaptive systems such as Self-RAG and Speculative RAG that retrieve only when necessary and verify drafts to stay both ef-

ficient and factual (Lv et al., 2024; Lewis et al., 2020). Finally, long-context methods like LongRAG show that retrieval can extend an LLM’s effective memory without quadratic attention overhead, enabling faithful reasoning over book-length evidence while keeping latency low (Jiang et al., 2024). In short, RAG equips LLMs with a dynamic, interpretable and cost-effective memory, tackling three core limitations—knowledge staleness, hallucination, and context length—that purely parametric models struggle to overcome.

2.2 Knowledge Graph based RAG

GraphRAG (Edge et al., 2024) explicitly builds an entity-level knowledge graph, run community detection, and then retrieve and summarize the most query-relevant communities, thereby improving the relevance of the top- k chunks with respect to a query. Follow-up studies highlight three practical drawbacks of this design: (i) the multi-pass entity/relation extraction and community-summary generation make preprocessing costly in both tokens and compute, (ii) traversing and summarising the graph at inference time adds 2–3 \times higher end-to-end latency, and (iii) the graph index and its summaries grow super-linearly with corpus size, complicating incremental updates and ballooning memory usage (Wang et al., 2025; Chen et al., 2025; Peng et al., 2024).

LightRAG (Guo et al., 2024) tackles these limitations by folding relational signals into a standard dense index and introducing a dual-level, coarse-to-fine retriever that first selects cluster representatives and then expands to their ego networks; this removes explicit graph traversal, supports incremental index patches, and reduces indexing token cost by $\approx 60\%$ while roughly halving median query latency, all without hurting answer quality on UltraDomain, QFS and other multi-hop QA benchmarks (Guo et al., 2024; Chen et al., 2025).

2.3 Long-context narrative benchmarks

Recent work has begun to probe whether LLMs can reason over book-length inputs without truncation. XL²BENCH (Ni et al., 2024), LOONG (Wang et al., 2024), and LONGGENBENCH (Liu et al., 2024) extend QA or generation tasks to 100 K-token contexts, while the very recent FICTION.LIVEBENCH (Fiction.live, 2025) packages full user-written stories into the prompt and asks multi-step comprehension questions that require tracking characters and foreshadowing across tens

of thousands of tokens. Because every benchmark above *gives the model the entire story up front*, they measure intrinsic long-context reasoning rather than the retrieval quality.

3 New Benchmark: ChronoQA

Although the datasets mentioned above for long context reasoning tasks can be adapted to benchmarking the RAG methods, there are very limited resources that can be used to evaluate an RAG method specifically on the temporal-causal consistency question: (i) Most existing datasets focus on general long context tasks, but have limited focus on these challenging reasoning tasks related to temporal-causal consistency. For example, a rare existing task requires reasoning over a character’s time-specific state (e.g., "late-story Hermione" versus "early-story Hermione"). (ii) Existing benchmarks typically provide *no* passage-level evidence for their ground-truth answers. Different from the long context tasks, which only evaluate the model answers, evaluating the retrieval snippets is also a common metric when benchmarking the RAG systems. However, without the passage-level evidence, extra effort may be required to determine the retrieval stage’s correctness. (iii) Although some datasets (e.g., Fiction.live (2025)) also focus tasks based on narrative documents, they rely on a entirely private evaluation framework with no publicly available data. These existing issues make it difficult for researchers to verify reported results or to probe where and why current RAG systems fail.

To test whether different RAG systems can reason over a document’s *temporal* and *causal consistency*, we introduce **ChronoQA**—a QA benchmark built from nine narrative works spanning novels, musical scripts and children’s stories.¹ For each story we automatically generate questions that probe eight fine-grained reasoning facets:

- *Causal Consistency* – cause and effect, logical sequences, or explanations of how events unfolded.
- *Character & Behavioural Consistency* – character motivations, development, or psychology.
- *Setting, Environment & Atmosphere* – physical locations, time periods, or mood/atmosphere
- *Symbolism, Imagery & Motifs* – symbolic elements, recurring imagery, or metaphorical representations.

¹The underlying literary works are in the public domain in the United States. The e-book files were obtained from Project Gutenberg and are redistributed under the Project Gutenberg-tm License that accompanies each file.

- *Thematic, Philosophical & Moral* – deeper meanings, philosophical ideas, or ethical implications
- *Narrative & Plot Structure* – story organization, plot devices, or narrative techniques.
- *Social, Cultural & Political* – societal contexts, cultural elements, or political dimensions.
- *Emotional & Psychological* – emotional responses, psychological states, or mental processes.

The questions from each category require solid understanding grounded on the progression of the story and cause-and-effect of the elements involved in the plot. Our **ChronoQA** closes the aforementioned gap with:

- **Focus on temporal and causal consistency** — each query hinges on a character’s *specific temporal facet* (e.g. "after the troll incident"), so retrieving an early-story snippet fails.
- **Passage-level supervision** — every answer is paired with exact start/end byte offsets, making verification straightforward and reliable.
- **Unrestricted accessibility** — every narrative is drawn from Project Gutenberg, putting the full texts in the public domain. Unlike fully private benchmark, This guarantees that anyone can obtain, inspect, and redistribute the benchmark data without restrictions.

Thus ChronoQA tests whether a RAG system can (1) fetch the right snippet when the full story cannot fit into context and (2) reason over evolving entity states that standard KG de-duplication erases.

Dataset statistics. The final release contains 9 stories, 497 question–answer pairs. The actual number of question–answer pairs for each story and categories can be found in Table 1 and 2. Other details can be found in Appendix A.

Table 1: Questions per story (total = 497).

Story	# Questions
The Wonderful Wizard of Oz	82
The Hound of the Baskervilles	72
The Phantom of the Opera	70
A Study in Scarlet	67
The Sign of the Four	62
Harry Potter and the Chamber of Secrets	55
The Adventures of Sherlock Holmes	34
Harry Potter and the Sorcerer’s Stone	30
Les Misérables	25
Total	497

Generation pipeline. Because every story is short enough to fit within the oracle’s context window, we do not chunk the text. Instead we use a *two-stage* pipeline:

Table 2: Questions per reasoning category (total = 497).

Category	# Questions
Character Consistency	229
Causal Consistency	96
Symbolism, Imagery & Motifs	56
Thematic, Philosophical & Moral	36
Narrative & Plot Structure	31
Setting, Environment & Atmosphere	25
Social, Cultural & Political	22
Emotional & Psychological	2
Total	497

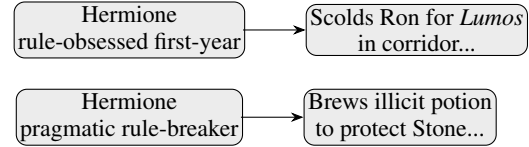


Figure 1: Illustration of two temporally distinct Hermione mentions and their associated events.

• **Q–A candidates generation:** The entire story is provided to the oracle model (GPT-o1-pro (Jaeche et al., 2024) and Claude-3.7-Sonnet (Anthropic, 2025)) together with a category list and an instruction to propose up to diverse question–answer pairs per category (prompt template in Appendix D.1). The model’s answer is retained as provisional ground truth.

• **Verification, Filtering and de-duplication:** For each ground truth, the oracle model is required to output the starting sentence and ending sentence of the excerpt where the answer can be inferred. Together with the starting sentence and the ending sentence, we also provide the excerpt and its exact start/end byte offsets in the dataset. We discard candidate questions that are duplicates or for which the excerpt extracted does not substantiate the answer.

4 Methodology

Standard KG-RAG preprocessing merge every mention of entity extracted from chunked document (e.g., *Hermione Granger*, *Professor Quirrell*, or the *Sorcerer’s Stone*) into a single node—obliterating the time-specific information. Such nuance is needed to answer questions such as Query 2. We therefore keep **each** mention—a concrete, context-specific instance of an entity as it appears in a single chunk of text (e.g., "Hermione" in Chapter 9 versus "Miss Granger" in Chapter 16) — distinct and tether it to the *event snippet* in which it occurs. Intuitively, the structure is as shown in Figure 1. The left column holds *entity mentions* and their contextual descriptions; the right column holds *events* and their descriptions. Treating the

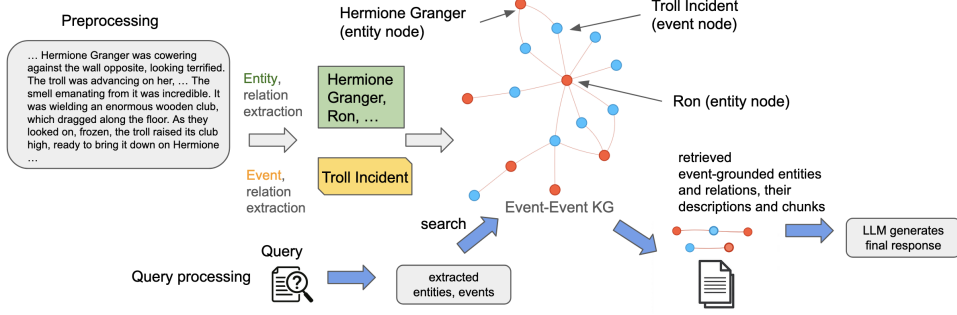


Figure 2: Overall architecture of the E²RAG framework.

entity as two disjoint vertex sets ensures temporal facets never collapse.

Formal definition. Let $\mathcal{G}_{\text{ent}} = (V_{\text{ent}}, E_{\text{ent}})$ be the directed graph of entity mentions and $\mathcal{G}_{\text{evt}} = (V_{\text{evt}}, E_{\text{evt}})$ the graph of events. We connect the two with a bipartite edge set

$$\mathcal{B} = \{(v_i, e_j) \mid v_i \in V_{\text{ent}}, e_j \in V_{\text{evt}}, \text{Name}(v_i) \subset \text{Desc}(e_j)\}. \quad (1)$$

so that an edge exists if and only if the entity’s name $\text{Name}(v_i)$ appears in the event description $\text{Desc}(e_j)$.

Preprocessing. E²RAG consists of the following preprocessing steps to build knowledge graphs².

1. *Chunking.* The document is split into chunks $\{c_\ell\}_{\ell=1}^L$ and stored in a key–value store.

2. *Creating entity and event sets.* Each chunk c_ℓ is fed twice to an LLM to extract *entities* $V_{\text{ent}}^{(\ell)}$ and *events* $V_{\text{evt}}^{(\ell)}$ respectively. For every entity and event extracted, we also ask the LLM for a one-sentence description. Note that in entity’s case, this description will be *context-aware* so the same entity extracted from different document chunks will have different descriptions³. This description is stored in the node’s description field and is kept *distinct* for every mention. The extractions form the node sets $V_{\text{ent}} = \bigcup_\ell V_{\text{ent}}^{(\ell)}$ and $V_{\text{evt}} = \bigcup_\ell V_{\text{evt}}^{(\ell)}$. Co-mentions inside the same chunk yield the directed edge sets E_{ent} and E_{evt} exactly as in LightRAG.

3. *Link entities to events (\mathcal{B}).* For every entity $v \in V_{\text{ent}}^{(\ell)}$ we scan the events $e \in V_{\text{evt}}^{(\ell)}$; if the node’s name $\text{Name}(v)$ appears in $\text{Desc}(e)$ we add the bipartite edge (v, e) to \mathcal{B} (Eq. 1).

²Pseudocode can be found as Algs. 1 and 2 in the appendix.

³e.g. "Hermione—rule-obsessed first-year" versus "Hermione—pragmatic rule-breaker".

4. *Embed and index.* Each node’s name and description is embedded with embedding model $f(\cdot)$ and written to two vector stores—one for entities, one for events—while the full graph $\mathcal{G}_{\text{ent}} \cup \mathcal{G}_{\text{evt}} \cup \mathcal{B}$ is saved for hop-limited traversals.

Figure 2 shows the overall pipeline of E²RAG and Figure 4 shows a subgraph of the Entity-Event KG of *Harry Potter and the Sorcerer’s Stone* and an event node’s meta data.

Query-time retrieval. Given a query q , the retrieval mechanism conducts the following steps.

1. *Cue extraction.* A pair of entity and event phrase sets are first extracted by an LLM extractor $g(\cdot)$, namely $(S_{\text{ent}}, S_{\text{evt}}) = g(q)$

2. *Embedding generation.* The $(S_{\text{ent}}, S_{\text{evt}})$ will be fed into the embedding model $f(\cdot)$ and generate a embedding z used for retrieval, namely $z = f(g(q))$.

3. *Seed nodes retrieval.* With the embedding z , the next step is to retrieve seed nodes vector stores $V_q \subseteq V_{\text{ent}} \cup V_{\text{evt}}$.

4. *One-hop expansion.* We take one step over the bipartite edges, $V_q^+ = V_q \cup \{e \mid \exists v \in V_q : (v, e) \in \mathcal{B}\} \cup \{v \mid \exists e \in V_q : (v, e) \in \mathcal{B}\}$, guaranteeing that every retrieved entity comes packaged with its time-specific events and vice-versa.

5. *Similarity ranking.* All passages are embedded offline; at query time we select only the vectors of nodes in V_q^+ —typically top- k based on their similarity instead of the full index. When multiple entity nodes correspond to the *same* entity, their similarity scores are re-ranked according to the similarity ranking of their *associated event* nodes. This grounds each entity in the events’ evidence, preserving temporal–causal consistency.

6. *Context assembly.* We collect (i) the raw passages behind the top- k nodes; (ii) a linearised dump

of the subgraph $\mathcal{G}_{\text{sub}} = (V_q^+, (E_{\text{ent}} \cup E_{\text{evt}} \cup \mathcal{B}) \cap (V_q^+ \times V_q^+))$; and (iii) the metadata (descriptions, labels, etc.) associated with every selected node and edge. This enriched, structured context is fed to the backbone LLM, allowing it to answer questions that demand fine-grained temporal and causal reasoning using the truly relevant document chunks.

Together, these two routines turn the intuition from Figure 1 into a fully operational retrieval pipeline, with \mathcal{B} acting as the critical "glue" that preserves evolving entity states across time.

Hypothetical response coupling. One potential limitation of E²RAG is that the retrieval effectiveness heavily depends on how much information can be extracted from the query. If the query lacks details, the extraction might not be able to capture enough information, preventing the core mechanism from being fully utilized. To make the *entity-event* approach more effective, we incorporate the core idea in HyDE (Gao et al., 2023). HyDE first asks the backbone LLM to draft a *hypothetical response* without the document, then merges the embedding of the hypothetical response with the original query before performing similarity-based top- k retrieval. Even though the content might be factually off, the inclusion of hypothetical response provides a richer context for similarity matching.

We introduce four variants of hypothetical response mechanism in E²RAG for retrieving more accurate and comprehensive information. To clearly present the them, we use h denote hypothetical response and $[q; h]$ as text-level concatenation of the original query q and hypothetical response h .

1. *Combined extraction (Comb. extraction).* Instead of providing the the query q to the extractor, this approach gives $[q; h]$ as the input to the extractor. Thus, the generated embedding can be represented as $z = f(g[q; h])$.

2. *Hypothetical extraction (Hyp. extraction).* Similarly, this variance replaces the original query q with the hypothetical answer h to the extractor, and use the output to generate embeddings, i.e., $z = f(g(h))$.

3. *Combined embedding (Comb. embedding).* This method omits the extractor, and directly embeds the concatenated the hypothetical answer and the original query, i.e., $f([q; h])$.

4. *Hypothetical embedding (Hype. embedding).* Similar to the above one in terms of removing the extractor step, we directly generate embeddings with the hypothetical documents, i.e., $z = f(h)$.

We experiment with these four variants to investigate the importance of two key factors: incorporating the original query and employing the extraction step. Specifically, we examine how the presence or absence of the original query and the extraction mechanism affect retrieval effectiveness. After generating embeddings using each approach, the following steps are identical to the query-time retrieval steps 3 to 6 introduced earlier.

5 Experiments

Baselines. We compare the five variants of E²RAG (four hypothetical response variants, one without hypothetical response) against the three modes of LightRAG (Guo et al., 2024) (local, global, hybrid), three modes of GraphRAG (Edge et al., 2024) (local, global, drift), RQ-RAG (Chan et al., 2024), vanilla HyDE (Gao et al., 2023) as well as vanilla RAG. In addition, we also give the hybrid mode of LightRAG the same four variants with the hypothetical response (thirteen in total excluding all the variants of our own methods). All methods use GPT-4o-mini (Hurst et al., 2024) as the backbone LLM for preprocessing, inference, or both, and text-embedding-3-small (OpenAI, 2024) as the embedding model $f(\cdot)$.

Evaluation. For each query, we give an LLM judge the query, the ground truth and the responses of the variants of E²RAG and all baselines. Each answer is graded independently by $J = 3$ LLM judges (Claude-3.7-Sonnet, GPT-4o, GPT-4.1-mini (OpenAI, 2025)) on a 1–10 Likert scale using the rubric in Appendix D.2. The overall quality score for a system is the mean of those ratings across all samples and judges, computed as

Table 3: Overall average scores

Rank	Mode	Avg Score	Total
1	E ² RAG (comb. extraction)	7.125 7	10 603
2	E ² RAG (comb. embedding)	7.071 9	10 523
3	E ² RAG (hyp. extraction)	6.983 2	10 391
4	E ² RAG (hyp. embedding)	6.939 5	10 326
5	LightRAG hybrid	6.880 4	10 238
6	GraphRAG drift	6.820 6	10 149
7	GraphRAG local	6.799 7	10 118
8	E ² RAG (vanilla)	6.708 3	9 982
9	vanilla RAG	6.602 2	9 824
10	LightRAG local	6.549 7	9 746
11	GraphRAG global	6.508 7	9 685
12	LightRAG global	6.458 3	9 610
13	vanilla HyDE	6.355 5	9 457
14	LightRAG hybrid (comb. embedding)	5.696 9	8 477
15	LightRAG hybrid (comb. extraction)	5.681 5	8 454
16	LightRAG hybrid (hyp. embedding)	5.671 4	8 439
17	LightRAG hybrid (hyp. extraction)	5.655 2	8 415
18	RQ-RAG	3.514 1	5 229

Table 4: Preprocessing runtimes for the three KG RAG systems

System	Stage	Time (s)
LightRAG	Total pipeline	101.1723
E ² RAG	Entity subgraph construction	108.3549
	Event subgraph construction	103.6091
	Bipartite-mapping	0.5076
	Total pipeline	108.8625
GraphRAG	Total pipeline	208.9904

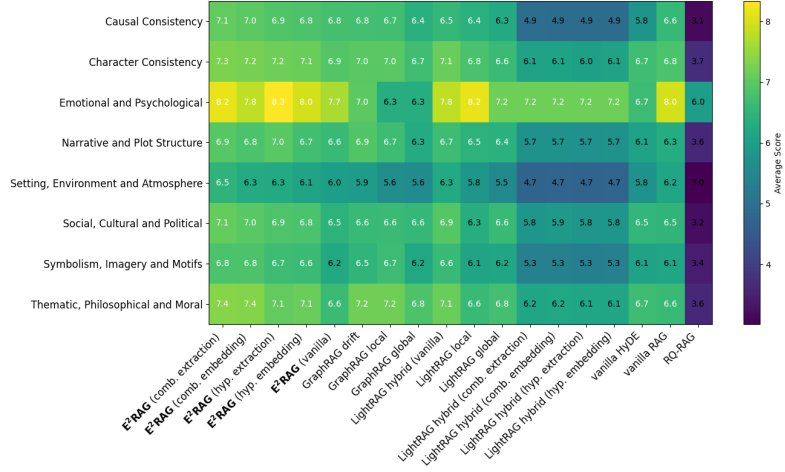


Figure 3: Heatmap of the scoring for each category and mode.

$$\text{Score} = \frac{1}{J} \sum_{j=1}^J \frac{1}{N} \sum_{i=1}^N s_{ij}, \quad (2)$$

where N is the number of question–answer pairs in the test set; J is the number of LLM judges; and $s_{ij} \in \{1, \dots, 10\}$ is the score assigned by judge j to sample i .

We report the (*mean score*) in Table 3; per-judge scores are provided in Appendix G.

Result. In the overall ranking, the top-3 modes are all hypothetical response variants of E²RAG, while E²RAG (vanilla) loses to GraphRAG local, GraphRAG drift and LightRAG hybrid mode. This corroborates that hypothetical responses play a major role in improving retrieval quality. In particular, the superior performance of the combined extraction variant, followed by the combined embedding variant, confirms that both incorporating the original query and employing the extraction step further enhance the retrieval effectiveness. In contrast, all hypothetical response variants of LightRAG hybrid mode performed poorly. We speculate that because E²RAG does not employ de-duplication, it synergizes particularly well with the hypothetical response, achieving top performance regardless of the specific variant used. Furthermore, the hypothetical response provides a richer context and an abundant number of entity and event candidates. This can be beneficial to the KG matching when there are event nodes to provide grounding to the entities to mitigate hallucination. Figure 3 plots the category-average scores (averaged over the three evaluator LLMs). In every category, the top-scoring model is again an E²RAG hypothetical

response variant. For more details, refer to Appendix G.

Table 5: Average query time per method (seconds)

Method	Avg. Time
GraphRAG drift	93.154 7
GraphRAG global	26.232 6
GraphRAG local	18.014 1
RQ-RAG	9.926 4
E ² RAG (comb. extraction)	8.720 0
E ² RAG (hyp. embedding)	8.085 0
E ² RAG (comb. embedding)	8.054 4
E ² RAG (hyp. extraction)	8.011 6
E ² RAG (vanilla)	7.506 5
LightRAG hybrid (comb. extraction)	7.492 4
LightRAG hybrid (hyp. extraction)	7.137 9
LightRAG hybrid (hypo. embedding)	6.714 2
LightRAG hybrid (comb. embedding)	6.695 1
LightRAG hybrid (vanilla)	6.156 4
LightRAG local	5.235 1
vanilla HyDE	2.910 6
LightRAG global	1.903 1
vanilla RAG	1.422 4

Preprocessing time cost. Table 4 shows the preprocessing time of the novel *The Phantom of the Opera* for three KG based RAG approaches used in our experiments.

Compared to LightRAG, although E²RAG needs to construct two subgraphs ($\mathcal{G}_{\text{ent}}, \mathcal{G}_{\text{evt}}$) as opposed to just event KG, and form \mathcal{B} (Bipartite mapping), the two subgraphs’ construction are independent and can be carried out in parallel. Forming \mathcal{B} is sequential after the subgraphs’ construction, but for ($\mathcal{G}_{\text{ent}}, \mathcal{G}_{\text{evt}}$) of reasonable size such as in this example, its cost is negligible in practice. As a result, E²RAG has preprocessing time comparable to LightRAG (subject to API query traffic fluctuation),

Table 6: GPT-4.1-mini’s verdicts (selected) of the responses to “Consider the dinner scene in which Holmes details his reasoning about the Lauriston Gardens mystery. How does Watson’s narration highlight Holmes’s eagerness to explain the logic step by step, and what rhetorical strategies (quoted or paraphrased) does Holmes use to underscore each clue’s significance?”

Mode	Average Score	Reason
E ² RAG (comb. extraction)	7.0000	Very thorough and accurate, captures Watson’s narration of Holmes’s eagerness and detailed rhetorical strategies with direct quotes and logical progression, closely matching the ground truth.
LightRAG hybrid (comb. extraction)	6.3333	Captures Holmes’s eagerness and rhetorical strategies well, including analogies, contrasts, and causal reasoning, with some direct quotes. However, it misses some of the specific step-by-step pacing and the theatrical, emphatic phrasing that the ground truth emphasizes, resulting in a somewhat less vivid depiction.
E ² RAG (vanilla)	6.0000	Detailed and faithful to the ground truth, includes Holmes’s systematic approach, direct quotes, and rhetorical strategies, capturing both Watson’s narration and Holmes’s methodical explanation.
LightRAG hybrid	6.0000	Provides a thorough analysis of Watson’s narration and Holmes’s rhetorical strategies with examples and direct quotes, closely aligning with the ground truth’s emphasis on step-by-step logic and Holmes’s delight in explaining.
GraphRAG drift	6.0000	Mentions Comprehensive and detailed, covers Watson’s narration and Holmes’s rhetorical strategies including analogies, rhetorical questions, and historical context, closely matching the ground truth.
LightRAG local	5.3333	Describes Watson’s narration and Holmes’s rhetorical strategies like analogies and highlighting contradictions, but lacks specific direct quotes and detailed step-by-step logic as in the ground truth.
vanilla RAG	4.3333	Captures Holmes’s eagerness and details several rhetorical strategies with some direct quotes and examples, but includes some inaccuracies and extraneous content not directly related to the Lauriston Gardens dinner scene.
vanilla HyDE	3.6667	Captures Holmes’s eagerness and rhetorical strategies with some direct quotes and analogies, but less focused on the step-by-step logic and specific clues as in the ground truth.
RQ-RAG	3.0000	Very Very brief and vague; mentions Holmes’s eagerness and rhetorical questions but lacks detail, examples, or direct quotes to support the answer.

whereas GraphRAG is noticeably slower.

Query time cost. Table 5 shows the average query time for each mode. GraphRAG drift takes the longest time. E²RAG (vanilla) takes marginally longer time compared to LightRAG hybrid mode, likely due to the time taken for searching the additional events on KG. The hypothetical response variants all took slightly longer time than their respective base form due to the additional step of generating the hypothetical response. In particular, the combined extraction variant took the longest time due to the extraction process. Nevertheless, the overall query time difference for LightRAG and E²RAG is insignificant with or without hypothetical response, and much shorter compared to all GraphRAG modes.

Token Cost Analysis. When KG construction phase, since event extraction is done separately from entity extraction, E²RAG has two times the number of API calls of LightRAG with same max token count C_{extract} , which doubles its worst-case token cost during preprocessing. On the other hand, as the author of LightRAG pointed out, GraphRAG is still has much higher token consumption when comes to KG construction due to massive, repeatedly generated community reports (Guo et al., 2024) (refer to Table 7 for the cost on preprocessing *The Phantom of the Opera*). During retrieval, to ensure a fair comparison, we set the max token C_{output} for the retrieved chunks to be the same for every KG-based RAG mode; therefore, they all have the same token cost. For every query, the hypothetical response step contributes to exactly one additional API call with the same C_{output} allowed

per API call, which doubles the output token count in the worst case.

Case study. In Table 6, we showcase the responses selected modes used for evaluation for a particular query from *The Hound of the Baskervilles* and the verdicts given by GPT-4.1-mini. The complete verdicts of all three LLM judges can be found in Appendix F.3. Overall, pairing with hypothetical response, E²RAG is able to retrieve contextually relevant chunks—rather than surface-level matches—yielding more thorough and accurate answers. The details on the extraction and retrieval of E²RAG (comb. extraction) mode can be found in Appendix F. Due to the space constraints, complete responses of other modes are provided in the supplementary materials.

6 Conclusion

E²RAG keeps every entity mention separate and anchors it to the exact events in which it appears, restoring the temporal and causal context that vanilla RAG and deduplicated KG variants fail to capture. On the new ChronoQA benchmark—designed specifically to test narrative, causal, and character-consistency reasoning—it delivers the best overall and category-specific scores while matching LightRAG’s preprocessing cost and latency. Because it builds on off-the-shelf extraction prompts and vector stores, requires no model fine-tuning, and pairs naturally with hypothetical response, E²RAG offers a drop-in upgrade for more faithful retrieval and a foundation for future work on RAG for complex question-answering tasks.

Limitations

The proposed E²RAG framework focuses on improving the response quality of the retrieval-augmented generation on documents that possess inherent temporal or causal structures—for instance novels, short stories, scripts, and other narrative-style texts in which characters evolve and events unfold in sequence. Outside of this scope, where information lacks such structures, E²RAG may not offer noticeable benefits compared to methods specifically designed for those contexts.

References

- Chenxin An, Jun Zhang, Ming Zhong, Lei Li, Shansan Gong, Yao Luo, Jingjing Xu, and Lingpeng Kong. 2024. Why does the effective context length of llms fall short? *arXiv preprint arXiv:2410.18745*.
- Anthropic. 2025. Claude 3.7 sonnet system card. <https://www.anthropic.com/claude-3-7-sonnet-system-card>. System card, accessed 2025-05-18.
- Patrice Béchard and Orlando Marquez Ayala. 2024. Reducing hallucination in structured outputs via retrieval-augmented generation. *arXiv preprint arXiv:2404.08189*.
- Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, et al. 2022. Improving language models by retrieving from trillions of tokens. In *International conference on machine learning*, pages 2206–2240. PMLR.
- Chi-Min Chan, Chunpu Xu, Ruibin Yuan, Hongyin Luo, Wei Xue, Yike Guo, and Jie Fu. 2024. Rq-rag: Learning to refine queries for retrieval augmented generation. *arXiv preprint arXiv:2404.00610*.
- Boyu Chen, Zirui Guo, Zidan Yang, Yuluo Chen, Junze Chen, Zhenghao Liu, Chuan Shi, and Cheng Yang. 2025. Pathrag: Pruning graph-based retrieval augmented generation with relational paths. *arXiv preprint arXiv:2502.14902*.
- Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. 2024. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*.
- Fiction.live. 2025. *Fiction.livebench: The first real-world long context benchmark for writers*. Accessed: 2025-05-10.
- Luyu Gao, Xueguang Ma, Jimmy Lin, and Jamie Callan. 2023. Precise zero-shot dense retrieval without relevance labels. In *Proceedings of the 61st Annual*

Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 1762–1777.

- Muhan Gao, TaiMing Lu, Kuai Yu, Adam Byerly, and Daniel Khashabi. 2024. Insights into llm long-context failures: When transformers know but don’t tell. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7611–7625.
- Zirui Guo, Lianghao Xia, Yanhua Yu, Tu Ao, and Chao Huang. 2024. *Lightrag: Simple and fast retrieval-augmented generation*. *arXiv preprint arXiv:2410.05779*. ArXiv:2410.05779 [cs.IR].
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Ziyan Jiang, Xueguang Ma, and Wenhua Chen. 2024. Longrag: Enhancing retrieval-augmented generation with long-context llms. *arXiv preprint arXiv:2406.15319*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474.
- Tianle Li, Ge Zhang, Quy Duc Do, Xiang Yue, and Wenhua Chen. 2024. Long-context llms struggle with long in-context learning. *arXiv preprint arXiv:2404.02060*.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. Lost in the middle: How language models use long contexts. *arXiv preprint arXiv:2307.03172*.
- Xiang Liu, Peijie Dong, Xuming Hu, and Xiaowen Chu. 2024. Longgenbench: Long-context generation benchmark. *arXiv preprint arXiv:2410.04199*.
- Qitan Lv, Jie Wang, Hanzhu Chen, Bin Li, Yongdong Zhang, and Feng Wu. 2024. Coarse-to-fine highlighting: Reducing knowledge hallucination in large language models. *arXiv preprint arXiv:2410.15116*.
- Xuanfan Ni, Hengyi Cai, Xiaochi Wei, Shuaiqiang Wang, Dawei Yin, and Piji Li. 2024. *Xl² bench: A benchmark for extremely long context understanding with long-range dependencies*. *arXiv preprint arXiv:2404.05446*.

- OpenAI. 2024. text-embedding-3-small: Openai embedding model. <https://platform.openai.com/docs/models/text-embedding-3-small>. Model release announced 25 Jan 2024.
- OpenAI. 2025. Introducing gpt-4.1 in the api. <https://openai.com/index/gpt-4-1>. Accessed: 20 May 2025.
- Boci Peng, Yun Zhu, Yongchao Liu, Xiaohe Bo, Haizhou Shi, Chuntao Hong, Yan Zhang, and Siliang Tang. 2024. Graph retrieval-augmented generation: A survey. *arXiv preprint arXiv:2408.08921*.
- Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. Retrieval augmentation reduces hallucination in conversation. *arXiv preprint arXiv:2104.07567*.
- Chandana Sree Mala, Gizem Gezici, and Fosca Gianotti. 2025. Hybrid retrieval for hallucination mitigation in large language models: A comparative analysis. *arXiv e-prints*, pages arXiv-2504.
- Minzheng Wang, Longze Chen, Cheng Fu, Shengyi Liao, Xinghua Zhang, Bingli Wu, Haiyang Yu, Nan Xu, Lei Zhang, Run Luo, et al. 2024. Leave no document behind: Benchmarking long-context llms with extended multi-doc qa. *arXiv preprint arXiv:2406.17419*.
- Shu Wang, Yixiang Fang, Yingli Zhou, Xilin Liu, and Yuchi Ma. 2025. Archrag: Attributed community-based hierarchical retrieval-augmented generation. *arXiv preprint arXiv:2502.09891*.

A More ChronoQA Details

Format and release. Each record is a JSON line with fields { "story_id", "story_title", "question_id", "category", "question", "ground_truth", "passages" }, where passages has subfields { "start_sentence", "end_sentence", "start_byte", "end_byte", "excerpt" } which pinpoint the relevant passage(s) from the story that provide the supporting evidence of the ground truth. The generation prompt can be found in Appendix D.1.

B Preprocessing Token Cost

Table 7: Token usage statistics for the three KG-RAG systems

System	Input tokens	Output tokens	Total tokens
LightRAG	343 437	112 840	456 277
E ² RAG	598 822	208 797	807 619
GraphRAG	880 804	440 401	1 321 205

C Hyperparameters

We implement E²RAG within LightRAG’s existing codebase. For all experiments, we standardize the chunk size, max token for backbone LLM generation, and chunks appended to be the same as the default setting as LightRAG. We experimented on changing their values but observed no noticeable benefit.

D Prompts

D.1 Data Generation Prompt

I want to test a few rag systems on their reasoning capabilities and the capability to pick up nuanced details. Use the document shown below, design queries for it, also provide the ground truth for each query. In addition, it should focus on testing the RAG system's causal consistency. for example, 1. it can ask the RAG system how would a certain character behavior in a given context/even from the document, and see if the answer is consistent with the character's personality/traits at that specific point in time (character can experience development so their response to the same thing can vary). 2. it can test if the RAG system confuse the event that has not happened but in the document with the query (the future events that has not happen should not be account for in the answer). Give the queries and ground truth in json format.the document is here:

DOCUMENT

Focus on Causal Consistency: Apart from character behavior over time and future event confusion, other causal consistency tests (e.g., testing how well it understands cause-and-effect relationships in the plot) can also be included. Query Difficulty: query should involve complex, nuanced reasoning/understanding of the document provided.

Format example:

```
{
  "queries": [
    {
      "query": "During Harry's first night at Hogwarts, the Gryffindor students climb the moving staircases. How does the text describe the corridors and stairways' magical behavior, and what is the immediate impact on Harry's sense of direction?",
      "ground_truth": "The staircases sometimes change direction, doors can vanish or move, and some require a password or a specific tickle of a doorknob to open. This constant shifting confuses new students like Harry, making it easy to get lost early on.",
      "start_sentence": "The staircases at Hogwarts are famous for moving unexpectedly, often depositing unwary students on entirely different floors than intended.",
      "end_sentence": "This enchantment leaves first-years such as Harry feeling hopelessly lost during their first nights in the castle.",
      "type": "Causal Consistency"
    }
  ]
}
```

Rules for the `start_sentence` and `end_sentence` strings:

- They must be *identical substrings* of the document (case-sensitive, byte-for-byte).

- Preserve every original character: spaces, line-breaks, hyphens, quotation marks, etc.
- If the document contains line breaks, represent them in JSON as the two-character sequence.
- Do not add, delete, or normalise any characters—copy-paste only.
- The passage between the two sentences must support the ground-truth answer you give.

It is also important to note that do not explicitly disclose the title/chapter/section number from which the context of the question is used. Simply described the related event and the characters involved to make the context clear.

D.2 Responses Evaluation Prompt

You are an expert evaluator of retrieval-augmented generation (RAG) answers.

Scoring rubric (10-point scale):

- 10 - Matches ground truth exactly or with faithful paraphrase.
- 7 - Mostly correct; minor omissions or wording differences.
- 5 - Partially correct; major missing points or inaccuracies.
- 3 - Mostly incorrect; small overlap.
- 1 - Off-topic or hallucinated.

Return **only** a valid JSON array, no markdown fences, in this exact shape:

```
[
  {"mode": "mode_name", "reason": "short rationale", "score": 9},
  ...
]
```

If you cannot produce the JSON array, return an object like:
{"error": "description"}.

D.3 Entity, Event Extraction Prompt

---Role---

You are a helpful assistant tasked with identifying entities and events in the user's query.

---Goal---

Given the query, list both entities and events. Entities are people, places, organizations, or objects mentioned in the query, while events are actions, occurrences, or happenings that take place.

---Instructions---

- Output the entities and events in JSON format.
- The JSON should have two keys:
 - "entities" for people, places, organizations, or objects.
 - "events" for actions, occurrences, or happenings.

#####

-Examples-

#####

Example 1:

Query: "How did Napoleon's invasion of Russia affect his empire's strength?"

#####

Output:

```
{
  "entities": ["Napoleon", "Russia", "Napoleon's empire"],
  "events": ["invasion of Russia", "empire's decline"]
}
```

#####

Example 2:

Query: "What role did MIT scientists play in the Manhattan Project?"

#####

Output:

```
{
  "entities": ["MIT", "MIT scientists", "Manhattan Project"],
  "events": ["scientific research", "atomic bomb development"]
}
```



```
#####
Example 3:

Query: "How did the Industrial Revolution change London's population?"
#####
Output:
{{
  "entities": ["London", "London's population", "Industrial Revolution"],
  "events": ["population growth", "urbanization", "industrial development"]
}}
#####
-Real Data-
#####
Query: {query}
#####
Output:

"""

PROMPTS["naive_rag_response"] = """---Role---

You are a helpful assistant responding to questions about documents provided.

---Goal---

Generate a response of the target length and format that responds to the user's question,
summarizing all information in the input data tables appropriate for the response length and format,
and incorporating any relevant general knowledge.
If you don't know the answer, just say so. Do not make anything up.
Do not include information where the supporting evidence for it is not provided.

---Target response length and format---

{response_type}

---Documents---

{content_data}

Add sections and commentary to the response as appropriate for the length and format. Style the
response in markdown.
```

E Entity-Event KG Visualization

769

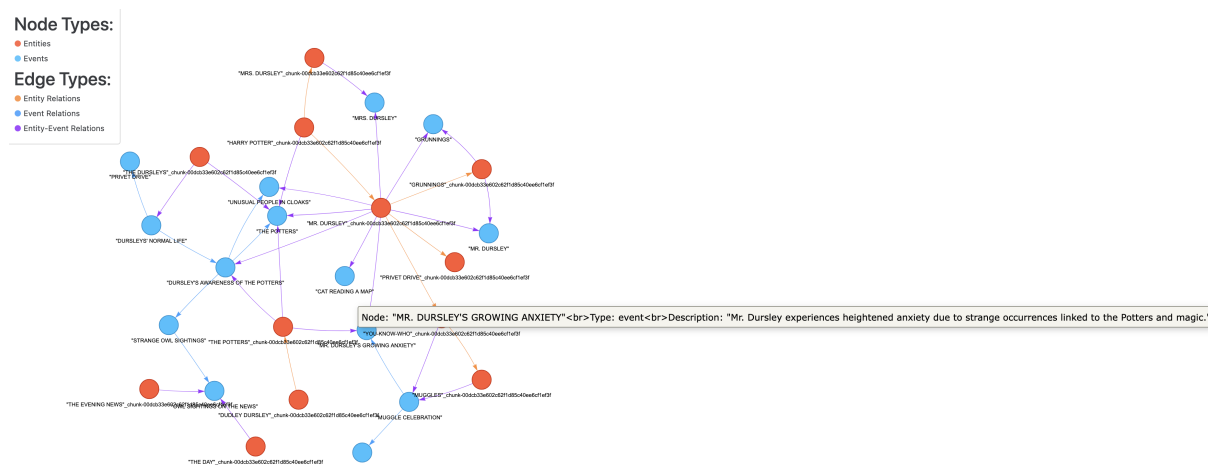


Figure 4: A snapshot of a part of the Entity-Event KG of *Harry Potter and the Sorcerer's Stone*.

F More Details on Case Study on Mode: E²RAG (comb. extraction)

F.1 Hypothetical Response

In the dinner scene from "The Hound of the Baskervilles," Watson's narration captures Holmes's fervent enthusiasm for elucidating his reasoning behind the Lauriston Gardens mystery. As they sit at the table, Watson observes Holmes's animated demeanor, noting how "his eyes shone with a keen light" and how he "leaned forward with an eagerness that betrayed his excitement." This vivid description not only illustrates Holmes's passion for deduction but also emphasizes his desire to engage Watson in the unraveling of the case, making him feel like a partner in the detective's intellectual adventure. Holmes employs various rhetorical strategies to accentuate the significance of each clue as he meticulously lays out his thought process. For instance, he often uses analogies to simplify complex ideas, comparing the clues to pieces of a puzzle that, when assembled correctly, reveal the bigger picture. He states, "The world is full of obvious things which nobody by any chance ever observes," underscoring how critical attention to detail is in investigative work. This rhetorical device serves to remind Watson and the reader of the importance of keen observation in deciphering the mystery. Moreover, Holmes frequently employs the technique of building suspense through his methodical breakdown of the evidence. He carefully dissects each piece of information, pausing for effect before revealing its implications. By saying, "And yet, it is the most obvious thing in the world," he not only highlights the clarity of his deductions but also invites the audience to share in the thrill of discovery. This step-by-step approach not only showcases Holmes's logical prowess but also enhances Watson's narrative by creating a sense of shared enlightenment, as they both traverse the landscape of clues together. Ultimately, Watson's narration, rich in detail and observation, amplifies Holmes's eagerness to teach and elucidate, while Holmes's rhetorical strategies serve to elevate the weight of each clue, transforming the dinner scene into a dynamic exploration of deduction and insight.

F.2 Extracted Data

Extracted entity keywords:

```
['Watson', 'dinner scene', 'The Hound of the Baskervilles', 'Lauriston Gardens', 'Holmes']
```

Extracted event keywords:

```
["Holmes's enthusiasm", 'dinner scene', 'discussion of the Lauriston Gardens mystery', 'breakdown of evidence', 'building suspense', "underscoring clues' significance", "narration of Holmes's reasoning", 'detailing reasoning', 'highlighting eagerness', 'explaining logic', 'using rhetorical strategies', 'shared enlightenment', 'deduction process']
```

Retrieved nodes data:

```
[{'type': 'PERSON', 'description': 'Holmes is a character portrayed as a detective who engages in solving a mystery surr...and expresses confidence in his theories.', 'source_id': 'chunk-e16ab5c062904b8574746080e70db172', 'node_type': 'entity', 'display_color': '#FF5733', 'entity_name': 'HOLMES', 'chunk-e16ab5c062904b8574746080e70db172', 'rank': 9}, {'type': 'EVENT', 'description': 'Holmes elaborates on his strategy to catch the assassin while navigating the complexities of the situation.', 'source_id': 'chunk-fbfff6479d4cfc0843851b4ff06a86927', 'node_type': 'event', 'display_color': '#33C1FF', 'entity_name': 'HOLMES'S STRATEGY DISCUSSION', 'rank': 2}]
```

Retrieved edge data:

```
[{'src_tgt': (...), 'rank': 37, 'weight': 8.0, 'description': 'I engages in dialogue with Holmes, seeking understanding of the detective's methodologies and deductions.', 'keywords': 'narration, learning', 'source_id': 'chunk-9d367c150b511e67b26c6a5701e7fcb9', 'edge_type': 'entity_relation', 'display_color': '#FF8C33'}, {'src_tgt': (...), 'rank': 25, 'edge_type': 'entity_event_relation', 'description': 'Entity Holmes is a skilled detective known for his sharp powers of observation and ...s wheels as a crucial piece of evidence.', 'weight': 1.0, 'keywords': 'ENTITY_EVENT_RELATION', 'source_id': 'chunk-9d367c150b511e67b26c6a5701e7fcb9<SEP>chunk-9d367c150b511e67b26c6a5701e7fcb9', 'display_color': '#9933FF'}]
```

Retrieved context data (relevant chunks):

\n-----Sources-----\n```\n\nid,content\r\n0,"and that the depravity of the victim was no condonement in the eyes of the law.\n\nThe more I thought of it the more extraordinary did my companion's hypothesis, that the man had been poisoned, appear. I remembered how he had sniffed his lips, and had no doubt that he had detected something which had given rise to the idea. Then, again, if not poison, what had caused the man's death, since there was neither wound nor marks of strangulation? But, on the other hand, whose blood was that which lay so thickly upon the floor? There were no signs of a struggle, nor had the victim any weapon with which he might have wounded an antagonist. As long as all these questions were unsolved, I felt that sleep would be no easy matter, either for Holmes or myself. His quiet self-confident manner convinced me that he had already formed a theory which explained all the facts, though what it was I could not for an instant conjecture. ...

Final response:

778

In the dinner scene where Sherlock Holmes discusses the Lauriston Gardens mystery with Dr. Watson, Watson's narration reveals Holmes's enthusiastic and methodical approach to detective work. This moment underscores not only Holmes's intellectual prowess but also his eagerness to engage Watson in the reasoning process that leads to his deductions. Watson's observations during this discussion reflect his genuine admiration for Holmes's analytical mind, which illuminates the depth of the detective's insight and reasoning.

Watson's Narrative Perspective

Watson, as the narrator, captures the intensity of Holmes's engagement with the case. He notes how Holmes displays a "quiet self-confident manner," suggesting that Holmes is not only certain of his deductions but also eager to articulate the thought process that led him there. For instance, Watson recounts moments where Holmes transitions on topics with excitement, indicating how much he relishes sharing the intricacies of his method.

The plot thickens, and he exclaims at key junctures, revealing his thrill at unraveling the mystery. This narrative style enhances Holmes's character as not just a brilliant detective but also a passionate educator, inviting Watson and, by extension, the reader into his deductive world.

Rhetorical Strategies Used by Holmes

Holmes employs several rhetorical strategies to underscore the significance of each clue in his explanation. These strategies include:

- Logical Sequencing:** Holmes emphasizes the importance of reasoning backwards and forwards, which he articulately illustrates when he examines the physical evidence. For example, he reasons from the cab's ruts and horse's hoof prints left on the wet soil to deduce that two men must have visited the house. Here, he insists that "it is a capital mistake to theorize before you have all the evidence," highlighting his disciplined approach to deduction that relies on empirical observations rather than mere speculation.
- Detailed Observation:** Holmes's acute observations serve as pivotal points in his reasoning. He recognizes that the "very first thing" he observed was the impressions left by the cab wheels. By meticulously detailing this observation, he conveys to Watson and the audience the fundamental role of minute details in solving a case: "There is no branch of detective science which is so important and so much neglected as the art of tracing footsteps."
- Analytical Clarity:** As he breaks down how he calculated the height and nature of the suspects based on their footprints and style of dress, Holmes uses clear, methodical language. He moves from discussing the physical evidence to the psychological implications, illustrating how the spatial and temporal context plays a critical role in understanding the crime. For example, he reflects on the footprints and concludes, "It is simple enough I could read all that in the dust."
- Connecting to Broader Concepts:** Holmes often draws on broader concepts of criminality and human behavior, thereby situating his deductions within a larger framework of detective work. He references historical cases and theories that bolster his arguments, explaining his reasoning not just as a technical endeavor, but as a complex interplay of human motivations and actions.

Conclusion

Holmes's eagerness to share his deductive reasoning exemplifies the detective's character as someone deeply engaged in the pursuit of knowledge and understanding. Through Watson's narration, readers see an interplay of admiration, curiosity, and insight as Holmes articulates the significance of each clue. The combination of logical sequencing, detailed observation, analytical clarity, and broader contextual reasoning allows Holmes to construct a compelling narrative that not only uncovers the truth but also elevates the art of detection itself. This dinner scene, rich with multi-layered reasoning, highlights the intricacies of Holmes's mind, drawing readers closer to the fascinating world of deduction.

F.3 Complete Verdicts from All Three LLM Judges

Table 8: GPT-4.1-mini’s verdicts (part 1) of the responses to “Consider the dinner scene in which Holmes details his reasoning about the Lauriston Gardens mystery. How does Watson’s narration highlight Holmes’s eagerness to explain the logic step by step, and what rhetorical strategies (quoted or paraphrased) does Holmes use to underscore each clue’s significance?”

Mode	Average Score	Reason
E ² RAG (comb. extraction)	7.0000	Very thorough and accurate, captures Watson’s narration of Holmes’s eagerness and detailed rhetorical strategies with direct quotes and logical progression, closely matching the ground truth.
E ² RAG (comb. embedding)	7.0000	Comprehensive and well-aligned with the ground truth, includes detailed analysis of Watson’s narration, Holmes’s eagerness, and rhetorical strategies with direct quotes.
E ² RAG (hyp. embedding)	7.0000	Detailed and faithful, includes Holmes’s eagerness, rhetorical strategies, direct quotes, and Watson’s narration, effectively capturing the essence of the ground truth.
LightRAG hybrid (hyp. extraction)	6.6667	Highly detailed and faithful to the ground truth, includes Holmes’s systematic reasoning, direct quotes, rhetorical questions, and Watson’s narration highlighting Holmes’s eagerness.
LightRAG hybrid (comb. extraction)	6.3333	Captures Holmes’s eagerness and rhetorical strategies well, including analogies, contrasts, and causal reasoning, with some direct quotes. However, it misses some of the specific step-by-step pacing and the theatrical, emphatic phrasing that the ground truth emphasizes, resulting in a somewhat less vivid depiction.
E ² RAG (hyp. extraction)	6.3333	Provides a detailed and faithful account of Watson’s narration and Holmes’s eagerness, including direct quotes and rhetorical strategies such as deductive reasoning, metaphor, and analogy. However, it lacks some of the exact emphatic openings, instructive questions, and dismissive flourishes explicitly cited in the ground truth, making it slightly less precise.
LightRAG hybrid (comb. embedding)	6.3333	Providing a solid overview of Holmes’s eagerness and rhetorical strategies with direct quotes and logical reasoning. However, it does not fully capture the breathless pacing and specific emphatic or dismissive phrases that characterize the ground truth.
LightRAG hybrid (hyp. embedding)	6.3333	Nearly identical to hybrid_combined_extraction in content and style, with good coverage of Watson’s narration and Holmes’s rhetorical strategies but lacking the full range of direct quotes and the theatrical emphasis present in the ground truth.
E ² RAG (vanilla)	6.0000	Detailed and faithful to the ground truth, includes Holmes’s systematic approach, direct quotes, and rhetorical strategies, capturing both Watson’s narration and Holmes’s methodical explanation.
GraphRAG global	6.0000	Addresses Strong answer highlighting Holmes’s eagerness and rhetorical strategies like analogies and emphatic language, with direct quotes and good alignment to the ground truth.
LightRAG hybrid	6.0000	Provides a thorough analysis of Watson’s narration and Holmes’s rhetorical strategies with examples and direct quotes, closely aligning with the ground truth’s emphasis on step-by-step logic and Holmes’s delight in explaining.
GraphRAG drift	6.0000	Mentions Comprehensive and detailed, covers Watson’s narration and Holmes’s rhetorical strategies including analogies, rhetorical questions, and historical context, closely matching the ground truth.
GraphRAG local	5.6667	Good summary of Watson’s narration and Holmes’s rhetorical strategies including logical reasoning and analogies, with some direct quotes, but less detailed than the ground truth.
LightRAG global	5.6667	Effectively highlights Holmes’s eagerness and rhetorical strategies such as analogies and rhetorical questions with some direct quotes, closely matching the ground truth’s focus on Holmes’s methodical explanation.

Table 8: GPT-4.1-minis verdicts (continued)

Mode	Average Score	Reason
LightRAG local	5.3333	Describes Watson’s narration and Holmes’s rhetorical strategies like analogies and highlighting contradictions, but lacks specific direct quotes and detailed step-by-step logic as in the ground truth.
vanilla RAG	4.3333	Captures Holmes’s eagerness and details several rhetorical strategies with some direct quotes and examples, but includes some inaccuracies and extraneous content not directly related to the Lauriston Gardens dinner scene.
vanilla HyDE	3.6667	Captures Holmes’s eagerness and rhetorical strategies with some direct quotes and analogies, but less focused on the step-by-step logic and specific clues as in the ground truth.
RQ-RAG	3.0000	Very Very brief and vague; mentions Holmes’s eagerness and rhetorical questions but lacks detail, examples, or direct quotes to support the answer.

Table 9: Claude-3.7-Sonnet’s verdicts (part 1) of the responses to *"Consider the dinner scene in which Holmes details his reasoning about the Lauriston Gardens mystery. How does Watson’s narration highlight Holmes’s eagerness to explain the logic step by step, and what rhetorical strategies (quoted or paraphrased) does Holmes use to underscore each clue’s significance?"*

Mode	Average Score	Reason
E ² RAG (comb. extraction)	7.0000	Mentions footprints and cab marks, and discusses Holmes’s eagerness to explain his reasoning. However, it doesn’t capture the specific rhetorical pattern of short declarations followed by challenges to Watson.
E ² RAG (comb. embedding)	7.0000	Discusses Holmes’s eagerness and mentions footprints and stride length, but doesn’t fully capture the rhetorical pattern of short declarations followed by challenges to Watson described in the ground truth.
E ² RAG (hyp. embedding)	7.0000	Mentions cab marks and footprints, and discusses Holmes’s eagerness to explain his reasoning. However, it doesn’t capture the specific rhetorical pattern of short declarations followed by challenges to Watson.
LightRAG hybrid (hyp. extraction)	6.6667	Mentions key elements like cab wheel ruts and stride length, and discusses Holmes’s eagerness to explain his reasoning. Includes some specific rhetorical strategies but doesn’t fully capture the rapid, orderly exposition and the pattern of declarations followed by challenges to Watson described in the ground truth.
LightRAG hybrid (comb. extraction)	6.3333	Accurately mentions cab tracks and Holmes’s eagerness to explain his reasoning step by step. References some rhetorical strategies like analogies and contrasts, but doesn’t fully capture all the specific elements (fingernails, puddle width) or the rhetorical pattern of emphatic openings and dismissive flourishes in the ground truth.
E ² RAG (hyp. extraction)	6.3333	Mentions cab marks and footprints but doesn’t specifically address Holmes’s eagerness to explain step by step or his rhetorical pattern of making declarations followed by challenging Watson.
LightRAG hybrid (comb. embedding)	6.3333	Mentions cab tracks and Holmes’s eagerness to explain his reasoning. Includes some rhetorical strategies but doesn’t fully capture all the specific elements or the rhetorical pattern described in the ground truth.
LightRAG hybrid (hyp. embedding)	6.3333	Mentions cab tracks and Holmes’s eagerness to explain his reasoning. Includes some rhetorical strategies but doesn’t fully capture all the specific elements or the rhetorical pattern described in the ground truth.
E ² RAG (vanilla)	6.0000	Mentions some key elements like footprints and cab marks but doesn’t specifically address Holmes’s eagerness to explain step by step or his rhetorical pattern of making declarations followed by challenging Watson.
GraphRAG global	6.0000	Addresses Holmes’s eagerness to explain his reasoning but lacks specific references to footprints, stride lengths, and cab marks. The rhetorical strategies mentioned are not specific to the text.
LightRAG hybrid	6.0000	Provides a detailed analysis of Holmes’s eagerness and rhetorical strategies but doesn’t specifically mention the footprints, stride lengths, and cab marks that are central to the ground truth. The rhetorical strategies are more generalized than specific.
GraphRAG drift	6.0000	Mentions Holmes’s eagerness and some rhetorical strategies but lacks specific references to footprints, stride lengths, and cab marks. The analysis is somewhat generic rather than text-specific.

Table 9: Claude-3.7-Sonnet’s verdicts (continued)

Mode	Average Score	Reason
GraphRAG local	5.6667	Discusses Holmes’s eagerness and some rhetorical strategies but lacks specific references to the footprints, stride lengths, and cab marks mentioned in the ground truth. The analysis is somewhat generic.
LightRAG local	5.3333	Discusses Holmes’s eagerness and methodical reasoning but lacks specific references to footprints, stride lengths, and cab marks mentioned in the ground truth. The rhetorical strategies are generalized rather than specific to the text.
LightRAG global	5.6667	Addresses Holmes’s eagerness to explain his logic but lacks specific references to the footprints, stride lengths, and cab marks. Mentions rhetorical questions but doesn’t capture Holmes’s pattern of short declarations followed by challenges to Watson.
vanilla RAG	4.3333	Contains some relevant elements about Holmes’s analytical reasoning but misquotes Holmes and includes fabricated dialogue. The analysis lacks specific references to Watson’s narration of Holmes’s eagerness and the step-by-step logic mentioned in the ground truth.
vanilla HyDE	3.6667	Discusses a different story (‘The Hound of the Baskervilles’) and doesn’t mention the specific elements from the ground truth like footprints, stride lengths, and cab marks.
RQ-RAG	3.0000	Very brief and generic. While it mentions Holmes’s eagerness and rhetorical strategies, it lacks specific details about footprints, stride lengths, and cab marks. The answer is too vague and underdeveloped.

Table 10: GPT-4o’s verdicts (part 1) of the responses to *“Consider the dinner scene in which Holmes details his reasoning about the Lauriston Gardens mystery. How does Watson’s narration highlight Holmes’s eagerness to explain the logic step by step, and what rhetorical strategies (quoted or paraphrased) does Holmes use to underscore each clue’s significance?”*

Mode	Average Score	Reason
E ² RAG (comb. extraction)	7.0000	Partially correct; discusses Holmes’s eagerness and some rhetorical strategies but lacks specific examples.
E ² RAG (comb. embedding)	7.0000	Partially correct; discusses Holmes’s eagerness and some rhetorical strategies but lacks specific examples.
E ² RAG (hyp. embedding)	7.0000	Partially correct; discusses Holmes’s eagerness and some rhetorical strategies but lacks specific examples.
LightRAG hybrid (hyp. extraction)	6.6667	Partially correct; discusses Holmes’s eagerness and some rhetorical strategies but lacks specific examples.
LightRAG hybrid (comb. extraction)	6.3333	Partially correct; discusses Holmes’s eagerness and some rhetorical strategies but lacks specific examples.
E ² RAG (hyp. extraction)	6.3333	Partially correct; discusses Holmes’s eagerness and some rhetorical strategies but lacks specific examples.
LightRAG hybrid (comb. embedding)	6.3333	Partially correct; discusses Holmes’s eagerness and some rhetorical strategies but lacks specific examples.
LightRAG hybrid (hyp. embedding)	6.3333	Partially correct; discusses Holmes’s eagerness and some rhetorical strategies but lacks specific examples.
E ² RAG (vanilla)	6.0000	Partially correct; discusses Holmes’s enthusiasm and some rhetorical strategies but lacks specific examples.
GraphRAG global	6.0000	Addresses Partially correct; discusses Holmes’s eagerness and some rhetorical strategies but lacks specific examples.
LightRAG hybrid	6.0000	Partially correct; mentions Holmes’s eagerness and some rhetorical strategies but lacks specific examples.
GraphRAG drift	6.0000	Mentions Partially correct; discusses Holmes’s eagerness and some rhetorical strategies but lacks specific examples.
GraphRAG local	5.6667	Partially correct; discusses Holmes’s eagerness and some rhetorical strategies but lacks specific examples.
LightRAG local	5.3333	Partially correct; mentions Holmes’s enthusiasm and some rhetorical strategies but lacks specific examples.
LightRAG global	5.6667	Partially correct; discusses Holmes’s eagerness and rhetorical strategies but lacks specific examples from the scene.
vanilla RAG	4.3333	Mostly incorrect; lacks specific details about Holmes’s rhetorical strategies and Watson’s narration.
vanilla HyDE	3.6667	Off-topic; discusses ‘The Hound of the Baskervilles’ instead of Lauriston Gardens.
RQ-RAG	3.0000	Very Mostly incorrect; brief mention of rhetorical strategies without specific examples or context.

G Additional Results

Table 11 to Table 13 show the ranking results of each individual LLM evaluator. It can be observed each mode’s ranking remains relatively stable. On the other hand, different evaluators have different standards, as the scores given by Claude-3-7-sonnet are generally low while the scores given by GPT-4.1-mini are generally high across the board. Table 14 to Table 21 show the ranking results for the questions based on category, averaged across all three evaluator LLMs.

Table 11: Average scores for Claude-3-7-sonnet

Rank	Mode	Avg Score	Total
1	E ² RAG (comb. extraction)	5.356 9	2 657
2	E ² RAG (comb. embedding)	5.266 1	2 612
3	E ² RAG (hyp. extraction)	5.137 1	2 548
4	LightRAG hybrid	5.135 1	2 547
5	E ² RAG (hyp. embedding)	5.106 9	2 533
6	vanilla RAG	5.024 2	2 492
7	GraphRAG drift	4.975 8	2 468
8	GraphRAG local	4.965 7	2 463
9	E ² RAG (vanilla)	4.927 4	2 444
10	LightRAG global	4.828 6	2 395
11	LightRAG local	4.715 7	2 339
12	GraphRAG global	4.598 8	2 281
13	LightRAG hybrid (comb. extraction)	4.360 9	2 163
14	LightRAG hybrid (hyp. embedding)	4.354 8	2 160
15	LightRAG hybrid (comb. embedding)	4.350 8	2 158
16	vanilla HyDE	4.320 6	2 143
17	LightRAG hybrid (hyp. extraction)	4.312 5	2 139
18	RQ-RAG	2.647 2	1 313

Table 12: Average scores for GPT-4o

Rank	Mode	Avg Score	Total
1	LightRAG hybrid	6.917 3	3 431
2	E ² RAG (comb. extraction)	6.856 9	3 401
3	GraphRAG local	6.808 5	3 377
4	E ² RAG (comb. embedding)	6.808 5	3 377
5	E ² RAG (hyp. extraction)	6.782 3	3 364
6	GraphRAG drift	6.774 2	3 360
7	LightRAG local	6.689 5	3 318
8	E ² RAG (hyp. embedding)	6.689 5	3 318
9	E ² RAG (vanilla)	6.633 1	3 290
10	GraphRAG global	6.596 8	3 272
11	LightRAG global	6.568 5	3 258
12	vanilla RAG	6.532 3	3 240
13	LightRAG hybrid (comb. embedding)	6.360 9	3 155
14	LightRAG hybrid (comb. extraction)	6.344 8	3 147
15	LightRAG hybrid (hyp. extraction)	6.340 7	3 145
16	LightRAG hybrid (hyp. embedding)	6.334 7	3 142
17	vanilla HyDE	6.258 1	3 104
18	RQ-RAG	3.987 9	1 978

Table 13: Average scores for GPT-4.1-mini

Rank	Mode	Avg Score	Total
1	E²RAG (comb. extraction)	9.163 3	4 545
2	E²RAG (comb. embedding)	9.141 1	4 534
3	E²RAG (hyp. extraction)	9.030 2	4 479
4	E²RAG (hyp. embedding)	9.022 2	4 475
5	GraphRAG drift	8.711 7	4 321
6	GraphRAG local	8.625 0	4 278
7	LightRAG hybrid	8.588 7	4 260
8	E²RAG (vanilla)	8.564 5	4 248
9	vanilla HyDE	8.487 9	4 210
10	GraphRAG global	8.330 6	4 132
11	vanilla RAG	8.250 0	4 092
12	LightRAG local	8.244 0	4 089
13	LightRAG global	7.977 8	3 957
14	LightRAG hybrid (comb. embedding)	6.379 0	3 164
15	LightRAG hybrid (comb. extraction)	6.338 7	3 144
16	LightRAG hybrid (hyp. embedding)	6.324 6	3 137
17	LightRAG hybrid (hyp. extraction)	6.312 5	3 131
18	RQ-RAG	3.907 3	1 938

Table 14: Causal Consistency category: average scores

Rank	Mode	Avg Score	Total
1	E²RAG (comb. extraction)	7.059 0	2 033
2	E²RAG (comb. embedding)	6.989 6	2 013
3	E²RAG (hyp. extraction)	6.899 3	1 987
4	E²RAG (hyp. embedding)	6.840 3	1 970
5	E²RAG (vanilla)	6.767 4	1 949
6	GraphRAG drift	6.753 5	1 940
7	GraphRAG local	6.732 6	1 933
8	vanilla RAG	6.569 4	1 892
9	LightRAG hybrid	6.548 6	1 886
10	LightRAG local	6.420 1	1 859
11	GraphRAG global	6.395 8	1 842
12	LightRAG global	6.267 4	1 805
13	vanilla HyDE	5.774 3	1 663
14	LightRAG hybrid (comb. extraction)	4.930 6	1 411
15	LightRAG hybrid (hyp. embedding)	4.902 8	1 413
16	LightRAG hybrid (comb. embedding)	4.902 8	1 413
17	LightRAG hybrid (hyp. extraction)	4.899 3	1 411
18	RQ-RAG	3.145 8	907

Table 15: Character Consistency category: average scores

Rank	Mode	Avg Score	Total
1	E²RAG (comb. extraction)	7.279 9	4 994
2	E²RAG (comb. embedding)	7.239 1	4 966
3	E²RAG (hyp. extraction)	7.157 4	4 910
4	E²RAG (hyp. embedding)	7.142 9	4 900
5	LightRAG hybrid	7.121 0	4 885
6	GraphRAG drift	6.975 2	4 783
7	GraphRAG local	6.962 1	4 772
8	E²RAG (vanilla)	6.909 6	4 743
9	vanilla RAG	6.819 2	4 687
10	LightRAG local	6.797 4	4 664
11	GraphRAG global	6.702 6	4 595
12	vanilla HyDE	6.701 2	4 594
13	LightRAG global	6.644 3	4 550
14	LightRAG hybrid (comb. embedding)	6.105 0	4 195
15	LightRAG hybrid (comb. extraction)	6.080 2	4 178
16	LightRAG hybrid (hyp. embedding)	6.070 0	4 171
17	LightRAG hybrid (hyp. extraction)	6.045 2	4 153
18	RQ-RAG	3.718 7	2 552

Table 16: Emotional and Psychological category: average scores

Rank	Mode	Avg Score	Total
1	E²RAG (hyp. extraction)	8.333 3	50
2	LightRAG local	8.166 7	49
3	E²RAG (comb. extraction)	8.166 7	49
4	E²RAG (hyp. embedding)	8.000 0	48
5	vanilla RAG	8.000 0	48
6	LightRAG hybrid	7.833 3	47
7	E²RAG (comb. embedding)	7.833 3	47
8	E²RAG (vanilla)	7.666 7	46
9	LightRAG hybrid (hyp. embedding)	7.166 7	43
10	LightRAG hybrid (comb. extraction)	7.166 7	43
11	LightRAG global	7.166 7	43
12	LightRAG hybrid (hyp. extraction)	7.166 7	43
13	LightRAG hybrid (comb. embedding)	7.166 7	43
14	GraphRAG drift	7.000 0	42
15	vanilla HyDE	6.666 7	40
16	GraphRAG local	6.333 3	38
17	GraphRAG global	6.333 3	38
18	RQ-RAG	6.000 0	36

Table 17: Narrative and Plot Structure category: average scores

Rank	Mode	Avg Score	Total
1	E²RAG (hyp. extraction)	6.978 5	649
2	E²RAG (comb. extraction)	6.946 2	646
3	GraphRAG drift	6.881 7	640
4	E²RAG (comb. embedding)	6.828 0	635
5	E²RAG (hyp. embedding)	6.720 4	625
6	LightRAG hybrid	6.688 2	622
7	GraphRAG local	6.677 4	621
8	E²RAG (vanilla)	6.612 9	615
9	LightRAG local	6.462 4	601
10	LightRAG global	6.419 4	597
11	vanilla RAG	6.301 1	586
12	GraphRAG global	6.290 3	585
13	vanilla HyDE	6.075 3	565
14	LightRAG hybrid (hyp. embedding)	5.720 4	532
15	LightRAG hybrid (comb. extraction)	5.720 4	532
16	LightRAG hybrid (comb. embedding)	5.720 4	532
17	LightRAG hybrid (hyp. extraction)	5.709 7	531
18	RQ-RAG	3.612 9	336

Table 18: Setting, Environment and Atmosphere category: average scores

Rank	Mode	Avg Score	Total
1	E²RAG (comb. extraction)	6.479 5	473
2	E²RAG (comb. embedding)	6.287 7	459
3	LightRAG hybrid	6.260 3	457
4	E²RAG (hyp. extraction)	6.260 3	457
5	vanilla RAG	6.150 7	449
6	E²RAG (hyp. embedding)	6.123 3	447
7	E²RAG (vanilla)	6.000 0	438
8	GraphRAG drift	5.945 2	434
9	LightRAG local	5.835 6	426
10	vanilla HyDE	5.753 4	420
11	GraphRAG local	5.643 8	412
12	GraphRAG global	5.589 0	409
13	LightRAG global	5.547 9	404
14	LightRAG hybrid (hyp. embedding)	4.739 7	346
15	LightRAG hybrid (comb. embedding)	4.739 7	346
16	LightRAG hybrid (comb. extraction)	4.726 0	345
17	LightRAG hybrid (hyp. extraction)	4.726 0	345
18	RQ-RAG	3.013 7	220

Table 19: Social, Cultural and Political category: average scores

Rank	Mode	Avg Score	Total
1	E ² RAG (comb. extraction)	7.075 8	467
2	E ² RAG (comb. embedding)	6.969 7	460
3	LightRAG hybrid	6.924 2	457
4	E ² RAG (hyp. extraction)	6.863 6	453
5	E ² RAG (hyp. embedding)	6.833 3	451
6	GraphRAG local	6.636 4	438
7	LightRAG global	6.621 2	437
8	GraphRAG drift	6.590 9	435
9	GraphRAG global	6.560 6	433
10	vanilla HyDE	6.530 3	431
11	vanilla RAG	6.484 8	428
12	E ² RAG (vanilla)	6.454 5	426
13	LightRAG local	6.333 3	418
14	LightRAG hybrid (comb. embedding)	5.878 8	389
15	LightRAG hybrid (hyp. embedding)	5.818 2	384
16	LightRAG hybrid (comb. extraction)	5.803 0	383
17	LightRAG hybrid (hyp. extraction)	5.787 9	382
18	RQ-RAG	3.212 1	212

Table 20: Symbolism, Imagery and Motifs category: average scores

Rank	Mode	Avg Score	Total
1	E ² RAG (comb. embedding)	6.839 3	1 149
2	E ² RAG (comb. extraction)	6.821 4	1 146
3	E ² RAG (hyp. extraction)	6.684 5	1 123
4	GraphRAG local	6.654 8	1 118
5	E ² RAG (hyp. embedding)	6.648 8	1 117
6	LightRAG hybrid	6.619 0	1 112
7	GraphRAG drift	6.500 0	1 092
8	E ² RAG (vanilla)	6.250 0	1 050
9	GraphRAG global	6.202 4	1 042
10	LightRAG global	6.166 7	1 037
11	vanilla RAG	6.131 0	1 020
12	LightRAG local	6.113 1	1 022
13	vanilla HyDE	6.071 4	1 020
14	LightRAG hybrid (comb. embedding)	5.339 3	898
15	LightRAG hybrid (hyp. embedding)	5.321 4	894
16	LightRAG hybrid (comb. extraction)	5.321 4	894
17	LightRAG hybrid (hyp. extraction)	5.315 5	894
18	RQ-RAG	3.446 4	580

Table 21: Thematic, Philosophical and Moral category: average scores

Rank	Mode	Avg Score	Total
1	E ² RAG (comb. extraction)	7.361 1	795
2	E ² RAG (comb. embedding)	7.351 9	794
3	GraphRAG local	7.185 2	776
4	GraphRAG drift	7.185 2	776
5	LightRAG hybrid	7.148 1	772
6	E ² RAG (hyp. embedding)	7.111 1	768
7	E ² RAG (hyp. extraction)	7.055 6	762
8	GraphRAG global	6.842 6	738
9	LightRAG global	6.750 0	729
10	vanilla HyDE	6.675 9	721
11	E ² RAG (vanilla)	6.648 1	718
12	vanilla RAG	6.601 9	713
13	LightRAG local	6.601 9	713
14	LightRAG hybrid (comb. embedding)	6.213 0	671
15	LightRAG hybrid (comb. extraction)	6.166 7	667
16	LightRAG hybrid (hyp. embedding)	6.148 1	664
17	LightRAG hybrid (hyp. extraction)	6.138 9	664
18	RQ-RAG	3.601 9	389

Algorithm 1 Entity–Event KG Insertion (Pre-processing)

Require: document D ; entity, event extractor $g(\cdot)$;
text encoder $f(\cdot)$; top- L chunk size m

Ensure: graphs $\mathcal{G}_{\text{ent}}, \mathcal{G}_{\text{evt}}$; bipartite edge set \mathcal{B} ; vector stores $\mathcal{I}_{\text{ent}}, \mathcal{I}_{\text{evt}}$

- 1: **procedure** INSERTDOCUMENT(D)
- 2: **Step 1: chunking**
- 3: $\{c_\ell\}_{\ell=1}^L \leftarrow \text{CHUNK}(D, m)$
- 4: **Step 2: node extraction**
- 5: **for** $\ell \leftarrow 1$ **to** L **do**
- 6: $V_{\text{ent}}^{(\ell)}, V_{\text{evt}}^{(\ell)} \leftarrow g(c_\ell)$
- 7: $V_{\text{ent}} \leftarrow \bigcup_\ell V_{\text{ent}}^{(\ell)}; V_{\text{evt}} \leftarrow \bigcup_\ell V_{\text{evt}}^{(\ell)}$
- 8: **Step 3: build edge sets**
- 9: $E_{\text{ent}}, E_{\text{evt}} \leftarrow \text{INTRACHUNKEDGES}(\{c_\ell\}, V_{\text{ent}}, V_{\text{evt}})$
- 10: $\mathcal{B} \leftarrow \text{BUILDBIPARTITE}(V_{\text{ent}}, V_{\text{evt}})$
- 11: **Step 4: embed and index**
- 12: **for all** $v \in V_{\text{ent}}$ **do** ▷ entities
- 13: $\mathcal{I}_{\text{ent}}.\text{ADD}(v, f(\text{CANON}(v)))$
- 14: **for all** $e \in V_{\text{evt}}$ **do** ▷ events
- 15: $\mathcal{I}_{\text{evt}}.\text{ADD}(e, f(\text{CANON}(e)))$
- 16: **return** $\mathcal{G}_{\text{ent}} = (V_{\text{ent}}, E_{\text{ent}}), \mathcal{G}_{\text{evt}} = (V_{\text{evt}}, E_{\text{evt}}), \mathcal{B}$

Algorithm 2 BUILDBIPARTITE: connect entity mentions to events

Require: entity nodes V_{ent} , event nodes V_{evt}

Ensure: bipartite edge set \mathcal{B}

- 1: **function** BUILDBIPARTITE($V_{\text{ent}}, V_{\text{evt}}$)
- 2: $\mathcal{B} \leftarrow \emptyset$
- 3: */* surface-form string match inside the same chunk */*
- 4: **for all** $v \in V_{\text{ent}}$ **do**
- 5: $c \leftarrow \text{CHUNKID}(v); S \leftarrow \text{NAME}(v)$
- 6: **for all** $e \in V_{\text{evt}}$ **with** $\text{CHUNKID}(e) = c$ **do**
- 7: **if** $S \subset \text{DESC}(e)$ **then**
- 8: $\mathcal{B} \leftarrow \mathcal{B} \cup \{(v, e)\}$
- 9: **return** \mathcal{B}

Algorithm 3 Entity–Event KG RETRIEVE (Inference-time)

Require: query q (or hypothetical response transformed text),

- 1: entity store \mathcal{I}_{ent} , event store \mathcal{I}_{evt} ,
- 2: bipartite edge set \mathcal{B} , similarity encoder $f(\cdot)$,
- 3: entity, event extractor $g(\cdot)$, top- k parameter k

Ensure: ranked context bundle C (passages + linearised subgraph)

```
4: procedure RETRIEVE( $q$ )
5:   Cue extraction
6:    $S_{\text{ent}}, S_{\text{evt}} \leftarrow g(q)$ 
7:    $V_q \leftarrow \text{LOOKUP}(S_{\text{ent}}, \mathcal{I}_{\text{ent}}) \cup \text{LOOKUP}(S_{\text{evt}}, \mathcal{I}_{\text{evt}})$ 
8:   One-hop expansion across  $\mathcal{B}$ 
9:    $V_q^+ \leftarrow V_q \cup \{e \mid (v, e) \in \mathcal{B}, v \in V_q\} \cup \{v \mid (v, e) \in \mathcal{B}, e \in V_q\}$ 
10:  Similarity ranking
11:   $\mathcal{P} \leftarrow \{\text{CHUNK}(x) \mid x \in V_q^+\}$  ▷ candidate passages
12:   $\mathbf{z} \leftarrow f(q)$ 
13:   $\text{scores}[p] \leftarrow \langle f(p), \mathbf{z} \rangle \quad \forall p \in \mathcal{P}$ 
14:   $\mathcal{P}_k \leftarrow \text{TOPK}(\text{scores}, k)$ 
15:  Context assembly
16:   $G_{\text{sub}} \leftarrow (V_q^+, (E_{\text{ent}} \cup E_{\text{evt}} \cup \mathcal{B}) \cap (V_q^+ \times V_q^+))$ 
17:   $C \leftarrow \text{FORMATCONTEXT}(\mathcal{P}_k, G_{\text{sub}})$ 
18:  return  $C$ 
```
