REALIZE WHAT MATTERS: PRINCIPLED CONTEXT REPRESENTATION FOR LARGE-SCALE REASONING

Anonymous authors
Paper under double-blind review

ABSTRACT

Solving complex tasks in domains such as science, medicine, law, and finance often requires assembling interdependent information scattered across vast, heterogeneous sources far beyond model context limits. Existing approaches tackle this challenge by organizing information into more manageable representations over which models can reason, such as graphs, textual memories, and retrieval collections. These representations dictate what downstream reasoning is possible and, ultimately, whether it succeeds; yet their design and construction remain largely ad hoc. In this work, drawing on the cognitive theory of relevance realization, we propose concrete principles for designing AI systems that construct effective representations of very large contexts. We analyze existing approaches and show how their successes and failures map onto their alignment with these principles, and introduce R3CON, a harness designed to operationalize the principles more systematically. We evaluate R3CON against nine state-of-the-art baselines on two recent benchmarks of reasoning over large document corpora. On these benchmarks, R3CON substantially outperforms the strongest baseline, by 20 and 8.4 percentage points. It also enables smaller models to outperform much larger ones: R3CON with `4B` and `9B` models outperforms all evaluated `35B` baselines, while R3CON with a `35B-A3B` model outperforms Claude Code with `Claude-Sonnet-5` at $3.7\times$ lower cost. Our results show that context representations following our principled approach can reduce reliance on model scale, pointing toward a future of AI systems with frontier-level performance powered by smaller models. Our code is available at `https://anonymous.4open.science/r/r3con-evaluation-anon`.

1 INTRODUCTION

Many complex tasks require assembling reasoning from partial, interdependent information sources scattered across corpora far too large for LLMs to consume directly. Such tasks are common across many domains of human endeavor: from lawyers reconstructing events from evidence (Gay & Odom, 2020; Day, 2024; Taranukhin et al., 2026), to scientists synthesizing prior studies (Ioannidis, 2023), financial analysts forecasting from SEC filings (Gibbons et al., 2021), or clinicians assessing patients from diverse medical records (Fan et al., 2025).

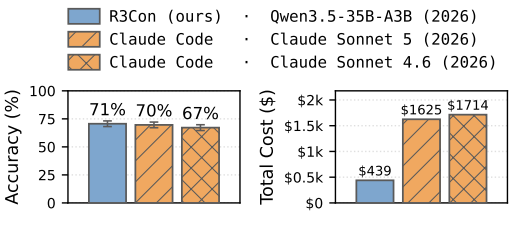


Figure 1: R3CON with a `3B`-active model outperforms Claude Code equipped with Anthropic's frontier Sonnet models on CorpusQA (Lu et al., 2026), while costing roughly $3.7\times$ to $3.9\times$ less.

To make such large contexts tractable for reasoning, much existing work first organizes the corpus into more manageable *representations*. In constructing these representations, methods make fundamentally different design decisions. Some pre-commit to sophisticated but fixed structures, such as graphs (Gutiérrez et al., 2025), retrieval pools (Sarthi et al., 2024), textual memories (Yu et al., 2026), or combinations thereof (Xia et al., 2026); others choose the structure dynamically (Li et al., 2025b). Some construct them offline (Edge et al., 2024; Lee et al., 2024; Zhuang et al., 2026), while others evolve them with the task (e.g., memory agents; see Hu et al., 2026). However, we lack a principled account of how representations should be constructed in the first place. Each of these many decisions

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

has profound implications for what downstream reasoning is possible (e.g., whether the solution is even accessible), yet it remains unclear which choices are necessary for robust reasoning, and why.

In this work, we provide a principled account of how to construct effective **R**epresentations of large **CON**texts, grounded in the cognitive theory of **R**elevance **R**ealization (Vervaeke et al., 2012). We first introduce this theory, which explains how humans arrive at a problem representation despite working-memory limits, and then translate the underlying cognitive processes into design principles for AI systems (§3). More specifically, such systems should ① extract relevant context in a *task-specific* and *enactive* manner. Then, they should ② structure the context using appropriate data structures, with the choice of structure informed by the relevant context and the task. Finally, the relevant context and structure should be brought together to ③ form the representation. Their interplay makes the representation easier to navigate and, in turn, the solution to the task more readily accessible.

Following these principles, we introduce **R3CON** (§4). R3CON uses recursive task-focused summarization to extract the relevant context, and then uses this information to organize the context into dynamically chosen data structures, e.g., tables, nested hierarchies, graphs. The relevant context and structure form the representation, which R3CON explores and reasons over using a coding agent.

We evaluate R3CON on two recent real-world benchmarks requiring the integration of information across large document corpora, against nine baselines ranging from state-of-the-art representation-construction systems to fully dynamic coding agents (§5). **R3CON improves over the strongest baseline by +19.88 percentage points on CorpusQA** (Lu et al., 2026) **and +8.42 points on Loong** (Wang et al., 2024a), as detailed in §6. Moreover, R3CON achieves the highest accuracy across all domains, including education, finance, real estate, science, and law. R3CON continues to outperform all baselines even with smaller, *cheaper* models. Specifically, R3CON with `Qwen3.5-4B` on Loong and `Qwen3.5-9B` on CorpusQA outperforms all baselines using the much larger `Qwen3.5-35B-A3B`. Finally, we find that R3CON with `Qwen3.5-35B-A3B` outperforms Claude Code equipped with Anthropic's frontier `Claude-Sonnet-5` while costing roughly 3.7× less (Fig. 1).

Our contributions are threefold: ① a set of cognitively grounded design principles for constructing effective representations of large contexts; ② an analysis of existing approaches showing how their failures map onto these principles; and ③ R3CON, an instantiation of these principles that clearly advances the state of the art on information integration tasks.

2 MOTIVATING EXAMPLE

Consider the task in Fig. 2 of identifying Dracula's victims in Bram Stoker's novel (1897). The challenge comes from the novel's epistolary form: the story unfolds from a corpus of documents written by its characters—letters, journals, newspaper articles, and others—each recording what its author knew at a particular point in time, often with incomplete information and without understanding the significance of what they observed. Answering correctly therefore requires piecing together information scattered across these incomplete accounts, much like an investigative journalist working through source materials (Stray, 2019; Taft, 2024; Spangher et al., 2025).

Task / Question:

On the voyage to England and during his stay there, how many people did Dracula kill, either directly or indirectly, and who were they?

Corpus:

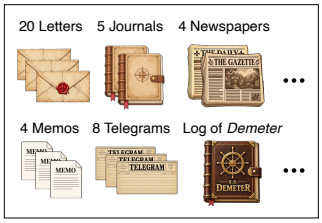


Figure 2: Running example based on Bram Stoker's *Dracula* (1897). Full details in §B.

Take, for example, one part of the story: the tragic events aboard the ship *Demeter* (see Fig. 3). On July 6, the ship sets sail for England, with its crew completely unaware that Dracula is hiding among the cargo. The captain establishes in the ship's log that nine people set sail—*"Crew, five hands ... two mates, cook, and myself"*—and then provides scattered reports of the crew disappearing one by one under mysterious circumstances. For instance, on July 16, he writes that *"one of crew, Petrofsky, was missing"*; eight days later, he reports *"another man lost"*; and so on, until only he remains.

After a month at sea, on August 7, the Demeter finally reaches the English coast, in the middle of a violent storm. Two days later, the strange arrival appears in the English newspaper *The Dailygraph*, which also reports an unidentified corpse *"lashed to the helm."* Taking a closer look at the logs, the

captain's last entry notes that he plans to tie his *"hands to the wheel"* until his *"strength begins to fail."* The newspaper therefore confirms the captain's death, the last of the crew of the Demeter.

Connecting all these deaths to Dracula requires piecing together the crew's accounts of a *"tall, thin man"*, *"ghastly pale"* and *"not like any of the crew"* roaming the ship with descriptions found elsewhere in the corpus. For example, Seward's diary later describes the vampire as a *"tall, thin man, clad in black"*. Together with the timeline of events, these accounts reveal that the mysterious man aboard the ship was Dracula, and he was the one responsible for the deaths of its crew (see Fig. 3).

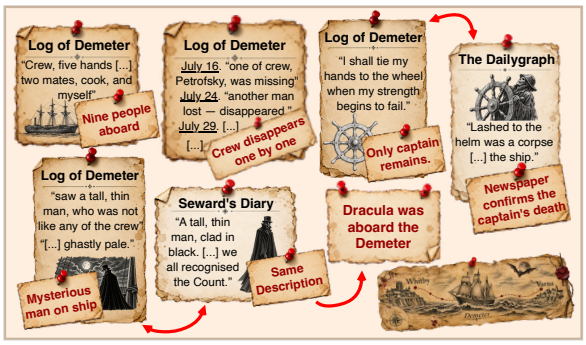


Figure 3: An investigation of the events aboard the ship Demeter. Figuring out that the entire crew of nine died, and that Dracula was behind it, requires connecting evidence across partial and incomplete accounts.

All nine baselines we evaluate fail on the Demeter deaths, and therefore fail the overall task in Fig. 2. We later trace these failures back to missing support for context representation principles (see §6.5). For example, HippoRAG2 (Gutiérrez et al., 2025) is unable to infer Dracula's involvement because the crew never identify him explicitly, leaving their deaths disconnected from him in the graph. MemAgent (Yu et al., 2026) also fails to make this connection: it forgets the deaths from its textual memory before it encounters the context needed to infer that Dracula was aboard (see §6.5). We trace both failures to limitations in how these methods interact with the context (i.e., *enactivity*; §3.1). Across the whole story, only R3CON manages to identify all 13 victims and hold Dracula fully accountable.

3 RELEVANCE REALIZATION

In this section, we introduce the theory of relevance realization (Vervaeke et al., 2012), and then translate its underlying cognitive processes into *computational principles for organizing large contexts*.

3.1 COGNITIVE THEORY

Humans routinely reason successfully through information that vastly exceeds our working memory limits (Baddeley, 2003; Jaeger et al., 2024). We do this by making good choices in organizing and relating information, i.e., choosing a problem representation conducive to solving the task. Yet, the space of all possible representations is enormous. In their seminal work, Vervaeke et al. (2012) put forward the theory of relevance realization, explaining how humans make this otherwise intractable problem tractable by selectively focusing on the relevant information and its relevant structure.

A central observation of Vervaeke et al. (2012) and Jackson et al. (2017) is that what is relevant depends on what we are trying to accomplish. The process of realizing relevance is *task-specific*. At the same time, our goals alone are not enough: relevance emerges through continued interaction with the environment, making the process *enactive*[1] (Jaeger et al., 2024). Through this interaction, what we learn changes our understanding, which in turn changes what we attend to and how we interpret information. As this understanding evolves, we zero in on the relevant information and its relevant structure, which form a representation that makes the problem tractable (Vervaeke et al., 2012).

3.2 COMPUTATIONAL DESIGN PRINCIPLES OPERATIONALIZING THE THEORY

We translate these cognitive processes into design principles for building AI systems capable of constructing effective representations of large contexts. To construct such representations, a system must first ① **extract relevant context** in a *task-specific* and *enactive* manner, allowing its understanding of both the task and the context to evolve. It should then ② **structure the context** guided by the extracted relevance and the task in a *flexible* way: the system should be free to choose both the structure's semantic organization (e.g., entities and relations, claims and evidence, etc.) and its data structure (e.g., tables, graphs, hierarchies, lists, etc.). Together, the extracted context and resulting

[1] *Enactive* refers to cognition emerging through continuous, reciprocal interaction between an agent and its environment, through which the agent actively makes sense of its world (Varela et al., 1991; Thompson, 2007).

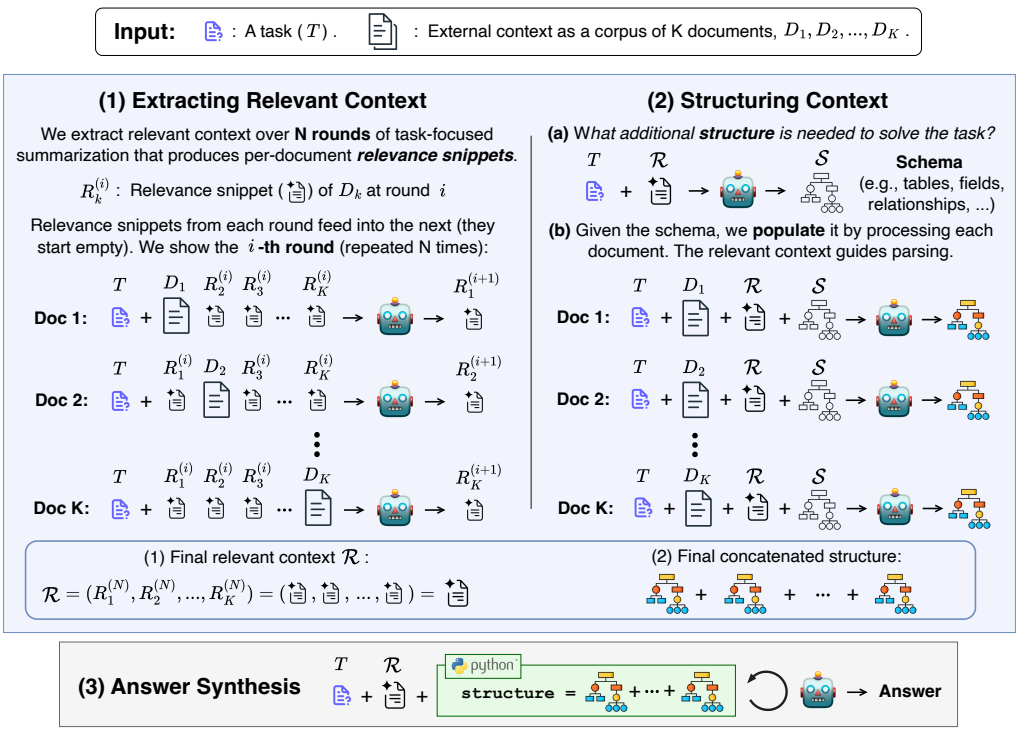


Figure 4: **R3CON** answers a task over a large context by first constructing an effective representation using our design principles (§3.2). It extracts the relevant context $\mathcal{R}$ (Phase 1) and then uses $\mathcal{R}$ to further structure the context (Phase 2). Together, the extracted relevant context and structured context form the representation, which is then used by a coding agent to produce the answer (Phase 3).

structure ③ **form the representation** that the system can then explore and reason over. Their interplay is important: the relevant information helps navigating the structure (i.e., contextualizes it), and the structure makes reasoning over relevant information easier.

4 R3CON

In this section, we introduce **R3CON** which instantiates the principles derived from the theory of **R**elevance **R**ealization for constructing an effective **R**epresentation of large **CON**texts (§3.2).

Problem Formulation. Let T be a task over a large context $\mathcal{D} = \{D_1, \ldots, D_K\}$, represented as a corpus of K text documents. The goal is to answer T using information available in the context.

4.1 EXTRACTING RELEVANT CONTEXT

In a sea of available information, the foremost challenge is extracting the context relevant to the task. Relevance, however, is not an intrinsic property of a passage or document: information that appears irrelevant in isolation may become crucial once we interact and consider other context.

Recall from our running example that *The Dailygraph* reports an unidentified corpse *"lashed to the helm"* (see Fig. 3). In isolation, this appears to be a tragic newspaper story and the relevance of these details to our task remains unclear. However, the final entry of the log of Demeter states the captain's plan to tie his *"hands to the wheel"* until his *"strength begins to fail."* Revisiting the newspaper account with this new context changes its interpretation and relevance: the unidentified corpse can now be recognized as the captain, and thus as Dracula's last victim on the voyage to England.

This illustrates the enactive nature of extracting relevance discussed in §3: as we interact with the broader environment (i.e., corpus), newly acquired context can change the interpretation—and therefore the relevance—of information encountered earlier. Thus, we *extract relevance* iteratively, revisiting information as our understanding of each document and of the corpus as a whole evolves.

	Round 1 - Relevance States	Round 2 - Relevance States
(Doc 37‡) Captain's log 	*A ship's log maintained by the captain of the Demeter [...] [...] **Specific victims** listed in the log include: Petrofsky (missing July 16), [...], **The captain survives the final entry** but anticipates his own death. This source provides the specific count and identities of the crew killed or driven to death during the voyage portion of the task.*	*A ship's log kept by the captain of the Demeter [...] The log documents the disappearance and **death of the nine crew members** one by one (naming Petrofsky, [...]) and **concludes with the Captain's death at the helm**. It contains [...] no information regarding fatalities during Dracula's subsequent **stay** in England.*
(Doc 12‡) The Dailygraph 	*A newspaper cutting from "The Dailygraph" [...] arrival of [...] Demeter at Whitby. [...] The report [...] **identifies the helmsman as a confirmed death**, found lashed to the wheel and dead for two days. [...] While this document [...] it does not list specific names or provide a definitive count of human victims.*	*A newspaper clipping from The Dailygraph [...] arrival of [...] Demeter in Whitby, [...] **It confirms the death of the ship's Captain**, found dead and lashed to the helm upon the ship's entry. [...] It contains [...] no information regarding fatalities that occur during the Count's subsequent stay on land.*

Figure 5: Evolution of the relevance snippets for the ship's log (D_{37}) and The Dailygraph newspaper (D_{12}) across two rounds of extracting relevance. For example, the top-left and top-right cells show $R_{37}^{(1)}$ and $R_{37}^{(2)}$, respectively. ‡Document ordering is arbitrary. The full snippets are shown in Fig. 12.

Instantiation Approach. More specifically, we revisit each document over N rounds using task-focused summarization (Dang, 2005). At each round i, for each document D_k, we construct a *relevance snippet* $R_k^{(i)}$. As illustrated in Phase 1 of Fig. 4, the language model receives the query T, the original document D_k, and the snippets of all *other* documents from the previous round (starting with $R_k^{(0)} = \emptyset$ for all k). It is then asked to summarize what the document contributes toward the task, retaining only task-relevant information (see Fig. 18 for the prompt). Here, we must be careful: what is relevant for the task might not be obvious in the current round—for example, necessary multi-hop connections may not yet have been discovered, as in the newspaper story about the unidentified dead helmsman. Therefore, when constructing D_k's relevance snippet at each round, we exclude its own previous state as it is potentially incomplete and distracting. This allows the model to re-interpret the document from scratch in light of the current corpus-wide understanding. Finally, after N rounds, we retain the final *relevant context* $\mathcal{R} = (R_1^{(N)}, R_2^{(N)}, ..., R_K^{(N)})$; we omit N from $\mathcal{R}$ for simplicity.

Fig. 5 illustrates this evolution across two rounds in our running example: by the second round, the broader context allows the necessary causal connections to emerge and relevance to be realized.

4.2 Structuring

Now that we have extracted the relevant context, we turn our focus to structuring it. We use $\mathcal{R}$ together with the task to guide both the choice of structure and how it is populated.

Schema. Given the task T and $\mathcal{R}$, the model is asked to propose a schema S that fully specifies how the information should be organized and related, including what information to represent, its attributes, relationships, nesting, and data types (see Fig. 19 for the prompt). For example, for the Dracula task, the model proposes a schema that organizes death events into two groups—those that happened during the voyage and those during Dracula's stay in England—and records attributes such as the name of the victim, cause of death, and supporting evidence (see Fig. 10 in §D.2.2). For other tasks, the schema can take more complex forms. For example, financial tasks may call for hierarchical or graph-like organizations of attributes (see Fig. 11 for one such chosen structure).

Parsing. Finally, given the schema S, we return to the original documents and populate it. Each document D_k is processed independently with the task T, the schema, and the relevant context $\mathcal{R}$ as context (see Fig. 4; the prompt is in Fig. 20). Crucially, $\mathcal{R}$ provides the contextual understanding needed to resolve ambiguities that commonly arise when structuring text (Martinelli et al., 2025).

4.3 Answer Synthesis

Together, the previous two stages construct the representation: the relevant context $\mathcal{R}$ and the resulting structure (§4.1 and §4.2). We use a standard agentic coding loop to reason over the representation (Zhang et al., 2025): this allows the model to explore and manipulate the structured data programmatically without having to fit the entire structure into its prompt. We provide $\mathcal{R}$ directly in the prompt (Fig. 21) and place the resulting structured data inside a Python Read-Eval-Print Loop (REPL) environment (see Fig. 4). We also study giving the full representation in-context in §D.2.4.

5 EXPERIMENTAL SETUP

5.1 BENCHMARKS

We select benchmarks that test reasoning over interdependent information distributed across large document collections from diverse real-world domains, including finance, education, real estate, science, and law. We evaluate on two complementary benchmarks. First, **CorpusQA** (Lu et al., 2026) focuses on scale and dispersion, with questions requiring the integration of information from hundreds of interconnected real-world documents at a time. Second, **Loong** (Wang et al., 2024a) offers diverse forms of reasoning across a collection of documents, with each question categorized by the reasoning it requires. This allows us to analyze performance separately on locating, comparing, clustering, and chaining evidence across the long context.

Metrics. Both benchmarks report accuracy and use an LLM-as-a-judge. We use the original prompts, with `gpt-5.4-mini` (OpenAI, 2026) as the judge. Further details are provided in Appendix C.1.

5.2 BASELINES

Reasoning over large contexts has been approached in several different ways. We evaluate state-of-the-art methods from different lines of work (see §C.2 for details).

First, *representation-based approaches* rely on transforming the context into more manageable representations that can then be used for reasoning. **HippoRAG2** (Gutiérrez et al., 2025) and **LinearRAG** (Zhuang et al., 2026) build sophisticated graphs, while **StructRAG** (Li et al., 2025b) chooses the structure dynamically from several options. Other work uses textual representations, as in **RAPTOR** (Sarthi et al., 2024), which builds a hierarchy of increasingly abstract summaries used for retrieval, and memory agents such as **ReadAgent** (Lee et al., 2024) and **MemAgent** (Yu et al., 2026).

On the other hand, *coding- and tool-calling agents* allow the model to dynamically adapt its reasoning strategy through multi-step loops based on the task and interactions with the context. We include **Claude Code** (Anthropic, 2025), **CodeAgent** (Roucher et al., 2025), and **Recursive Language Models (RLMs)** (Zhang et al., 2025), which allow models to recursively call themselves within a sandbox, enabling powerful decomposition patterns. Finally, we include **A-RAG** (Du et al., 2026), which equips the agent with multiple specialized retrieval tools.

5.3 MODELS

We use `Qwen3.5-35B-A3B` (Qwen Team, 2026) as the default model across R3CON and all baselines, unless otherwise specified. We additionally evaluate `Qwen3.5-4B` and `Qwen3.5-9B` with R3CON and the strongest baselines. For MemAgent (Yu et al., 2026), we use its RL-trained `RL-MemoryAgent-14B`. For experiments with Claude Code, we use `Claude-Sonnet-4.6` (Anthropic, 2026a) with `xhigh` reasoning and `Claude-Sonnet-5` (Anthropic, 2026b) with its default `high` reasoning. Finally, we use OpenAI's `text-embedding-3-small` for all embedding calls across baselines.

Cost. We report cost in USD ($) to provide a common unit across methods, including the cost of embedding models, using the lowest listed prices on OpenRouter[2] at the time of writing.

5.4 IMPLEMENTATION DETAILS

For R3CON, we use two rounds of extracting relevance ($N = 2$) throughout, unless otherwise specified. We analyze the effect of the number of rounds in §D.2.5. Further implementation details for each baseline are provided in §C.2. Infrastructure details are provided in §C.4.

6 RESULTS

6.1 OVERALL PERFORMANCE AND COST

R3CON outperforms baselines by large margins across benchmarks and domains. R3CON improves over the strongest baseline by +19.88 percentage points on CorpusQA and +8.42 points on Loong (Tab. 1). Moreover, R3CON achieves the highest accuracy all evaluated domains, including education, finance, real estate, academic papers and legal cases (Tab. 1). On CorpusQA, improvements

[2] `https://openrouter.ai/models`

Table 1: Accuracy (% ↑) across real-world domains requiring reasoning over large document corpora. R3CON achieves the highest accuracy across all domains, outperforming the strongest baseline by +19.88 pp on CorpusQA and +8.42 pp on Loong overall (SEs in Tab. 10). Accuracy is micro-averaged.

Method	**CorpusQA** (Lu et al., 2026)				**Loong** (Wang et al., 2024a)			
	Edu.	**Finance**	**Real-estate**	**Overall**	**Paper**	**Finance**	**Legal**	**Overall**
RAPTOR (2024)	8.00	11.11	17.95	12.06	0.00	22.17	5.71	12.12
ReadAgent (2024)	42.68	37.42	31.25	37.23	15.15	35.33	14.93	24.49
MemAgent (2026)	7.32	7.83	8.64	7.90	0.00	26.80	4.00	13.63
HippoRAG2 (2025)	7.32	7.83	20.99	10.94	0.32	17.42	6.29	10.06
LinearRAG (2026)	8.54	6.63	16.05	9.42	0.61	9.38	5.43	6.04
StructRAG (2025b)	53.66	55.42	38.27	50.76	19.70	32.66	12.32	23.72
CodeAgent (2025)	14.63	26.99	26.25	23.69	13.80	40.07	24.26	29.01
RLMs (2025)	13.41	37.58	20.51	27.38	12.77	41.67	22.77	28.96
A-RAG (2026)	8.54	7.23	18.52	10.33	0.61	18.43	2.86	9.55
R3CON (ours)	**75.61**	**82.42**	**41.25**	**70.64**	**27.58**	**50.59**	**24.29**	**37.43**

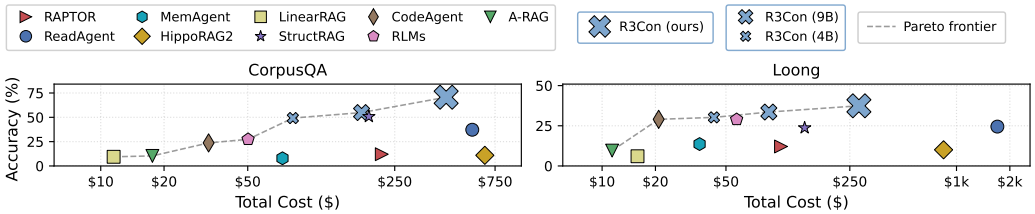


Figure 6: R3CON lies on the Pareto frontier of the *overall accuracy vs. cost* trade-off across model sizes. Smaller R3CON markers show runs with smaller models from the same family (`4B` and `9B`).

over the strongest baseline reach +27.00 points in finance and +21.95 points in education, while on Loong they reach +7.88 points on academic papers and +8.92 points in financial reports.

R3CON continues to outperform all baselines even with smaller, *cheaper* models. With the default `Qwen3.5-35B-A3B` used across R3CON and all baselines, R3CON is on the higher end of cost, although it remains less expensive than ReadAgent and HippoRAG2 (see Fig. 6). Importantly, we can reduce this cost by using smaller models while retaining superior accuracy over all baselines (see Fig. 6). More specifically, on Loong, R3CON with `Qwen3.5-4B` achieves 30.23% accuracy, surpassing all baselines using the default `35B` model, while reducing its cost by $6.58\times$. The same story holds on CorpusQA, where R3CON with `Qwen3.5-9B` achieves 54.88% accuracy, surpassing all baselines with $2.52\times$ lower cost. Notably, the strongest baselines become substantially less competitive when using these smaller models (see §D.1). Finally, beyond using smaller models, cost can be further reduced by decreasing the number of rounds N with little loss in accuracy (see §D.2.5).

6.2 FINE-GRAINED ANALYSIS OF REASONING FACETS

We now investigate how R3CON's performance varies across types of reasoning using Loong (2024a), since each example is annotated with the reasoning facet required to answer it, as shown in Fig. 7.

R3CON improves across all reasoning categories. R3CON achieves the highest accuracy in every category, outperforming the strongest baseline by +8.42 percentage points overall (Tab. 2).

R3CON's improvements increase with reasoning difficulty. Based on overall performance across methods, "chain of reasoning" and "clustering" are the most difficult facets (Tab. 2), and also where R3CON has the largest advantage over the strongest baseline. Its gains nearly double from +2.45 and +4.74 percentage points on "spotlight locating" and "comparison" to +8.89 and +8.03 on "chain of reasoning" and "clustering", respectively.

6.3 CAN'T CLAUDE CODE DO IT?

With the remarkable progress in coding agents (Clark, 2026), a natural question is whether a simpler approach might suffice: place all documents in a folder, provide the task, and *"ask Claude Code"*. We therefore test exactly this setup on CorpusQA with Claude Code (see §C.2.10 for details).

Table 2: Accuracy (% ↑) across Loong's reasoning facets.

Method	Spot.	Comp.	Chain	Clust.	Overall
RAPTOR (2024)	51.31	16.25	2.04	1.54	12.12
ReadAgent (2024)	60.11	33.77	18.49	10.70	24.49
MemAgent (2026)	42.13	35.00	0.64	0.95	13.63
HippoRAG2 (2025)	46.70	13.33	0.33	0.38	10.06
LinearRAG (2026)	32.49	3.75	0.64	0.38	6.04
StructRAG (2025b)	44.16	30.83	24.76	12.19	23.72
CodeAgent (2025)	63.54	45.26	22.80	12.90	29.01
RLMs (2025)	61.14	45.15	19.68	15.27	28.96
A-RAG (2026)	35.03	14.58	1.28	2.65	9.55
R3CON (ours)	**65.99**	**50.00**	**33.65**	**23.30**	**37.43**

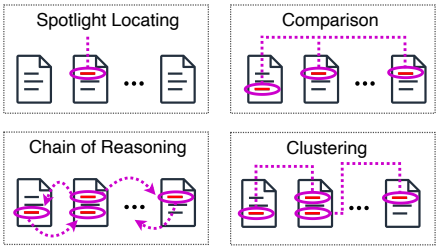


Figure 7: Reasoning facets for each example in Loong (Wang et al., 2024a).

R3CON with a 3B-active model surpasses Claude Code with frontier models in accuracy. R3CON achieves 71% accuracy using `Qwen3.5-35B-A3B` (3B active parameters). In comparison, Claude Code achieves 67% with `Claude-Sonnet-4.6` and 70% with `Claude-Sonnet-5` (Fig. 1).

R3CON costs 3.7× **to** 3.9× **less than Claude Code.** The improvements come at lower cost: R3CON costs $439, compared with $1,625 for Claude Code with `Sonnet-5` and $1,714 with `Sonnet-4.6`, making R3CON 3.7× and 3.9× cheaper, respectively, while achieving higher accuracy (Fig. 1).

6.4 ABLATION STUDIES

We ablate the two stages of R3CON which construct the representation used for reasoning (Fig. 4).

Removal of extracting relevance or structuring degrades performance. As shown in Tab. 3, removing extracting relevance has the largest effect. Notably, without this stage, R3CON becomes similar to StructRAG (Li et al., 2025b), and indeed achieves comparable performance (Tab. 1). Removing structuring also consistently degrades performance.

Table 3: Accuracy in the ablation of the two phases of R3CON (see Fig. 4).

Variant	Loong (↑)	CorpusQA (↑)
R3CON (default)	37.43%	70.64%
w/o extracting relevance	23.96% (-13.47)	52.58% (-18.1)
w/o structuring	32.97% (-4.46)	63.22% (-7.42)

6.5 UNDERSTANDING WHY R3CON SUCCEEDS WHILE OTHERS FAIL

In this section, we return to our running example (Fig. 2) and investigate why the nine baselines considered in this paper fail to attribute all deaths to Dracula, whereas R3CON succeeds.

Only R3CON satisfies all design principles. Recall from §3.2 that ① extracting relevance requires being both *task-specific* and *enactive* and ② structuring should be *flexible*. Our method directly instantiates each of these principles through its two stages, whose interplay allows R3CON to answer this task correctly, as illustrated throughout §4. We further analyze each baseline in §C.2 and find that none satisfy all principles, as shown in Tab. 4.

Table 4: Taxonomy of baselines according to the principles (§3.2) satisfied by their **representation-construction** phase. See §C.2 for detailed analysis. ◐ = limited support; N/A = not applicable.

Method	Task-specific	Enactive	Flexible Structuring
RAPTOR (2024)	✗	✓	✗
ReadAgent (2024)	✗	✗	✗
MemAgent (2026)	✓	✗	✗
HippoRAG2 (2025)	✗	◐	✗
LinearRAG (2026)	✗	◐	✗
StructRAG (2025b)	✓	✗	✓
CodeAgent (2025)	N/A	N/A	N/A
RLMs (2025)	N/A	N/A	N/A
Claude Code (2025)	N/A	N/A	N/A
A-RAG (2026)	✗	✗	✗
R3CON (ours)	✓	✓	✓



Why did Dracula get away with it? We focus on several representative baselines to illustrate how their failures relate to missing or limited support for particular principles. First, RAPTOR (2024) creates its representations (i.e., hierarchical summaries) **task-agnostically**. With no signal for what is relevant to the task, it discards details such as the specific crewmen who died aboard the Demeter and how many initially set sail (Fig. 14). Once discarded, the solution becomes inaccessible.

StructRAG (2025b) is task-specific, but has **no enactive mechanism** for reinterpreting earlier documents in light of new information. It discards crucial details about the captain's plan to tie his *"hands to the wheel"* until his *"strength begins to fail"*, as these only become relevant later when a newspaper reports a *"dead helmsman"* at the wheel. StructRAG therefore misses the captain's

death (Fig. 15). So do other graph-based approaches (e.g., HippoRAG2 does not even link Dracula to the Demeter events at all). A better, but still **limited**, instantiation of **enactivity** comes from MemAgent (2026), which maintains an evolving textual memory as it reads through the corpus. When it reads the ship's log, it does not yet know enough to attribute these deaths to Dracula, and information about the deaths gradually fades from its memory until it is forgotten altogether (Fig. 16).

Lastly, coding agents have **no constructed representation** at all, and instead rely on the model to dynamically discover an effective reasoning strategy. These agents all miss various multi-hop connections required by the task (§B.2). For example, the captain's death requires a two-hop connection, which even the best-performing coding agent, RLMs (2025), fails to make (Fig. 17).

7 Related Work

Work on reasoning over large contexts has largely followed two complementary directions: explicit representation construction to aid downstream reasoning and dynamic agentic reasoning.

Explicit Representation-Construction. In these approaches, the strategy for constructing a representation of the context is explicitly encoded in the system design (i.e., the harness). Some approaches construct and update textual memories as the context is processed (Xu et al., 2025; Hu et al., 2026; Yu et al., 2026), with other work constructing offline summaries at different levels of granularity (Sarthi et al., 2024; Lee et al., 2024). Other approaches choose to impose more explicit structure: GraphRAG and follow-up work use graphs (Edge et al., 2024; Gutiérrez et al., 2024; Gutiérrez et al., 2025; Zhuang et al., 2026), alongside broader efforts to improve RAG systems (Li et al., 2025a; Huang et al., 2025; Wu et al., 2026; Jin et al., 2026; Zhou et al., 2026). More recently, some methods make the structuring decision itself dynamic, including StructRAG (Li et al., 2025b), Structure-R1 (Wu et al., 2025), and LogicRAG (Chen et al., 2026). Taken together, these approaches provide effective but largely ad hoc answers to how context should be represented. Our work instead proposes principles that cut across these individual approaches, providing general guidance for how these representation-construction decisions should be made and why.

Agentic Reasoning. In contrast, coding and tool-calling agents (Yao et al., 2023; Wang et al., 2024b; Anthropic, 2025; Zhang et al., 2025; Li et al., 2026; Du et al., 2026) give models full freedom to structure their reasoning strategy at test time. Yet, this flexibility is a double-edged sword: the model must discover effective reasoning strategies on its own and there is increasing evidence that models struggle to do so in knowledge-intensive tasks (§6.3; Kargupta et al., 2025; Light et al., 2026; Theologitis et al., 2026; Xu et al., 2026). In our work, R3Con uses a coding agent only *after* an effective representation has been constructed, relieving the model of having to discover one on its own inside the coding loop. More broadly, our principles could also be used to create training trajectories that teach models how to organize context effectively at test time.

8 Conclusions

In this paper, we study the problem of organizing large heterogeneous contexts into representations that enable accurate downstream reasoning. Grounded in the cognitive theory of relevance realization, we propose a set of design principles to build rich task-specific and structured context representations. We then introduce R3Con which operationalizes them, and evaluate it against nine baselines on hard long-context benchmarks, outperforming prior state-of-the-art methods by up to 20 percentage points.

By keeping representation construction explicit, R3Con provides two important benefits. First, it keeps decisions about what is relevant and how information should be organized external to the model, where they can be inspected, adapted, and controlled by the data owner, rather than relying entirely on opaque relevance patterns internalized by the model during training. Second, better context representation reduces reliance on model scale: we demonstrate how R3Con enables models an order of magnitude smaller to outperform frontier systems using much larger models (§6.1 and §6.3). This could enable data owners to deploy smaller models on their own infrastructure, retaining control over sensitive knowledge while still achieving frontier-level performance. These benefits are especially valuable in domains like science, medicine, law, finance, and journalism, which often depend on sensitive knowledge that must be handled with care. Explicit representation construction and smaller, locally deployable models could provide a path toward AI systems that are not only capable, but also more customizable and amenable to inspection and control.

AI USE STATEMENT

In this work, we did not use generative AI tools for tasks requiring disclosure. Additionally, we used generative AI tools to polish the writing of the manuscript. We have reviewed all AI-assisted work. AI-assisted edits to the manuscript were reviewed by the authors for correctness and consistency with the intended meaning. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

We have made an effort to make this work as reproducible as possible. All prompts used by our method are provided in §E. Implementation details for all baseline methods are documented in §C.2. Benchmark details and additional evaluation metrics are provided in §C.1. Finally, all evaluation code is available at `https://anonymous.4open.science/r/r3con-evaluation-anon`.

REFERENCES

Anthropic. Claude code, 2025. URL `https://claude.com/product/claude-code`. Accessed: 2026-06-18.

Anthropic. Introducing Claude Sonnet 4.6, 2026a. URL `https://www.anthropic.com/news/claude-sonnet-4-6`. Accessed: 2026-08-18.

Anthropic. Introducing Claude Sonnet 5, 2026b. URL `https://www.anthropic.com/news/claude-sonnet-5`. Accessed: 2026-08-18.

Alan Baddeley. Working memory: looking back and looking forward. *Nature Reviews Neuroscience*, 4(10):829–839, 2003. doi: 10.1038/nrn1201. URL `https://doi.org/10.1038/nrn1201`.

Shengyuan Chen, Chuang Zhou, Zheng Yuan, Qinggang Zhang, Zeyang Cui, Hao Chen, Yilin Xiao, Jiannong Cao, and Xiao Huang. You don't need pre-built graphs for RAG: retrieval augmented generation with adaptive reasoning structures. In Sven Koenig, Chad Jenkins, and Matthew E. Taylor (eds.), *Fortieth AAAI Conference on Artificial Intelligence, Thirty-Eighth Conference on Innovative Applications of Artificial Intelligence, Sixteenth Symposium on Educational Advances in Artificial Intelligence, AAAI 2026, Singapore, January 20-27, 2026*, pp. 30270–30278. AAAI Press, 2026. doi: 10.1609/AAAI.V40I36.40278. URL `https://doi.org/10.1609/aaai.v40i36.40278`.

Kate Clark. The trillion dollar race to automate our entire lives, March 2026. URL `https://www.wsj.com/tech/ai/claude-code-cursor-codex-vibe-coding-52750531`. Accessed: 2026-08-26.

Hoa Trang Dang. Overview of DUC 2005. In *Proceedings of the Document Understanding Conference*, pp. 1–12, Vancouver, Canada, 2005. URL `https://duc.nist.gov/pubs/2005papers/OVERVIEW05.pdf`.

J. Brendan Day. "every little thing": Preparing your case for trial, 2024. URL `https://www.americanbar.org/groups/litigation/resources/litigation-journal/2024-winter/every-little-thing-preparing-your-case-trial/`. Accessed: 2026-08-24.

Mingxuan Du, Benfeng Xu, Chiwei Zhu, Shaohan Wang, Pengyu Wang, Xiaorui Wang, and Zhendong Mao. A-RAG: scaling agentic retrieval-augmented generation via hierarchical retrieval interfaces. *CoRR*, abs/2602.03442, 2026. doi: 10.48550/ARXIV.2602.03442. URL `https://doi.org/10.48550/arXiv.2602.03442`.

Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, and Jonathan Larson. From local to global: A graph RAG approach to query-focused summarization. *CoRR*, abs/2404.16130, 2024. doi: 10.48550/ARXIV.2404.16130. URL `https://doi.org/10.48550/arXiv.2404.16130`.

Hao Fan, Angela Hardi, and Po-Yin Yen. Supporting clinical reasoning through visual summarization and presentation of patient data: A systematic review. *Journal of the American Medical Informatics Association*, 32(9):1485–1498, 2025. doi: 10.1093/jamia/ocaf103.

Faith E. Gay and Mary Joy Odom. The only way to win, 2020. URL https://www.americanbar.org/groups/litigation/resources/litigation-journal/2020-fall/the-only-way-to-win/. Accessed: 2026-08-24.

Brian Gibbons, Peter Iliev, and Jonathan Kalodimos. Analyst information acquisition via EDGAR. *Management Science*, 67(2):769–793, 2021. doi: 10.1287/mnsc.2019.3465.

Bernal Jiménez Gutiérrez, Yiheng Shu, Yu Gu, Michihiro Yasunaga, and Yu Su. Hipporag: neurobiologically inspired long-term memory for large language models. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24, Red Hook, NY, USA, 2024. Curran Associates Inc. ISBN 9798331314385.

Bernal Jiménez Gutiérrez, Yiheng Shu, Weijian Qi, Sizhe Zhou, and Yu Su. From RAG to memory: Non-parametric continual learning for large language models. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, BC, Canada, July 13-19, 2025*, volume 267 of *Proceedings of Machine Learning Research*. PMLR / OpenReview.net, 2025. URL https://proceedings.mlr.press/v267/gutierrez25a.html.

Yuyang Hu, Shichun Liu, Yanwei Yue, Guibin Zhang, Boyang Liu, Fangyi Zhu, Jiahang Lin, Honglin Guo, Shihan Dou, Zhiheng Xi, Senjie Jin, Jiejun Tan, Yanbin Yin, Jiongnan Liu, Zeyu Zhang, Zhongxiang Sun, Yutao Zhu, Hao Sun, Boci Peng, Zhenrong Cheng, Xuanbo Fan, Jiaxin Guo, Xinlei Yu, Zhenhong Zhou, Zewen Hu, Jiahao Huo, Junhao Wang, Yuwei Niu, Yu Wang, Zhenfei Yin, Xiaobin Hu, Yue Liao, Qiankun Li, Kun Wang, Wangchunshu Zhou, Yixin Liu, Dawei Cheng, Qi Zhang, Tao Gui, Shirui Pan, Yan Zhang, Philip Torr, Zhicheng Dou, Ji-Rong Wen, Xuanjing Huang, Yu-Gang Jiang, and Shuicheng Yan. Memory in the age of ai agents, 2026. URL https://arxiv.org/abs/2512.13564.

Haoyu Huang, Yongfeng Huang, Junjie Yang, Zhenyu Pan, Yongqiang Chen, Kaili Ma, Hongzhi Chen, and James Cheng. Retrieval-augmented generation with hierarchical knowledge. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2025, Suzhou, China, November 4-9, 2025*, pp. 6044–6060. Association for Computational Linguistics, 2025. doi: 10.18653/V1/2025.FINDINGS-EMNLP.321. URL https://doi.org/10.18653/v1/2025.findings-emnlp.321.

John P. A. Ioannidis. Systematic reviews for basic scientists: a different beast. *Physiological Reviews*, 103(1):1–5, 2023. doi: 10.1152/physrev.00028.2022. URL https://doi.org/10.1152/physrev.00028.2022. PMID: 36049113.

Jade Jackson, Anina N. Rich, Mark A. Williams, and Alexandra Woolgar. Feature-selective attention in frontoparietal cortex: Multivoxel codes adjust to prioritize task-relevant information. *Journal of Cognitive Neuroscience*, 29(2):310–321, 02 2017. ISSN 0898-929X. doi: 10.1162/jocn_a_01039. URL https://doi.org/10.1162/jocn_a_01039.

Johannes Jaeger, Anna Riedl, Alex Djedovic, John Vervaeke, and Denis Walsh. Naturalizing relevance realization: why agency and cognition are fundamentally not computational. *Front Psychol*, 15: 1362658, June 2024.

Yiqiao Jin, Kartik Sharma, Vineeth Rakesh, Yingtong Dou, Menghai Pan, Mahashweta Das, and Srijan Kumar. SARA: Selective and adaptive retrieval-augmented generation with context compression. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14508–14528, San Diego, California, United States, July 2026. Association for Computational Linguistics. ISBN 979-8-89176-390-6. doi: 10.18653/v1/2026.acl-long.661. URL https://aclanthology.org/2026.acl-long.661/.

Priyanka Kargupta, Shuyue Stella Li, Haocheng Wang, Jinu Lee, Shan Chen, Orevaoghene Ahia, Dean Light, Thomas L. Griffiths, Max Kleiman-Weiner, Jiawei Han, Asli Celikyilmaz, and Yulia Tsvetkov. Cognitive foundations for reasoning and their manifestation in llms. *CoRR*, abs/2511.16660, 2025. doi: 10.48550/ARXIV.2511.16660. URL https://doi.org/10.48550/arXiv.2511.16660.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.

Kuang-Huei Lee, Xinyun Chen, Hiroki Furuta, John F. Canny, and Ian Fischer. A human-inspired reading agent with gist memory of very long contexts. In Ruslan Salakhutdinov, Zico Kolter, Katherine A. Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, volume 235 of *Proceedings of Machine Learning Research*, pp. 26396–26415. PMLR / OpenReview.net, 2024. URL https://proceedings.mlr.press/v235/lee24c.html.

Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and Zhicheng Dou. Search-o1: Agentic search-enhanced large reasoning models. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, EMNLP 2025, Suzhou, China, November 4-9, 2025*, pp. 5420–5438. Association for Computational Linguistics, 2025a. doi: 10.18653/V1/2025.EMNLP-MAIN.276. URL https://doi.org/10.18653/v1/2025.emnlp-main.276.

Zhuofeng Li, Haoxiang Zhang, Cong Wei, Pan Lu, Ping Nie, Yi Lu, Yuyang Bai, Shangbin Feng, Hangxiao Zhu, Ming Zhong, Yuyu Zhang, Jianwen Xie, Yejin Choi, James Zou, Jiawei Han, Wenhu Chen, Jimmy Lin, Dongfu Jiang, and Yu Zhang. Beyond semantic similarity: Rethinking retrieval for agentic search via direct corpus interaction. *CoRR*, abs/2605.05242, 2026. doi: 10.48550/ARXIV.2605.05242. URL https://doi.org/10.48550/arXiv.2605.05242.

Zhuoqun Li, Xuanang Chen, Haiyang Yu, Hongyu Lin, Yaojie Lu, Qiaoyu Tang, Fei Huang, Xianpei Han, Le Sun, and Yongbin Li. Structrag: Boosting knowledge intensive reasoning of llms via inference-time hybrid information structurization. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025b. URL https://openreview.net/forum?id=GhexuBLxbO.

Dean Light, Michael Theologitis, Kshitish Ghate, Shuyue Stella Li, Benjamin Newman, Chirag Shah, Aylin Caliskan, Pang Wei Koh, Dan Suciu, and Yulia Tsvetkov. Deep reasoning in general purpose agents via structured meta-cognition, 2026. URL https://arxiv.org/abs/2605.11388.

Zhiyuan Lu, Chenliang Li, Yingcheng Shi, Weizhou Shen, Ming Yan, and Fei Huang. Corpusqa: A 10 million token benchmark for corpus-level analysis and reasoning. *CoRR*, abs/2601.14952, 2026. URL https://arxiv.org/abs/2601.14952.

Giuliano Martinelli, Bruno Gatti, and Roberto Navigli. xCoRe: Cross-context coreference resolution. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 34264–34278, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.1737. URL https://aclanthology.org/2025.emnlp-main.1737/.

Xuying Ning, Katherine Tieu, Dongqi Fu, Tianxin Wei, Zihao Li, Yuanchen Bei, Jiaru Zou, Mengting Ai, Zhining Liu, Ting-Wei Li, Lingjie Chen, Yanjun Zhao, Ke Yang, Bingxuan Li, Cheng Qian, Gaotang Li, Xiao Lin, Zhichen Zeng, Ruizhong Qiu, Sirui Chen, Yifan Sun, Xiyuan Yang, Ruida Wang, Rui Pan, Chenyuan Yang, Dylan Zhang, Liri Fang, Zikun Cui, Yang Cao, Pan Chen, Dorothy Sun, Ren Chen, Mahesh Srinivasan, Nipun Mathur, Yinglong Xia, Hong Li, Hong Yan, Pan Lu, Lingming Zhang, Tong Zhang, Hanghang Tong, and Jingrui He. Code as agent harness, 2026. URL https://arxiv.org/abs/2605.18747.

OpenAI. Introducing GPT-5.4 mini and nano, 2026. URL https://openai.com/index/introducing-gpt-5-4-mini-and-nano. Accessed: 2026-08-17.

Richard Yuanzhe Pang, Alicia Parrish, Nitish Joshi, Nikita Nangia, Jason Phang, Angelica Chen, Vishakh Padmakumar, Johnny Ma, Jana Thompson, He He, and Samuel R. Bowman. Quality: Question answering with long input texts, yes! In Marine Carpuat, Marie-Catherine de Marneffe, and Iván Vladimir Meza Ruíz (eds.), *Proceedings of the 2022 Conference of the North American*

Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022, pp. 5336–5358. Association for Computational Linguistics, 2022. doi: 10.18653/V1/2022.NAACL-MAIN.391. URL https://doi.org/10.18653/v1/2022.naacl-main.391.

Qwen Team. Qwen3.5: Towards native multimodal agents, February 2026. URL https://qwen.ai/blog?id=qwen3.5.

Aymeric Roucher, Albert Villanova del Moral, Thomas Wolf, Leandro von Werra, and Erik Kaunismäki. 'smolagents': a smol library to build great agentic systems. https://github.com/huggingface/smolagents, 2025.

Parth Sarthi, Salman Abdullah, Aditi Tuli, Shubh Khanna, Anna Goldie, and Christopher D. Manning. RAPTOR: recursive abstractive processing for tree-organized retrieval. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL https://openreview.net/forum?id=GN921JHCRw.

Alexander Spangher, Tenghao Huang, Yiqin Huang, Lucas Spangher, Sewon Min, and Mark Dredze. A novel multi-document retrieval benchmark: Journalist source-selection in newswriting. In Weijia Shi, Wenhao Yu, Akari Asai, Meng Jiang, Greg Durrett, Hannaneh Hajishirzi, and Luke Zettlemoyer (eds.), *Proceedings of the 4th International Workshop on Knowledge-Augmented Methods for Natural Language Processing*, pp. 180–204, Albuquerque, New Mexico, USA, May 2025. Association for Computational Linguistics. ISBN 979-8-89176-229-9. doi: 10.18653/v1/2025.knowledgenlp-1.18. URL https://aclanthology.org/2025.knowledgenlp-1.18/.

Bram Stoker. *Dracula*. Project Gutenberg, 1897. URL https://www.gutenberg.org/ebooks/345. Project Gutenberg eBook No. 345.

Jonathan Stray. Making artificial intelligence work for investigative journalism. *Digital Journalism*, 7 (8):1076–1097, 2019. doi: 10.1080/21670811.2019.1630289. URL https://doi.org/10.1080/21670811.2019.1630289.

Isabelle Taft. This mississippi hospital transfers some patients to jail to await mental health treatment, 2024. URL https://www.propublica.org/article/baptist-desoto-hospital-civil-commitment-jail. Accessed: 2026-08-24.

Maksym Taranukhin, Shuyue Stella Li, Evangelos Milios, Geoff Pleiss, Yulia Tsvetkov, and Vered Shwartz. Infogatherer: Principled information seeking via evidence retrieval and strategic questioning, 2026. URL https://arxiv.org/abs/2603.05909.

Michael Theologitis and Dan Suciu. Thucy: An llm-based multi-agent system for claim verification across relational databases, 2026. URL https://arxiv.org/abs/2512.03278.

Michael Theologitis, Preetam Prabhu Srikar Dammu, Chirag Shah, and Dan Suciu. Claimdb: A fact verification benchmark over large structured data. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2026, San Diego, California, United States, July 2-7, 2026*, pp. 34428–34451. Association for Computational Linguistics, 2026. doi: 10.18653/V1/2026.ACL-LONG.1589. URL https://doi.org/10.18653/v1/2026.acl-long.1589.

Evan Thompson. *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press, Cambridge, MA, 2007.

Francisco J. Varela, Eleanor Rosch, and Evan Thompson. *The Embodied Mind: Cognitive Science and Human Experience*. The MIT Press, 09 1991. ISBN 9780262285476. doi: 10.7551/mitpress/6730.001.0001. URL https://doi.org/10.7551/mitpress/6730.001.0001.

John Vervaeke, Timothy P. Lillicrap, and Blake A. Richards. Relevance realization and the emerging framework in cognitive science. *Journal of Logic and Computation*, 22(1):79–99, 02 2012. ISSN 0955-792X. doi: 10.1093/logcom/exp067. URL https://doi.org/10.1093/logcom/exp067.

Minzheng Wang, Longze Chen, Fu Cheng, Shengyi Liao, Xinghua Zhang, Bingli Wu, Haiyang Yu, Nan Xu, Lei Zhang, Run Luo, Yunshui Li, Min Yang, Fei Huang, and Yongbin Li. Leave no document behind: Benchmarking long-context llms with extended multi-doc QA. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pp. 5627–5646. Association for Computational Linguistics, 2024a. doi: 10.18653/V1/2024.EMNLP-MAIN.322. URL https://doi.org/10.18653/v1/2024.emnlp-main.322.

Xingyao Wang, Yangyi Chen, Lifan Yuan, Yizhe Zhang, Yunzhu Li, Hao Peng, and Heng Ji. Executable code actions elicit better LLM agents. In Ruslan Salakhutdinov, Zico Kolter, Katherine A. Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, volume 235 of *Proceedings of Machine Learning Research*, pp. 50208–50232. PMLR / OpenReview.net, 2024b. URL https://proceedings.mlr.press/v235/wang24h.html.

Junlin Wu, Xianrui Zhong, Jiashuo Sun, Bolian Li, Bowen Jin, Jiawei Han, and Qingkai Zeng. Structure-r1: Dynamically leveraging structural knowledge in LLM reasoning through reinforcement learning. *CoRR*, abs/2510.15191, 2025. doi: 10.48550/ARXIV.2510.15191. URL https://doi.org/10.48550/arXiv.2510.15191.

Mingyan Wu, Zhenghao Liu, Xinze Li, Yuqing Lan, Yukun Yan, Shuo Wang, Cheng Yang, Minghe Yu, Zheni Zeng, and Maosong Sun. Finding what matters: Anchoring context knowledge with evolving indices for iterative retrieval. 2026. URL https://arxiv.org/abs/2601.16462.

Menglin Xia, Xuchao Zhang, Shantanu Dixit, Paramaguru Harimurugan, Rujia Wang, Victor Ruhle, Robert Sim, Chetan Bansal, and Saravan Rajmohan. Memora: A harmonic memory representation balancing abstraction and specificity, 2026. URL https://arxiv.org/abs/2602.03315.

Hanwen Xu, Xuyao Huang, Yuzhe Liu, and Zhijie Deng. TPS-bench: Evaluating AI agents' tool planning & scheduling abilities in compounding tasks. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 34949–34961, San Diego, California, United States, July 2026. Association for Computational Linguistics. ISBN 979-8-89176-390-6. doi: 10.18653/v1/2026.acl-long.1614. URL https://aclanthology.org/2026.acl-long.1614/.

Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. A-mem: Agentic memory for LLM agents. In Danielle Belgrave, Cheng Zhang, Laura N. Montoya, Hsuan-Tien Lin, Razvan Pascanu, Piotr Koniusz, Marzyeh Ghassemi, Nancy Chen, Iván Vladimir Meza Ruíz, and Arturo Loaiza-Bonilla (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2025, NeurIPS 2025, San Diego, CA, USA, December 2-7, 2025 / Mexico City, Mexico, November 30 - December 5, 2025*, 2025. URL http://papers.nips.cc/paper_files/paper/2025/hash/19909c36f51abc4856b4560aff3d36d6-Abstract-Conference.html.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R. Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL https://openreview.net/forum?id=WE_vluYUL-X.

Hongli Yu, Tinghong Chen, Jiangtao Feng, Jiangjie Chen, Weinan Dai, Qiying Yu, Ya-Qin Zhang, Wei-Ying Ma, Jingjing Liu, Mingxuan Wang, and Hao Zhou. Memagent: Reshaping long-context LLM with multi-conv RL-based memory agent. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=k5nIOvYGCL.

Alex L. Zhang, Tim Kraska, and Omar Khattab. Recursive language models. *CoRR*, abs/2512.24601, 2025. doi: 10.48550/ARXIV.2512.24601. URL https://doi.org/10.48550/arXiv.2512.24601.

Junzhe Zhou, Fulin Lin, Tairan Cheng, Shaowen Chen, and Hongwei Wang. AED-RAG: Continuous multi-granular context fusion for retrieval-augmented generation via adaptive ensemble decoding.

In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 22885–22899, San Diego, California, United States, July 2026. Association for Computational Linguistics. ISBN 979-8-89176-395-1. doi: 10.18653/v1/2026.findings-acl.1148. URL https://aclanthology.org/2026.findings-acl.1148/.

Luyao Zhuang, Shengyuan Chen, Yilin Xiao, Huachi Zhou, Yujing Zhang, Hao Chen, Qinggang Zhang, and Xiao Huang. LinearRAG: Linear graph retrieval augmented generation on large-scale corpora. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=mCtfkypdm6.

Appendix Contents

A Limitations **16**

B Running Example **16**

B.1 Bram Stoker's Dracula (1897) 16

B.2 Evidence Attribution 17

B.3 Scoring 19

B.4 Parametric Knowledge 19

C Experimental Details **19**

C.1 Benchmarks 19

C.1.1 CorpusQA 19

C.1.2 Loong 21

C.2 Baselines 21

C.2.1 RAPTOR 21

C.2.2 ReadAgent 21

C.2.3 MemAgent 22

C.2.4 HippoRAG2 22

C.2.5 LinearRAG 22

C.2.6 StructRAG 23

C.2.7 CodeAgent 23

C.2.8 RLMs 23

C.2.9 A-RAG 24

C.2.10 Claude Code 24

C.3 Models 24

C.4 Infrastructure 25

D Detailed Results **25**

D.1 Evaluating Baselines with Smaller Models 25

D.2 R3Con – Detailed Analyses 25

D.2.1 Size of Relevance and Structuring 25

D.2.2 Examples of Structures 25

D.2.3 Examples of Relevance Snippets 26
D.2.4 Downstream Reasoning 26
D.2.5 Effect of Rounds 27
D.3 Qualitative Results 27
D.4 Standard Errors for Main Results 28

E Prompts **31**

A LIMITATIONS

We discuss several limitations and open questions that also motivate directions for future work.

Reusing representations across similar tasks. We build each representation based on the task, following our first principle of being task-specific. This means that R3CON needs to recompute the representation (i.e., relevant context and structure) for each new query. In terms of efficiency, this is more expensive than methods that can construct a generic, task-agnostic representation and reuse it across different tasks; still, such an approach risks creating a subpar representation which can hurt downstream reasoning, as we see in §6. In future work, R3CON could embed and store relevance snippets and structures across queries, retrieving and reusing them for similar questions rather than recomputing everything from scratch. Over time, this could result in a rich collection of task-specific representations that could be combined and explored to provide an even better view of the context.

Budgeting attention and intelligence in relevance extraction. In every iteration of relevance realization in R3CON, the relevance snippets from all documents are fed back into the model. This assumes that the model's context window is large enough to accommodate them. In Appendix D.2.1, we show that, in practice, the combined size of all relevance snippets across all documents remains roughly 1% of the original input size for each example, giving substantial headroom as context windows continue to grow. However, scaling to unbounded corpora such as the open web will eventually require mechanisms that allocate attention more selectively. One avenue is more sophisticated enactive mechanisms that use the understanding developed in each round to discard documents deemed irrelevant, reducing the context that subsequent rounds need to process.

B RUNNING EXAMPLE

B.1 BRAM STOKER'S DRACULA (1897)

Dracula (1897) is written in epistolary style—the story is assembled from a collection of documents (e.g., journals, telegrams, letters, etc.) produced by its characters as events unfold. The story is told in a roughly chronological order by piecing together relevant snippets from the different text artifacts. For example, the events of September 17–18 are told by piecing together entries from Lucy Westenra's and Dr. Seward's diaries, a telegram from Van Helsing to Seward, a memorandum left by Lucy describing the events of the night, and a letter from Mina to Lucy.

No individual document contains the complete story. Each records what its author knew at a particular point in time, often without understanding the significance of what they observed and with incomplete information. For example, when Mina discovers two small wounds on Lucy's neck, she attributes them to accidentally pricking Lucy with a safety-pin. Only when this observation is considered alongside later accounts of Lucy's deteriorating health and Dracula's activities does its significance become clear. Similarly, Dracula's arrival in England is recorded independently through the ship's log, a newspaper account of the shipwreck, and Mina's observations in Whitby. Understanding what happened therefore requires connecting information written by different people, at different times, and for different purposes.

Documents. For our purposes, we reconstruct the individual text artifacts that make up the novel. For example, Jonathan Harker's journal appears throughout the book as a series of dated entries. We collect and concatenate these entries to reconstruct Jonathan Harker's journal—a single document.

Document type	# Docs.	Description
Journal / diary	5	Personal journals and diaries (e.g., Jonathan Harker's journal)
Letter	20	Letters between characters (e.g., Mina to Lucy)
Telegram	8	Telegrams between characters (e.g., Van Helsing to Seward)
Memorandum	4	Written accounts of events (e.g., Lucy's final memorandum)
Newspapers	4	Newspaper cuttings (e.g., The Pall Mall Gazette)
Note	2	Short written notes (e.g., Van Helsing's contingency note)
Ship's log	1	The log of the Demeter ship (from the captain)
Report	1	Medical report (Dr. Hennessey to Dr. Seward)
Phonograph	1	Transcribed recording of Van Helsing
Total #Docs	**46**	
Total Tokens	**220k**	

Table 5: The reconstructed document collection in Bram Stoker's Dracula (1897).

We do the same for the other recurring journals and diaries. Standalone artifacts, such as individual letters, telegrams, memoranda, and newspaper cuttings, are treated as separate documents. We end up with a collection of 46 documents of different types, summarized in Tab. 5.

Task. We consider the following query over the resulting corpus:

> *On the voyage to England and during his stay there, how many people did Dracula kill, either directly or indirectly, and who were they?*

Answer. The answer is 13: the 9 members of the *Demeter*'s crew, Mr. Swales, Mrs. Westenra, Lucy Westenra, and Renfield.

Of course, this answer does not appear explicitly in any individual document. In fact, each individual death follows a lot of highly non-trivial reasoning steps throughout the different documents (§B.2).

B.2 Evidence Attribution

We trace the relevant evidence for each of the 13 deaths below.

Demeter's crew (9). The main evidence is scattered throughout the ship's log. First, we must establish that the ship leaves Varna with nine people aboard: the captain says *"On 6 July we finished taking in cargo, silver sand and boxes of earth."* and then *"At noon set sail. East wind, fresh. Crew, five hands ... two mates, cook, and myself (captain)."* Over the course of the voyage, the captain records the crew disappearing or dying one by one: *"one of crew, Petrofsky, was missing"* (16 July); *"another man lost–disappeared"* (24 July); *"Another tragedy ... Are now without second mate"* (29 July); *"both man of watch and steersman missing. Only self and mate and two hands left to work ship"* (30 July); *"One more gone. Lord, help us!"* (2 August); and *"I went to relieve the man at the wheel, and when I got to it found no one there"* (3 August). Later that same day, the mate *"sprang on the bulwark and deliberately threw himself into the sea"*. By the end, only the captain remains.

We must also establish that Dracula is aboard the ship. The crew repeatedly report seeing a mysterious man, including *"a tall, thin man, who was not like any of the crew, ..., go along the deck forward, and disappear."*, *"I saw It, like a man, tall and thin, and ghastly pale."*, and, finally, *"in the dimness of the night I saw It–Him!"* Other accounts identify this figure as Dracula: Seward later encounters a *"tall, thin man, clad in black"* whom he and the others recognize as the Count.

The captain is the last remaining member of the crew. His final log entry states: *"I shall tie my hands to the wheel when my strength begins to fail."* His death, however, can only be confirmed by a different document: an English newspaper cutting from The Dailygraph describes the ship's arrival with *"lashed to the helm was a corpse, with drooping head, which swung horribly to and fro at each motion of the ship."* Together, the log and newspaper account identify the otherwise unnamed corpse as the captain.

Mr. Swales (1). The Demeter brings Dracula from Transylvania to England, arriving in the coastal town of Whitby. While staying there, Mina befriends Mr. Swales, an elderly local whom she frequently meets at their favorite seat in the churchyard overlooking the harbor. Shortly after the ship's arrival, Mina records that *"poor old Mr. Swales was found dead this morning on our seat, his neck being broken."* The circumstances of his death are unusual: according to the doctor, he had *"fallen back in the seat in some sort of fright,"* with a look of *"fear and horror on his face."* Mina herself wonders, *"Perhaps he had seen Death with his dying eyes!"*

Connecting his death to Dracula requires evidence surrounding the location and timing. When the Demeter arrives two days earlier, a newspaper account describes an *"immense dog"* leaping from the ship and *"making straight for the steep cliff, where the churchyard hangs over the laneway to the East Pier."* The day after Swales is found dead, Mina returns to the same seat and finds Lucy there asleep. She observes *"something dark"* bending over Lucy and, as she approaches, sees *"a white face and red, gleaming eyes."* It can be safely assumed that the mysterious dog that came from the ship was Dracula himself in one of his guises, and that it was responsible for knocking the old man down and causing his death.[3]

Mrs. Westenra (1). Mrs. Westenra is Lucy's mother, with whom Lucy lives at Hillingham. She suffers from a serious heart condition. And, Mina Murray notes (in her journal) that *"her heart is weakening"* and that *"a sudden shock would be almost sure to kill her."* On the night of September 17, while Mrs. Westenra is in Lucy's room, a wolf crashes through the window. In a memorandum left by Lucy, she records that *"in the aperture of the broken panes there was the head of a great, gaunt grey wolf."* Her mother *"cried out in a fright"* and then suddenly collapsed; Lucy realizes shortly afterward that *"her dear heart had ceased to beat."*

Connecting the wolf to Dracula requires a separate newspaper account. In an interview with the zookeeper published in The Pall Mall Gazette, a man matching Dracula's appearance visits the wolf enclosure earlier that day. The keeper describes him as *"a tall, thin chap, with a 'ook nose and a pointed beard"* with *"red eyes"* and *"a mouth full of white, sharp teeth."* The same newspaper account reports that a wolf named Bersicker escapes from the enclosure that night. After Mrs. Westenra's death, Bersicker returns to the zoo with his head *"all cut and full of broken glass"*—matching Lucy's memorandum describing the wolf crashing through her window. Together, these accounts connect Dracula to the escaped wolf whose attack frightens Mrs. Westenra and kills her indirectly.

Lucy Westenra (1). Lucy is Mina's close friend and becomes Dracula's primary victim after his arrival in England. The first attack occurs in the Whitby churchyard. Mina's journal describes finding Lucy asleep on their favorite seat with *"something dark"* bending over her and later notices *"two little red points like pin-pricks"* on Lucy's neck. Mina initially assumes that she caused them herself with a safety-pin, but later observes that *"the tiny wounds seem not to have healed. They are still open, and, if anything, larger than before."*

Over the following weeks, Lucy becomes increasingly weak and continues to lose blood with no apparent explanation. In her own diary, she records hearing *"scratching or flapping at the window"* during the night and waking *"horribly weak."* Dr. Seward's diary records that her condition becomes so severe that she receives blood transfusions from four different men, yet the blood repeatedly seems to disappear. Together with the persistent wounds on her neck and the nighttime encounters, these accounts establish that Dracula continues to feed on Lucy. Her condition eventually deteriorates, and on September 20, Dr. Seward's diary records her death: *"And then Lucy's breathing became stertorous again, and all at once it ceased. It is all over,' said Van Helsing. She is dead!"'*

Renfield (1). Renfield is a patient in Dr. Seward's asylum who initially regards Dracula as his *"Master"*, but eventually turns against him after realizing that Dracula has been attacking Mina. While mortally wounded, Renfield recounts their confrontation to Seward: *"it made me mad to know that He had been taking the life out of her"* and, when he tried to stop Dracula, *"He raised me up and flung me down."* Seward's diary records that Renfield's *"back is broken"* and that *"the real injury was a depressed fracture of the skull."* His death is later confirmed in the same diary when Arthur reports that *"the poor fellow is dead."*

[3] https://www.cliffsnotes.com/literature/d/dracula/summary-and-analysis/chapters-78

Table 6: Answers to the Dracula question using parametric knowledge alone and with the full novel provided in context. The ground-truth answer is 13, but also requires the correct people named. Full model responses are shown in Fig. 8.

Model	Parametric knowledge		With corpus in context	
	Answer	Correct	Answer	Correct
`Qwen3.5-35B-A3B` (2026)	31	✗	7	✗

Excluded Deaths. Identifying the victims also requires reasoning about which deaths should *not* be counted. Four deaths connected to Dracula fall outside the question's time window. Before the Demeter sets sail, at the castle in Transylvania, Dracula brings an infant to the three vampire sisters, followed by the child's mother, who comes to the castle crying *"Monster, give me my child!"* and is killed by the wolves he summons. After Dracula flees England, Petrof Skinsky, the agent hired to move his last box, is murdered—a death Mina explicitly attributes to Dracula, who *"blotted out his traces, as he thought, by murdering his agent"*. Quincey Morris also dies after Dracula's flight, and is killed by the Szgany guarding Dracula's cart rather than by Dracula himself. Two other deaths also happen within the time window but must still be excluded. Mr. Hawkins *"has died very suddenly"* in Exeter, with no evidence connecting his death to Dracula, while Arthur's father dies after a longer decline documented across several accounts. Finally, several apparent victims do not die at all: the Hampstead children survive their attacks, as does Mina, who remains alive in the closing note seven years later.

B.3 Scoring

We score answers using `gpt-5.4-mini` as an LLM judge. We count an answer as correct only if it identifies both the correct number of victims and their individual names: the nine victims aboard the Demeter, Mrs. Westenra, Mr. Swales, Lucy, and Renfield. Because Mr. Swales' death is more inferential and could be contested, we also evaluated scores when excluding him from the ground truth and found that this does not change the scores of either R3Con or any baseline.

B.4 Parametric Knowledge

Naturally, we expect *Dracula* (Stoker, 1897) to appear in the models' pre-training data and thus be part of their parametric knowledge. We therefore test whether models can answer the Dracula question (Fig. 2) under two settings: ① using parametric knowledge alone, where we ask the same question while specifying *"In Bram Stoker's 1897 novel Dracula, [...]"* to distinguish the novel from other versions; and ② with the entire novel provided in context. We conduct these experiments using `Qwen3.5-35B-A3B` (Qwen Team, 2026), the default model used throughout our experiments.

Results. We find that `Qwen3.5-35B-A3B` fails to answer the question correctly in both settings: when relying on parametric knowledge alone and when provided with the entire novel in context (Tab. 6; see also the raw model responses in Fig. 8).

C Experimental Details

C.1 Benchmarks

C.1.1 CorpusQA

CorpusQA (Lu et al., 2026) is a corpus-level analysis benchmark built around questions whose answers require reading and aggregating information over an entire collection of documents. The benchmark is designed so that answer-critical evidence is highly dispersed and compositional. It covers financial, education, and real-estate domains. We use the one-million-token question split, as evaluating larger splits across all nine baselines exceeds our computational budget (e.g., the 4M split).

Parametric knowledge

```
Based on Bram Stoker's 1897 novel *Dracula*, the count of
people killed by the vampire, either directly or indirectly,
can be broken down into three specific groups: the crew of the
*Demeter*, Lucy Westenra, and Dr. John Seward's patient, Mark
Renfield.

While exact numbers regarding the crew are slightly ambiguous
in the text, the most consistent count provided within the
narrative is **31 people**.

Here is the detailed breakdown:

### 1. The Crew of the *Demeter* (29 Men)
The novel begins with the arrival of the ship *Demeter* at
Whitby. Throughout the voyage, Dracula (traveling in the form
of a dog and later a man inside a crate) slaughtered the crew
to feed.
*   **Location:** On the voyage from Varna to Whitby, and the
ship itself.
*   **Cause of Death:** All crew members were found dead upon
the ship's arrival. The Captain's log entries describe the men
disappearing one by one, and the final log states, "We are all
dead."
*   **Number:** According to Renfield's journal in Chapter 22,
there were **"29 men on board"** the *Demeter* (a figure
generally accepted to include the Captain and his crew). The
Whitby report in Chapter 8 confirms that the entire crew,
including the captain, was found dead.
*   **Status:** **Direct Kills** (Slaughtered by Dracula).

### 2. Lucy Westenra (1 Woman)
*   **Location:** Whitby and London.
*   **Cause of Death:** Lucy was bitten by Dracula on multiple
occasions (directly), which turned her into a vampire. While
she was physically "killed" by being staked through the heart
by Arthur Holmwood and the other vampire hunters, her death
was the direct result of Dracula's transformation of her.
*   **Status:** **Indirect Kill** (Dracula's bite caused the
condition that led to her death; she was killed to prevent her
from harming others).

### 3. Mark Renfield (1 Man)
*   **Location:** London (The Crippled Asylum / Seward's
Asylum).
*   **Cause of Death:** Renfield was Dracula's primary human
contact in London. While he did not die from a physical attack
by Dracula's fangs in the asylum, Van Helsing and the
narrative confirm that Dracula was psychically "feeding" on
Renfield. This psychic draining drained Renfield's vitality,
leading to his death in his asylum bed.
*   **Status:** **Indirect Kill** (Psychic draining and
influence).

### Total Count
**29 (Crew) + 1 (Lucy) + 1 (Renfield) = 31 people.**

**Important Notes on Exclusions:**
*   **Mina Harker:** Although Dracula bit Mina and infected
her with blood, she was never turned and survived the novel.
Therefore, she is not counted as a kill.
*   **Other Characters:** No other minor characters (such as
Carfax servants or patients at the asylum) are confirmed to
have died as a result of Dracula's actions in the text.

**Final Answer:**
In Bram Stoker's novel, Dracula is responsible for the deaths
of **31 people** directly or indirectly. These are:
1.  **29 crew members** of the *Demeter* (slaughtered during
the voyage).
2.  **Lucy Westenra** (turned into a vampire by Dracula and
subsequently staked).
3.  **Mark Renfield** (died due to psychic draining by
Dracula).
```

With corpus in-context

```
Based on the provided documents, Dracula killed or caused the
death of the following people:

**On the Voyage to England (Ship** ***Demeter*):**
*   **Seven (7) Crew Members:** The Log of the "Demeter"
(Document 6) details the disappearance and death of the crew.
Specifically mentioned or implied victims include:
    *   **Petrofsky** (Hand, missing 16 July).
    *   **One man** (missing 24 July).
    *   **One man** (missing 29 July).
    *   **One man of watch and the Steersman** (missing 30
    July).
    *   **The Second Mate** (died/committed suicide 2 August).
    *   **The Captain** (died on board 4 August).
    *   *(Total Crew Start: 9 (5 hands, 2 mates, 1 cook, 1
    captain); Total Crew Dead: 7)*.

**During Stay in England:**
*   **Two (2) Adults:**
    *   **Mrs. Westenra** (Lucy's mother, died of fright/heart
    condition induced by the attack).
    *   **Lucy Westenra** (died 17 September, became an
    Un-Dead).
*   **Several (Unknown Number) Children:**
    *   Multiple newspaper clippings (Documents 12, 25, 35)
    report children being lured away by the "bloofer lady"
    (the Un-Dead Lucy/Dracula) with tiny wounds in the throat.
    *   Document 12 specifically mentions "Another child
    injured" and a child missed from Hampstead Heath.
    *   Document 25 mentions "Several cases" and "two children
    missed last night."
*   **Renfield (Highly Implied):**
    *   Dr. Seward's Diary (Document 14) describes Renfield
    found lying in a "glittering pool of blood" with a broken
    back and bruises, implying death by Dracula's attack
    (though the text focuses on the injury).

**Total Confirmed Dead:** 9 Humans (7 on *Demeter*, 2
Westenras).
**Total Injured/Implied Dead:** Several Children, Renfield.
```

Figure 8: Responses of our default model, `Qwen3.5-35B-A3B`, to the Dracula question (Fig. 2) using parametric knowledge alone (left) and with the full novel provided in context (right). Both responses are incorrect which showcases the difficulty of the Dracula task (see §B.2).

Metrics. We report accuracy and we follow the benchmark's llm-as-a-judge based scoring mechanism (Lu et al., 2026) using `gpt-5.4-mini` which decides whether the answer is correct or not (scoring each example 0 or 1).

C.1.2 LOONG

Loong (Wang et al., 2024a) is a long-context multi-document QA benchmark spanning three professional domains: academic papers, financial reports, and legal cases. Questions fall into four reasoning categories: *spotlight locating*, which requires finding evidence in a single document among similar decoys; *comparison*, which requires locating and comparing values across documents; *clustering*, which requires grouping or integrating evidence from several sources; and *chain-of-reasoning*, which requires multi-hop or temporal reasoning across documents. These categories let us examine not only whether a method handles long context, but also which kind of dispersed reasoning it supports. The benchmark also varies context length by increasing the number and size of real documents in the bundle. This makes Loong useful for studying how performance changes as the amount of necessary context grows. We use the three out of four context sizes (50k-100k, 100k-200k, and 200k-250k tokens) and exclude the smallest (10-50k).

Metrics. We report accuracy using an LLM-as-a-judge and follow the original evaluation prompts (Wang et al., 2024a). The judge assigns a score to each answer, and we count only answers receiving a perfect score as correct as in (Li et al., 2025b). We use `gpt-5.4-mini` as the judge.

C.2 BASELINES

C.2.1 RAPTOR

Method. To address tasks requiring a more holistic understanding of the corpus, Sarthi et al. (2024) propose RAPTOR, which recursively summarizes the corpus. It first segments the input into leaf chunks, embeds them, clusters semantically related nodes, and asks an LLM to summarize each cluster into a parent node. This cluster-and-summarize process repeats bottom-up, producing a tree in which leaves retain local detail and higher layers contain increasingly abstract summaries. At query time, RAPTOR collapses all tree nodes into a single retrieval pool and retrieves across every abstraction level before passing the selected nodes to a reader LLM.

Principles. RAPTOR constructs its textual representation (i.e., its tree of summaries) without knowledge of the downstream task, and is therefore not *task-specific*. Its construction is *enactive*: summaries are recursively constructed from and informed by the corpus as the hierarchy develops. Finally, its structure is fixed in advance to a hierarchy of textual summaries with a retrieval pool, and therefore is not *flexible*.

Evaluation Details. In our experimental evaluation, we replace the local SBERT encoder with the more powerful `text-embedding-3-small` from OpenAI. Furthermore, RAPTOR was proposed and evaluated on benchmarks far smaller than CorpusQA and Loong—for example, QuALITY (Pang et al., 2022), with only a few thousand tokens per example. On our benchmarks, the default 100-token leaf size yields on the order of 10k leaves for a single example, requiring thousands of summary calls and making it impractical to run with large reasoning models. We therefore scale the leaf size to roughly 2,000 tokens, the cluster-split ceiling from roughly 3,500 to 40,000 tokens, the summary length from 100 to 400 tokens, and retrieval from the top 10 to the top 20 nodes. This configuration scales RAPTOR's hyperparameters to the corpus sizes of Loong and CorpusQA so that they are computationally feasible to run.

C.2.2 READAGENT

Method. ReadAgent (Lee et al., 2024) constructs an offline gist memory that provides a compressed view of the corpus while selectively expanding back into the original documents. More specifically, it first divides the document(s) into semantically coherent pages using the model to choose where pages naturally end (pagination). Then, it summarizes each page into a short gist, and keeps the ordered list of gists as a single memory. At test time, the model reads this compact memory and selects a small number of pages to re-open in full which provide the context for answering.

Principles. ReadAgent constructs its textual representation (i.e., its gist memory) without knowledge of the downstream task, and is therefore not *task-specific*. Its construction is not *enactive*: each page

is summarized independently, without further interaction with the corpus. Finally, its structure is fixed in advance to an ordered list of textual gists, and therefore is not *flexible*.

Evaluation Details. CorpusQA is much larger than the benchmarks on which ReadAgent was originally evaluated. In practice, we found that the original hyperparameters result in an impractical number of individual calls to the large reasoning model, so we scale them with the corpus size. Specifically, we scale the page size from roughly 600 to 6,000 words, the pagination thresholds from 280/350 to 2,800/3,500, and the per-gist target from roughly 128 to 640 tokens. This scaling allows ReadAgent to finish within a practical compute budget on CorpusQA. Loong is evaluated using the original hyperparameters.

C.2.3 MEMAGENT

Method. MemAgent (Yu et al., 2026) maintains a task-focused textual memory by processing documents sequentially and giving the model the ability to retain, discard, or condense information. More specifically, it streams the document collection through the model in fixed-size chunks while maintaining a short textual memory. For each chunk, the model receives the question, the current memory, and the current chunk, then rewrites the memory by retaining information it predicts will be useful and discarding the rest. After the final chunk, the answer is generated from the question and the final memory alone.

Principles. MemAgent constructs its textual representation (i.e., its memory) with knowledge of the downstream task, and is therefore *task-specific*. However, its construction is limited in *enactivity*: although its memory evolves as it processes new chunks, it makes only a single linear pass through the corpus, without the ability to revisit or reinterpret earlier information as its understanding evolves. Finally, its structure is fixed in advance to a bounded textual memory, and therefore is not *flexible*.

Evaluation Details. We use the released `RL-MemoryAgent-14B`[4] model, which is based on `Qwen2.5-14B` and trained with RL in their seminal work (Yu et al., 2026).

C.2.4 HIPPORAG2

Method. HippoRAG2 (Gutiérrez et al., 2025) is a retrieval method inspired by associative human memory. It first constructs a knowledge graph from the corpus by extracting entity–relation triples from each passage and connecting semantically similar entities across passages. At test time, it identifies entities and passages relevant to the query and propagates their relevance through the graph using Personalized PageRank, allowing passages connected through intermediate concepts to surface. The top-ranked passages are finally read by an LLM to generate the answer.

Principles. HippoRAG2 constructs its graph representation without knowledge of the downstream task, and is therefore not *task-specific*. Its construction has limited *enactivity*: information extracted from each passage is progressively integrated with previously processed information through the shared graph. However, this interaction does not develop a broader contextual understanding that can be used to revisit and reinterpret earlier information as the corpus is processed. Finally, its structuring is not *flexible*: HippoRAG2 pre-commits to organizing information as a knowledge graph.

Evaluation Details. For embeddings, we use OpenAI's `text-embedding-3-small`. We segment CorpusQA and Loong into approximately 8,000- and 3,000–token chunks, respectively, chosen to keep the number of OpenIE calls computationally feasible for each benchmark.

C.2.5 LINEARRAG

Method. Traditional retrieval-based methods can struggle when tasks require following chains of evidence distributed across many documents. Zhuang et al. (2026) proposed LinearRAG, which first constructs a graph of entities from the corpus together with their relevant passages. At test time, it

[4] `https://huggingface.co/BytedTsinghua-SIA/RL-MemoryAgent-14B`

starts with entities mentioned in the query and progressively discovers relevant entities that extend the reasoning chain, extracting the passages needed to answer the query. Then, it ranks passages with semantic similarity to the query with the importance of the entities they contain (with some further refinements). The top-ranked passages are finally read by an LLM to generate the answer.

Principles. LinearRAG constructs its graph representation without knowledge of the downstream task, and is therefore not *task-specific*. Its construction has limited *enactivity*: information extracted from the corpus is integrated into a shared graph, but this accumulated information does not inform how the corpus is subsequently interpreted or represented. Finally, its structure is fixed in advance to a relation-free graph over entities, sentences, and passages, and therefore is not *flexible*.

Evaluation Details. In our evaluation, we use OpenAI's `text-embedding-3-small` in place of the paper's default sentence encoder.

C.2.6 StructRAG

Method. StructRAG[5] (Li et al., 2025b) dynamically induces a task-specific structure for the corpus. Given the query and document contents, it first selects one of several representations—such as a table, graph, or catalog—suited to the task. It then transforms each document into the selected representation. Finally, a utilizer decomposes the query into sub-questions, extracts relevant information from the structured representations, and combines the results into a final answer.

Principles. StructRAG constructs its representation with knowledge of the downstream task, and is therefore *task-specific*. Its construction is not *enactive* since each document is transformed into the selected structure independently, without further interaction with the corpus. Finally, its structure is *flexible*, as it dynamically chooses among several structures (e.g., table, graph, catalog) based on the task.

Evaluation Details. In our evaluation we increase the maximum generation length from 4,096 to 32,768 tokens to accommodate reasoning models.

C.2.7 CodeAgent

Method. Coding agents (Wang et al., 2024b; Roucher et al., 2025) allow the model to interact symbolically with large corpora and dynamically structure its reasoning through executable code.

Principles. CodeAgent does not construct an intermediate representation of the corpus.

Evaluation Details. We use the `CodeAgent` implementation from `smolagents`, a production-oriented implementation of the CodeAct paradigm (Wang et al., 2024b). To handle long document bundles, we place the documents directly into the executor state.

On CorpusQA, we increase the maximum number of steps from the default of 20 to 60, as we found in initial experiments that the agent often required more than 20 steps. On Loong, we retain the default of 20 steps for the larger `Qwen3.5-35B-A3B`, which did not encounter this limitation. However, for the the experiments with the smaller `Qwen3.5-4B` and `Qwen3.5-9B` models (see §D.1), we increase the limit to 60 steps on Loong as well.

C.2.8 RLMs

Method. Recursive Language Models (RLMs) (Zhang et al., 2025) allow an agent to recursively call itself within a Python environment. These recursive calls enable powerful task-decomposition patterns, including spawning sub-agents for individual sub-tasks.

Principles. RLMs do not construct an intermediate representation of the corpus.

[5] Despite its name, StructRAG is not a RAG-based method.

Table 7: Sampling parameters used for each model, following the recommendations from the official HuggingFace model cards.

Model	Temp.	Top-p	Top-k	Min-p	Pres. pen.	Rep. pen.
`Qwen3.5-4B`	1.0	0.95	20	0.0	1.5	1.0
`Qwen3.5-9B`	1.0	0.95	20	0.0	1.5	1.0
`Qwen3.5-35B-A3B`	1.0	0.95	20	0.0	1.5	1.0

Evaluation Details. We use the official `rlms` package from PyPI, version 0.1.2.

C.2.9 A-RAG

Method. The work of Du et al. (2026) on A-RAG equips an ReAct-style (Yao et al., 2023) agent with three tools: keyword search, semantic search, and full-chunk reading. This allows the model to refine its search strategy as new evidence is discovered and connect multi-hop evidence chains as needed.

Principles. A-RAG constructs its retrieval representation (i.e., passage embeddings and search indices) without knowledge of the downstream task, and is therefore not *task-specific*. Its construction is not *enactive*, as passages are indexed independently without further interaction with the corpus. Finally, its structure is fixed in advance to a retrieval index, and therefore is not *flexible*.

Evaluation Details. In our evaluation, we replace the original `Qwen3-Embedding-0.6B` embedding model with OpenAI's `text-embedding-3-small`.

C.2.10 Claude Code

Method. Claude Code (Anthropic, 2025) is Anthropic's coding agent harness, which allows the model to reason over files through a collection of coding and system tools.

Principles. Claude Code does not construct an intermediate representation of the corpus.

Evaluation Details. For each task, we write the document bundle to a temporary working directory and run the Claude Code CLI in headless mode, instructing the agent to answer using only the provided files. The agent has full permissions within this directory and can interact with the corpus using its standard toolset.

We use Claude Code version 2.1.181 with its default toolset,[6] except for `WebSearch`, `WebFetch`, and `AskUserQuestion`. We disable `WebSearch` and `WebFetch` because the source documents are public, and web access could allow Claude to find the answers online. We disable `AskUserQuestion` because evaluation runs are fully headless.

The remaining toolset comprises of the following available tools: `Bash`, `Read`, `Write`, `Edit`, `NotebookEdit`, `Task`, `TaskCreate`, `TaskGet`, `TaskList`, `TaskOutput`, `TaskStop`, `TaskUpdate`, `ToolSearch`, `Workflow`, `DesignSync`, `Monitor`, `EnterPlanMode`, `ExitPlanMode`, `EnterWorktree`, `ExitWorktree`, `CronCreate`, `CronDelete`, `CronList`, `RemoteTrigger`, `PushNotification`, and `ScheduleWakeup`.

C.3 Models

Tab. 7 lists the sampling parameters used for each model, following the recommendations from their official model cards.

[6] `https://code.claude.com/docs/en/tools-reference`

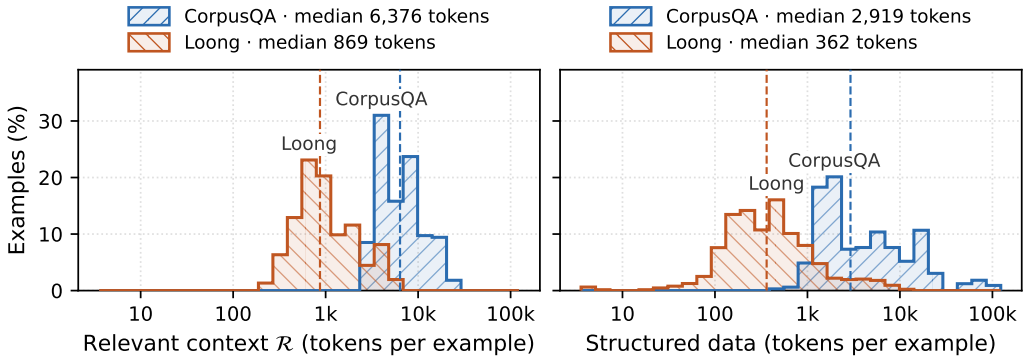


Figure 9: Distribution of the representation components (i.e., final relevant context and structuring) for the two benchmarks with the default `Qwen3.5-35B-A3B` model.

C.4 Infrastructure

Experiments with open-source models are conducted on a cluster of NVIDIA H200 SXM GPUs using vLLM (Kwon et al., 2023) for inference, while closed-source models are accessed via API.

D Detailed Results

D.1 Evaluating Baselines with Smaller Models

In this section, we evaluate the strongest baseline on each benchmark—StructRAG on CorpusQA and CodeAgent on Loong—with the smaller Qwen models (i.e., `Qwen3.5-9B` and `Qwen3.5-4B`) that we use to reduce R3Con's cost. We test whether these baselines can retain competitive performance with smaller models.

Tab. 8 reports results with the smaller `Qwen3.5-9B` and `Qwen3.5-4B` models for R3Con, StructRAG, and CodeAgent. As expected, performance generally decreases with model size. However, R3Con consistently outperforms both baselines across model sizes and benchmarks. Notably, even with `Qwen3.5-4B`, R3Con achieves 49.24% on CorpusQA and 30.23% on Loong, outperforming StructRAG and CodeAgent using the substantially larger `Qwen3.5-35B-A3B` model.

Table 8: Accuracy (↑) across benchmarks and model sizes of the strongest baselines.

Method	CorpusQA	Loong
`Qwen3.5-4B`		
StructRAG (2025b)	29.79%	16.84%
CodeAgent (2025)	10.92%	19.44%
R3Con (ours)	**49.24%**	**30.23%**
`Qwen3.5-9B`		
StructRAG (2025b)	40.43%	20.30%
CodeAgent (2025)	20.13%	22.23%
R3Con (ours)	**54.88%**	**33.54%**
`Qwen3.5-35B-A3B`		
StructRAG (2025b)	50.76%	23.72%
CodeAgent (2025)	23.69%	29.01%
R3Con (ours)	**70.64%**	**37.43%**

D.2 R3Con – Detailed Analyses

D.2.1 Size of Relevance and Structuring

Fig. 9 shows the distribution of representation sizes across individual examples. Median sizes are 6,376 and 2,919 tokens for the final relevant context and structured data on CorpusQA, respectively, and 869 and 362 tokens on Loong. Importantly, both components of the representation are around 1% of the *total* original context. The relevance snippet $\mathcal{R}$ averages 0.83% and 1.30% of the context on CorpusQA and Loong, respectively, while the structured data averages 0.93% and 0.74%.

Structures to Tokens. We measure the size of the structured data by serializing it to compact-JSON and counting tokens using the `Qwen3.5` tokenizer.

D.2.2 Examples of Structures

To make the structured intermediate representation concrete, we show the schema proposed for the running Dracula example in Figure 10. The top-level `Parse` object separates deaths during the voyage to England from deaths during Dracula's stay in England; both collections contain records with the

```python
from pydantic import BaseModel, Field

class DeathEvent(BaseModel):
    """A documented death of someone killed by Dracula during his
    voyage or stay in England."""
    victim_name: str | None = Field(
        description="Name of the deceased, or 'unnamed X' if not specified"
    )
    phase: str | None = Field(
        description="Whether the death occurred on 'voyage_to_england' or 'stay_in_england'"
    )
    cause_of_death: str | None = Field(
        description="How the person died--e.g., 'blood loss', 'fright', 'attack', 'unknown'"
    )
    confirmed_in_document: bool | None = Field(
        description="Whether this document explicitly confirms this person died"
    )
    phase_context: str | None = Field(
        description="Additional context about the death in relation to phase"
    )
    quote: str = Field(
        description="Short verbatim excerpt from the document"
    )

class Parse(BaseModel):
    voyage_victims: list[DeathEvent]
    stay_victims: list[DeathEvent]
```

Figure 10: Pydantic schema proposed by the model for structuring the running-example corpus. This is proposed by the model in Phase 2 of the R3CON (see Fig. 4)

shared `DeathEvent` schema. We also show an example of a more involved hierarchical schema in Fig. 11 for a CorpusQA task.

D.2.3 EXAMPLES OF RELEVANCE SNIPPETS

Fig. 12 shows the full relevance snippets from the running Dracula example corresponding to the abbreviated snippets shown in Fig. 5. As discussed in §4.1, the snippets evolve across rounds as we extract relevant information which provides additional information for re-interpretation of documents.

D.2.4 DOWNSTREAM REASONING

The first two stages of R3CON construct a representation that can be consumed by different downstream reasoning methods. In our default implementation, we use a coding agent that can programmatically explore the structured data without requiring the entire structure to fit within its context window (§4). Here, we study an alternative, simpler reasoning paradigm: directly providing the representation to the LLM and asking it to produce the answer.

Table 9: Accuracy (↑) with different downstream reasoning methods over the representation constructed by R3CON.

Downstream Reasoning	CorpusQA	Loong
LLM	65.43%	**40.17%**
Coding Agent (default)	**70.64%**	37.43%

Tab. 9 compares these two approaches. Neither reasoning paradigm consistently dominates the other. On CorpusQA, the coding agent improves accuracy from 65.43% to 70.64%, while on Loong, direct LLM reasoning improves accuracy from 37.43% to 40.17%. Direct LLM reasoning provides a simpler alternative when the resulting representation can be accommodated in context, while the coding agent provides a scalable mechanism for programmatically exploring larger structured representations.

A case for coding agents downstream. As the context grows and the resulting structured data becomes larger, programmatic reasoning may be more preferable. As shown in Fig. 9, R3CON produces substantially larger structured data on CorpusQA, which we hypothesize contributes to the roughly +5 point improvement over direct LLM reasoning. A coding agent avoids placing the entire structure in the model's context and can delegate symbolic operations to the Python runtime. Finally, programmatic reasoning provides an additional benefit: the generated code offers provenance and transparency on the reasoning process (Theologitis & Suciu, 2026; Ning et al., 2026).

```python
from pydantic import BaseModel, Field
from typing import List, Optional, Literal

class CashFlowPeriod(BaseModel):
    """One cash flow amount for a specific period"""
    period: Literal["March_31_2025", "March_31_2024"] = Field(
        description="The reporting period"
    )
    amount: float = Field(
        description="Cash flow amount as reported in the filing"
    )
    unit_scale: Literal["millions", "thousands", "dollars"] = Field(
        description="Scale of the reported amount"
    )

class CashFlowByActivity(BaseModel):
    """Cash flow data for one activity type across both periods"""
    activity_type: Literal["operating", "investing", "financing"] = Field(
        description="Type of cash flow activity"
    )
    period_2025: Optional[CashFlowPeriod] = Field(
        description="Cash flow for March 31, 2025"
    )
    period_2024: Optional[CashFlowPeriod] = Field(
        description="Cash flow for March 31, 2024"
    )

class CompanyCashFlowRecord(BaseModel):
    """All cash flow data extracted from one company's 10-Q filing"""
    company_name: str = Field(
        description="The company name as named in the filing"
    )
    operating: Optional[CashFlowByActivity] = Field(
        description="Operating cash flow data"
    )
    investing: Optional[CashFlowByActivity] = Field(
        description="Investing cash flow data"
    )
    financing: Optional[CashFlowByActivity] = Field(
        description="Financing cash flow data"
    )

class Parse(BaseModel):
    """All company cash flow records extracted from the document collection"""
    company_cash_flows: List[CompanyCashFlowRecord] = Field(
        description="List of all company cash flow records"
    )
```

Figure 11: Pydantic schema proposed by the model for the CorpusQA task: *"Rank all companies by the volatility of their cash flow sources, measured as the standard deviation of their net cash from operating, investing, and financing activities. List the top 3 companies with the highest volatility."* The schema is proposed in Phase 2 of R3CON (see Fig. 4). Notably, the schema organizes information hierarchically by company, activity type, and reporting period.

D.2.5 EFFECT OF ROUNDS

The number of rounds N provides a knob for trading off performance and cost when extracting relevance. As shown in Fig. 13, with `Qwen3.5-35B-A3B`, performance is already strong at $N = 1$ and remains similar at our default $N = 2$, which costs \$439 on CorpusQA and \$281 on Loong. Increasing to $N = 3$ raises these costs to \$574 and \$361 without consistent improvements in accuracy. We use $N = 2$ throughout our experiments, while noting that the appropriate number of rounds may depend on model capability, with smaller models potentially benefiting from additional rounds. Resolving relevant connections across the context can be challenging, and additional rounds provide more opportunities to progressively identify and refine them.

D.3 QUALITATIVE RESULTS

This section provides the detailed qualitative evidence underlying the analysis in §6.5. Specifically, we show RAPTOR's hierarchical summaries (Fig. 14), StructRAG's task-specific tables (Fig. 15), MemAgent's evolving textual memory (Fig. 16), and RLMs' reasoning trajectory (Fig. 17).

‡Document ordering is arbitrary	Round 1 - Relevance States	Round 2 - Relevance States
(Doc 37‡) Captain's log 	$R_{37}^{(1)}$ *A ship's log maintained by the captain of the Demeter during its voyage from Varna to Whitby. Relevant to the task, this document accounts for the deaths and disappearances of crew members occurring during the sea voyage, attributed to the vampire's presence on board. **Specific victims** listed in the log include: Petrofsky (missing July 16), an unnamed crewman (missing July 24), the Second Mate (missing July 29), a man of the watch and the steersman (missing July 30), another watchman (missing August 2), and the First Mate (who committed suicide by jumping overboard to escape the entity on August 3). The log does not explicitly state the fate of the cook, though he was part of the original crew. **The captain survives the final entry but anticipates his own death.** This source provides the specific count and identities of the crew killed or driven to death during the voyage portion of the task.*	$R_{37}^{(2)}$ *A ship's log kept by the captain of the Demeter during the voyage from Varna to Whitby. Relevant to the task: It provides the detailed, chronological account of fatalities occurring specifically on the **voyage to England**. The log documents the disappearance and **death of the nine crew members** one by one (naming Petrofsky, the Second Mate, and the First Mate, alongside the Steersman and Man of the Watch and several unnamed hands) and **concludes with the Captain's death at the helm**. It establishes the count and identities of victims who died before landing, but it contains no information regarding fatalities during Dracula's subsequent **stay** in England.*
(Doc 12) The Dailygraph 	$R_{12}^{(1)}$ *A newspaper cutting from "The Dailygraph" (August 8–9) reprinted in Mina Murray's Journal, detailing the arrival of the derelict Russian schooner Demeter at Whitby. Relevant to the task: The report confirms the vessel's arrival in England and **identifies the helmsman as a confirmed death,** found lashed to the wheel and dead for two days. It cites the ship's log which lists "facts of missing men," implying additional crew casualties during the voyage to England. It also records the death of a mastiff near the pier. While this document establishes the helmsman as a victim and implies further crew losses, it does not list specific names or provide a definitive count of human victims.*	$R_{12}^{(2)}$ *A newspaper clipping from The Dailygraph (dated 8–9 August), transcribed into Mina Murray's Journal. Relevant to the task: This report details the stormy arrival of the Russian schooner Demeter in Whitby, England, marking the completion of Dracula's voyage. **It confirms the death of the ship's Captain**, found dead and lashed to the helm upon the ship's entry While the article references a ship's log containing "facts of missing men" (implying other crew casualties during the voyage), this specific text does not enumerate the crew's names or provide a total casualty count for the voyage. It contains no information regarding fatalities that occur during the Count's subsequent stay on land.*

Figure 12: Evolution of the relevance snippets for the log of the Demeter (D_{37}) and The Dailygraph newspaper (D_{12}) across two rounds of extracting relevance. For example, the top-left and top-right cells show $R_{37}^{(1)}$ and $R_{37}^{(2)}$, respectively. ‡Document ordering is arbitrary.

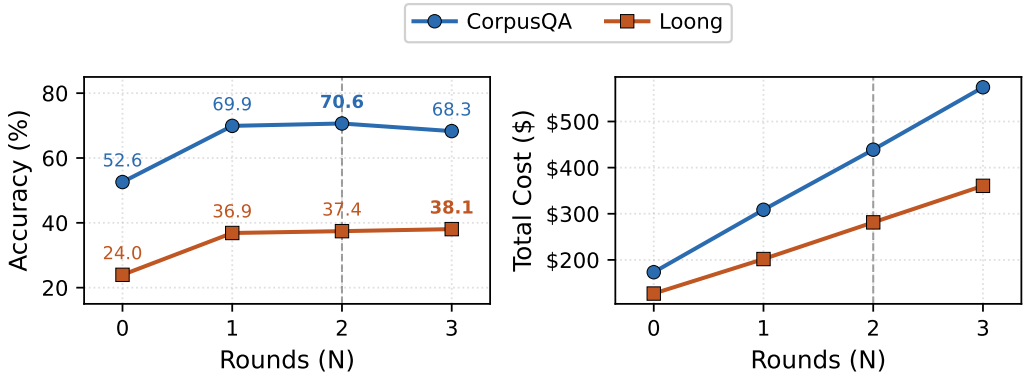


Figure 13: Effect of the number of relevance-extraction rounds N on accuracy and cost.

D.4 Standard Errors for Main Results

Tables 10 and 11 report standard errors across domains and reasoning facets, respectively.

The improvements are large relative to their standard errors. On CorpusQA, R3Con achieves 70.64 ± 2.52 accuracy, compared with 50.76 ± 2.76 for the strongest baseline, a 19.88 percentage-point improvement. On Loong, R3Con achieves 37.43 ± 1.35, compared with 29.01 ± 1.28, an 8.42

Layer 1 – Summary Node #106 – 1st chunk of the *Demeter*'s log

Mina Harker writes Van Helsing (25 Sep) awaiting Jonathan. **The *Demeter* log (18 July-1 Aug) documents Dracula's voyage from Varna to Whitby. He hides in earth-filled boxes, killing the crew; the ship grounds at Tate Hill Pier with the captain dead, lashed to the wheel.**

In October, Van Helsing, Seward, Harker, Morris, and Godalming track Dracula via the *Czarina Catherine* to Galatz. Mina enters hypnotic trances to locate him ("lapping waves"), showing vampiric traits (lethargy, sharper teeth, forehead scar) that terrify Seward, who fears euthanasia. Mina provides train timetables. At Galatz, they board the *Czarina*; Captain Donelson reports the crew died mysteriously under a guiding fog. A box was transferred to Hildesheim, then Petrof Skinsky, who is found dead with his throat torn.

Dracula escaped via the Sereth/Bistritza rivers toward his Carpathian castle. The group pursues in a steam launch. Van Helsing explains Mina's trance links to Dracula, risking his counter-intelligence. They arm with Winchesters, wild rose branches, stakes, and Kukri knives, preparing to stake and decapitate him in daylight. They acknowledge the risks to Mina and themselves, driven to end the evil.

Layer 1 – Summary Node #112 – 2nd chunk of the *Demeter*'s log

The text compiles journal entries, letters, and notes from Bram Stoker's *Dracula*. **On 3 August, the captain of the *Demeter* records the crew dying in thick fog. The mate, hallucinating a "tall, thin, ghastly pale" figure in the bow, stabs empty air before leaping overboard. The captain vows to tie himself to the helm to protect the vessel and honor.**

In Whitby, Mina Murray writes to Lucy Westenra about Jonathan's return from Transylvania. She practices shorthand to assist him. At Whitby, old sailor Swales scoffs at local ghost legends. Mina transcribes Jonathan's journal; Jonathan sees a young Dracula on Piccadilly. Mr. Hawkins dies, leaving a fortune to Jonathan, burdening him. Lucy Westenra and her mother die following Van Helsing's care.

The group chases Dracula's coffins on a Romanian river. Godalming manhandles the launch while Mina writes. When the launch breaks, they switch to horses, fearing for Mina. Letters confirm Lucy's engagement to Arthur Holmwood and Dr. Van Helsing guarding her flowers.

Seven years later, the group revisits Transylvania; the castle remains. They conclude no proof is needed to believe their deeds. Mina and Jonathan name their son Quincey after the fallen Morris, born on his birthday. Van Helsing notes the boy will know his mother's bravery. Solicitors confirm Count de Ville bought Carfax via London agents, sending fifty boxes invoiced as "clay" to Purfleet.

Figure 14: RAPTOR (Sarthi et al., 2024) iteratively clusters and summarizes passages, forming a hierarchical tree with increasingly abstract summaries. At Layer 0, we have raw chunked text from the corpus. The log of the Demeter gets chunked into two pieces that end up in different clusters at Layer 1 and are therefore summarized separately. We show these two summaries above. Notably, the summarization happens task-agnostically, so the model is not inclined to retain minor details such as background crew members dying or how many crew members the Demeter sailed with. Indeed, as illustrated above, these details disappear from the summaries, and even more so as we go higher in the tree. Subsequent reasoning cannot recover this information, and RAPTOR therefore fails to identify the deaths and answer the task in Fig. 2.

Structure built from the log of the Demeter – Verbatim

Table Title: Crew Disappearances and Deaths During Voyage to England (Varna to Whitby)
Source: Raw Content (Demeter's Log, 6 July - 4 August)

Date	Crew Member	Status/Event	Details/Notes
16 July	Petrofsky	Missing	Reported missing. Taken larboard watch eight bells last night; relieved by Abramoff, did not go to bunk.
24 July	Unknown Crew Member	Disappeared	"Already a hand short... another man lost–disappeared." Came off watch, not seen again.
29 July	Second Mate	Missing	Found no one except steersman when morning watch came on deck. Crew in a panic.
30 July	Man of Watch	Missing	Found no one except steersman. Later text says "both man of watch and steersman missing". Only self, mate, and two hands left.
30 July	Steersman	Missing	Found no one except steersman... later text confirms missing.
2 August	Man on Watch	Missing	Mate heard cry, ran on deck, no sign of man on watch.
3 August	First Mate	Died (Disappeared/Threw himself)	Became "raging madman", claimed "He" is there, said "The sea will save me from Him", then "sprang on the bulwark and deliberately threw himself into the sea".

Status of Inquiry	Information Available	Limitation
Killer Identity	Text refers to "strange man", "It", "Him", "madman". Name "Dracula" not explicitly mentioned.	Query asks about "Dracula", text implies but does not name him.
Location Scope	Voyage only (Varna to Whitby).	Query asks about "stay there" (England). Raw Content ends with Captain alone on ship in fog.
Total Deaths Mentioned	**At least 7 crew members/Officers (Petrofsky, 2 men lost, 2nd Mate, Steersman, Man of watch, Mate).**	**Captain remains alone as of 4 August.**

Structure built from the newspaper The Dailygraph – Verbatim

Table Title: Details of Deceased Found on the Schooner *Demeter*
Source: From a Correspondent (*Whitby*)

Name/Description	Status	Details (Voyage & Findings)
Dead Steersman / Seaman	Deceased	Found lashed to the helm of the schooner with hands tied over one another. A crucifix and beads were fastened around wrists and wheel. Cords had cut flesh to bone. Declared dead for quite two days by Surgeon J. M. Caffyn.
Crew Members	Missing	Log-book of the *Demeter* was in order up to within three days, but contained "facts of missing men". No specific names or count provided in the article.
Total People Killed / Confirmed Dead	Unspecified	The document does not contain the name "Dracula", nor a confirmed total count. Only the dead steersman is explicitly detailed as found.

Figure 15: StructRAG (Li et al., 2025b) creates task-specific structures for each document separately. For this task, the model chooses a table structure. For the log of the Demeter, it proposes and populates the schema shown in the first box. Notably, at the end of the log, the captain is still alive. His death can only be established by connecting his final entry, where he plans to tie his *"hands to the wheel"* until his *"strength begins to fail"*, with the separate newspaper account in The Dailygraph, which reports a dead seaman found at the helm. We must therefore infer that this dead seaman is the captain. However, there is no enactive mechanism during representation construction to make this connection as broader understanding develops. Worse, the structured representation of the log does not retain the captain's plan to tie himself to the wheel, making the connection almost impossible for downstream reasoning to recover. StructRAG therefore misses the captain's death (among others).

Table 10: Accuracy (% ↑) across real-world domains requiring reasoning over large document corpora. R3CON achieves the highest accuracy across all domains, outperforming the strongest baseline by +19.88 pp on CorpusQA and +8.42 pp on Loong overall. Overall accuracy is micro-averaged. We report standard errors (SE) as ± values.

Method	**CorpusQA** (Lu et al., 2026)				**Loong** (Wang et al., 2024a)			
	Edu.	**Finance**	**Real-estate**	**Overall**	**Paper**	**Finance**	**Legal**	**Overall**
RAPTOR (2024)	8.00 ± 3.13	11.11 ± 2.47	17.95 ± 4.35	12.06 ± 1.84	0.00 ± 0.00	22.17 ± 1.71	5.71 ± 1.24	12.12 ± 0.92
ReadAgent (2024)	42.68 ± 5.46	37.42 ± 3.79	31.25 ± 5.18	37.23 ± 2.68	15.15 ± 2.08	35.33 ± 2.03	14.93 ± 1.95	24.49 ± 1.25
MemAgent (2026)	7.32 ± 2.88	7.83 ± 2.09	8.64 ± 3.12	7.90 ± 1.49	0.00 ± 0.00	26.80 ± 1.81	4.00 ± 1.05	13.63 ± 0.96
HippoRAG2 (2025)	7.32 ± 2.88	7.83 ± 2.09	20.99 ± 4.52	10.94 ± 1.72	0.32 ± 0.32	17.42 ± 1.55	6.29 ± 1.30	10.06 ± 0.85
LinearRAG (2026)	8.54 ± 3.09	6.63 ± 1.93	16.05 ± 4.08	9.42 ± 1.61	0.61 ± 0.43	9.38 ± 1.19	5.43 ± 1.21	6.04 ± 0.67
StructRAG (2025b)	53.66 ± 5.51	55.42 ± 3.86	38.27 ± 5.40	50.76 ± 2.76	19.70 ± 2.19	32.66 ± 1.92	12.32 ± 1.76	23.72 ± 1.19
CodeAgent (2025)	14.63 ± 3.90	26.99 ± 3.48	26.25 ± 4.92	23.69 ± 2.36	13.80 ± 1.91	40.07 ± 2.01	24.26 ± 2.33	29.01 ± 1.28
RLMs (2025)	13.41 ± 3.76	37.58 ± 3.77	20.51 ± 4.57	27.38 ± 2.47	12.77 ± 1.84	41.67 ± 2.03	22.77 ± 2.25	28.96 ± 1.28
A-RAG (2026)	8.54 ± 3.09	7.23 ± 2.01	18.52 ± 4.32	10.33 ± 1.68	0.61 ± 0.43	18.43 ± 1.59	2.86 ± 0.89	9.55 ± 0.82
R3CON (ours)	**75.61 ± 4.74**	**82.42 ± 2.96**	**41.25 ± 5.50**	**70.64 ± 2.52**	**27.58 ± 2.46**	**50.59 ± 2.05**	**24.29 ± 2.29**	**37.43 ± 1.35**

Table 11: Accuracy (% ↑) across Loong's reasoning facets. We show standard errors as ± values.

Method	**Spot.**	**Comp.**	**Chain**	**Clust.**	**Overall**
RAPTOR (2024)	51.31 ± 3.62	16.25 ± 2.38	2.04 ± 0.82	1.54 ± 0.54	12.12 ± 0.92
ReadAgent (2024)	60.11 ± 3.67	33.77 ± 3.13	18.49 ± 2.27	10.70 ± 1.40	24.49 ± 1.25
MemAgent (2026)	42.13 ± 3.52	35.00 ± 3.08	0.64 ± 0.45	0.95 ± 0.42	13.63 ± 0.96
HippoRAG2 (2025)	46.70 ± 3.55	13.33 ± 2.19	0.33 ± 0.33	0.38 ± 0.27	10.06 ± 0.85
LinearRAG (2026)	32.49 ± 3.34	3.75 ± 1.23	0.64 ± 0.45	0.38 ± 0.27	6.04 ± 0.67
StructRAG (2025b)	44.16 ± 3.54	30.83 ± 2.98	24.76 ± 2.45	12.19 ± 1.43	23.72 ± 1.19
CodeAgent (2025)	63.54 ± 3.47	45.26 ± 3.27	22.80 ± 2.39	12.90 ± 1.46	29.01 ± 1.28
RLMs (2025)	61.14 ± 3.51	45.15 ± 3.23	19.68 ± 2.26	15.27 ± 1.57	28.96 ± 1.28
A-RAG (2026)	35.03 ± 3.40	14.58 ± 2.28	1.28 ± 0.64	2.65 ± 0.70	9.55 ± 0.82
R3CON (ours)	**65.99 ± 3.38**	**50.00 ± 3.23**	**33.65 ± 2.68**	**23.30 ± 1.84**	**37.43 ± 1.35**

point improvement. In both cases, the performance gap is substantially larger than the standard errors, giving high confidence in the observed improvements.

E PROMPTS

We provide the prompts used in each stage of R3CON. Fig. 18 shows the prompt for extracting relevant context through task-focused summarization. Fig. 19 shows the prompt for proposing the schema, and Fig. 20 shows the prompt for populating that schema from each document. Finally, Fig. 21 shows the prompt used by the coding agent to reason over the resulting representation and synthesize the final answer.

Memory – Round #2.

The provided section does not contain any direct information about Dracula's killings during his voyage to England or his stay there. However, it does offer some context about Dracula's movements and activities leading up to his arrival in England. Specifically, it mentions that Dracula arrived in Galatz and arranged for the transportation of a box containing himself via Slovak traders. The section also notes that Dracula might have murdered his agent, Skinsky, to blot out his traces.

Given the content, the updated memory should include the following:

- The log of the "Demeter" describes the mysterious disappearance of several crew members during the voyage, including Petrofsky, Olgaren, and others, which the crew attributes to a "strange man" or "It" aboard the ship.

- The captain and remaining crew members express fear and superstition regarding these disappearances.

- The log ends with the captain's determination to tie himself to the wheel to avoid being harmed by the entity responsible for the crew members' disappearances.

- The text does not explicitly state that Dracula killed these individuals or provide a list of people killed by Dracula during his voyage or stay in England.

- Dracula arrived in Galatz and arranged for the transportation of a box containing himself via Slovak traders.

- Dracula might have murdered his agent, Skinsky, to blot out his traces.
Updated memory:

- The log of the "Demeter" describes the mysterious disappearance of several crew members during the voyage, including Petrofsky, Olgaren, and others, which the crew attributes to a "strange man" or "It" aboard the ship.

- The captain and remaining crew members express fear and superstition regarding these disappearances.

- The log ends with the captain's determination to tie himself to the wheel to avoid being harmed by the entity responsible for the crew members' disappearances.

- The text does not explicitly state that Dracula killed these individuals or provide a list of people killed by Dracula during his voyage or stay in England.

[...]

Memory – Round #3.

The provided section does not contain any direct information about Dracula's killings during his voyage to England or his stay there. Instead, it focuses on the interactions between Mina Harker and various characters, particularly Dr. Seward, as they deal with the supernatural events surrounding Lucy and Count Dracula.

However, the section does mention Mina Harker's journal entries, which could potentially contain more information about Dracula's activities and the deaths related to him. Given the context, **it is important to note that the journal entries might provide additional insights into the mysterious disappearances of the crew members on the "Demeter" and the overall impact of Dracula's presence in England.**

Updated memory:

- The log of the "Demeter" describes the mysterious disappearance of several crew members during the voyage, including Petrofsky, Olgaren, and others, which the crew attributes to a "strange man" or "It" aboard the ship.

- The captain and remaining crew members express fear and superstition regarding these disappearances.

- The log ends with the captain's determination to tie himself to the wheel to avoid being harmed by the entity responsible for the crew members' disappearances.

- The text does not explicitly state that Dracula killed these individuals or provide a list of people killed by Dracula during his voyage or stay in England.

[...]

Memory – Round #19.

The provided section of Dr. Seward's diary does not directly address the number of people Dracula killed during his voyage to England or his stay there. [...]

- The section describes the group's encounter with Dracula at Carfax, where they attempt to stop him but fail to capture him. Dracula escapes after a brief confrontation.

- Mina Harker begins to show signs of transformation into a vampire, similar to what happened to Lucy Westenra. This transformation is attributed to Dracula's bite and the subsequent "Vampire's baptism of blood."

- Dr. Van Helsing and Dr. Seward discuss the possibility that Dracula could compel Mina to reveal information to him through hypnosis, which adds to their concerns about her safety and the broader threat posed by Dracula.

- The text does not provide a list of people killed by Dracula during his voyage or stay in England but emphasizes the ongoing threat and the group's efforts to protect Mina and stop Dracula.

- The mysterious disappearances of the crew members on the "Demeter" remain unexplained in terms of direct attribution to Dracula, but the text suggests that these events are related to his presence and activities in England.

[...]

Figure 16: MemAgent (Yu et al., 2026) gradually forgets the deaths aboard the Demeter. MemAgent processes the corpus sequentially. In round 2, it records the mysterious deaths of the Demeter's crew, but fails to attribute them to Dracula. As it continues through the corpus, these deaths remain "unexplained" (round 19) and gradually disappear from its memory. By the final chunk (round 42), they are no longer retained at all which, meaning that the LLM reasoning over this textual representation no longer has the information needed to answer the task in Fig. 2.

Turn 1 – REPL code

```
print("Context length:", len(context))
print(context[:5000])
```

Turn 1 – stdout

```
=== Document 6 ===
LOG OF THE “DEMETER.”
[...]
On 6 July we finished taking in cargo, silver sand and boxes of earth.
At noon set sail. East wind, fresh. Crew, five hands ... two mates,
cook, and myself (captain).
```

Turn 5 – REPL code

```
headers = ['LOG OF THE "DEMETER."', 'LOG OF THE DEMETER',
'Log of the Demeter', 'DIARY OF THE DEMETER']
for header in headers:
idx = context.find(header)
if idx != -1:
print(f"Found '{header}' at index {idx}") [...]
else:
print("Header not found. Searching for date patterns in voyage...")
date_patterns = re.findall(r'On \d+ [A-Z][a-z]+ [...]', context, re.DOTALL)
```

Turn 5 – stdout

```
Header not found. Searching for date patterns in voyage...
Found 1 date-stamped entries
— Entry 1 —
On 6 July we finished taking in cargo, silver sand and boxes of earth.
At noon set sail. East wind, fresh. Crew, five hands ... two mates,
cook, and myself (captain).
```

Turn 12 – REPL code

```
demeter_log_start = context.find('On 6 July we finished taking in cargo')
demeter_log_section = context[demeter_log_start : demeter_log_start+10000]
demeter_query = """Based on the Log of the Demeter below, list:
1. How many crew members died during the voyage
2. Their names if mentioned  3. How they died [...]
Context:""" + demeter_log_section
result = llm_query(demeter_query, "List all Demeter crew deaths specifically")
```

Turn 12 – sub-model reply

1. How many crew members died during the voyage
8 crew members (5 hands, 2 mates, and 1 cook). *Reasoning:* The log begins stating the crew consisted of “five hands ... two mates, cook, and myself (captain).” **By the end of the voyage (4 August), only the Captain remains alive.** Therefore, 8 crew members were lost.

[...]

Turn 29 – REPL code

```
demeter_crew_deaths = """
DEMETER VOYAGE DEATHS (8 crew members killed):
[...]
8. One unnamed remaining crew member - Presumed lost before ship arrived

ENGLAND VICTIMS (3 confirmed deaths):
[...]
TOTAL: 11 people killed directly or indirectly by Dracula
- 8 on the Demeter voyage
- 3 in England (Mrs. Westenra, Lucy Westenra, Renfield)
"""
print(demeter_crew_deaths)
# Set the answer
answer["content"] = f"""Based on the text, Count Dracula killed the following
people during his voyage to England on the Demeter [...]"""
answer["ready"] = True
```

Figure 17: RLMs (Zhang et al., 2025) place the corpus inside a Python REPL allowing the agent to interact with it via code and also spawn sub-agents of itself. Here, the agent searches for the Demeter log using local string and regex queries, then asks a sub-agent to reason over the extracted section. Since this section does not include the later newspaper account from The Dailygraph, the sub-model incorrectly concludes that the captain survived the voyage. This conclusion is carried forward and ultimately causes the agent to miss captain’s death. It also fails to recover Mr. Swales’ death due to a similar pattern. Both deaths require connecting evidence scattered across the corpus, yet the agent’s fragmented search strategy never brings the necessary evidence together.

```
You are given one document from a collection, and a single **task** (between `<task>` and `</task>` tags) that the whole collection will be used to answer. Write a **task-conditioned summary** of this one document: what does it contribute toward the task?

Produce a summary, **not an answer**. Even when the task is phrased as an instruction ("list. . . ", "rank. . . ", "construct the chain. . . ", "how many. . . "), do not try to carry it out or commit to a final answer here — you are looking at only one document, and the answer is assembled later from all the summaries together. Just record, faithfully and in your own words, what THIS document offers toward the task; do not invent.

Open with a brief, high-level orientation of what this document broadly is (e.g. "a municipal budget memo", "a 10-Q filing", "a chapter of a noir novel", "a government press release") — this latent, big-picture read helps the collection be understood as a whole — then note the task-relevant specifics. Keep it tight: a larger document still gets a short summary; include only what is durable and bears on the task.

<task>
{{ task }}
</task>

## Task-conditioned summaries of the other documents (previous round)

Below are the other documents' task-conditioned summaries for this same task, from the previous round. Use them to read THIS document in light of the collection — a document often only becomes useful and understandable for the task once you know information from the other documents. Still write only about THIS document (the one the user provides).

{{ other_summaries }}

## Output
Reply with the task-conditioned summary of this document only — no preamble, no restating the task. If this document contributes nothing relevant to the task, it is fine to say so briefly (or reply with nothing).

Examples:

Input:
<task>
Across the candidates' platform papers, who proposes the largest cut to the city transit budget?
</task>
Document:
[Councilmember Dana Reyes — "A Leaner, Fairer City": Platform on City Services]

Introduction. Our city stands at a crossroads. After three years of rising costs and flat revenue, the next council must make hard choices about where every dollar goes. This paper lays out my priorities across transit, public safety, parks, and city administration.

Public safety. I will protect frontline policing and fire services from cuts; response times have already slipped and residents feel it.

Transit. The current transit budget stands at $80 million. In my view it has grown bloated by underused routes — three of our crosstown lines run at less than 20% capacity outside rush hour. I propose reducing the transit budget by $12 million and redirecting those savings to road maintenance, where our backlog of potholes and failing bridges is now a genuine safety issue. To be clear, this is a reallocation, not an attack on riders; the core commuter lines stay funded.

City administration. I will impose a hiring freeze at City Hall and cut the outside-consulting budget by 30%.

Parks and quality of life. I support a modest downtown bike-lane pilot, funded from the parks capital reserve — not from transit.

Closing. These choices reflect one principle: spend where residents see the result.

Output (task-conditioned summary):
A candidate's campaign platform paper (Councilmember Dana Reyes). Relevant to the task: Reyes proposes cutting the city transit budget by $12M (from $80M), reallocating it to road maintenance. Her hiring freeze, consulting cut, and bike-lane pilot are not transit cuts. As just one candidate's paper, it gives only Reyes's proposed cut — not enough on its own to say who proposes the largest.
```

Figure 18: Prompt for extracting relevance via task-focused summarization (Dang, 2005). It is used in Phase 1 of R3CON as illustrated in Figure 4. Text enclosed in double curly braces denotes a template variable populated at runtime.

```
You design a Pydantic schema that captures exactly the information needed to complete a single **task** over a collection of documents you have not seen. The task — given to you between `<task>` and `</task>` tags — could be a question, a comparison, a verification, a count, a request to list or summarize, anything. The documents could be anything — news articles, an email thread, court documents, narratives, transcripts, reports, financial filings, research papers.

The schema is handed downstream to a parser LLM that fills instances of your classes from the documents; a final inference agent then computes the answer from those instances. You don't need to think about those downstream stages — just design the schema that captures the right information.

Use `Literal[...]` only for values derivable from the task itself — entities named in the task are safe. Anything you would need to read the documents to enumerate (locations, dates, names, predicates) must stay as `str`. Model only what is required to compute the answer.

Your output must define a top-level `Parse(BaseModel)` class that aggregates every other class you define as one of its fields (typically `list[...]` of each).

## Task-conditioned document summaries

In order to know what the documents contain, we have summarized each document for you upstream in a task-specific way: we summarize the useful information learned from each document so you have enough understanding to propose a good schema. Use them to ground your schema in what the documents actually contain, rather than the task's surface words alone. (Do not bake any answer into the schema; just let the summaries inform which information is worth a field.)

{{ summaries }}

## Output format

Reply with two parts:

1. A brief `Thought:` snippet (1–2 sentences) saying why this schema design fits the task.
2. A `<schema>...</schema>` block containing only Python source. The runtime extracts the body between the tags and `exec`s it.

## Rules

**The parser sees ONE document at a time.** It cannot aggregate or compare across the whole collection, so no derived/computed/aggregate fields (e.g. `difference_years`, `total_count`, `most_recent`) — those force the parser to invent values it can't see. Put extractable facts as rows in `list[...]` fields; the inference agent aggregates downstream.

**Every top-level field of `Parse` must itself be a `list[...]`.**

**Model the underlying domain, not the task's surface form.** The task may rest on a premise that doesn't hold in the text. If the task is *"How does X react when Y does Z?"* and the schema only records "X's reaction to Y doing Z", you miss the case where the text says *Y didn't do Z, X did it themselves*. Model the underlying event class (e.g. *"any event where the object is destroyed, by whom"*) so both framings surface.

**Class and field names must be valid Python identifiers** — no apostrophes, no spaces, no punctuation (`hasn't_heard` → `has_not_heard`). The schema is `exec`'d as Python.

Examples:

Input:
<task>
How many of these companies reported a net loss for their most recent quarter?
</task>

Task-conditioned document summaries:
### Document 1
Aeva Technologies 10-Q — bottom line is a "Net loss" of $(43.2) million for the quarter.
### Document 2
ClearOne 10-Q — the line is labeled "Net income (loss)", reported as $(1,196) thousand.
### Document 3
Cross Timbers Royalty Trust 10-Q — reports positive "Distributable income" of $4.1 million.

Thought: "How many . . . net loss" is a count computed downstream — the parser sees one filing at a time and cannot count across them. Each filing labels the bottom line differently (the summaries tell us which), so capture each company's net result as a signed amount plus its as-written label and unit; inference counts how many are negative.

<schema>
from pydantic import BaseModel, Field

class CompanyNetResult(BaseModel):
    """The bottom-line net result reported in one company's filing for the period in question."""
    company: str | None = Field(description="The company the filing belongs to, as named")
    label: str | None = Field(description="The line-item label as written (e.g. 'Net loss', 'Net income (loss)')")
    net_result: float | None = Field(description="The amount as reported; a loss is negative")
    unit_scale: str | None = Field(description="How the amount is expressed: 'dollars', 'thousands', 'millions'")
    period: str | None = Field(description="The period the figure covers")
    quote: str = Field(description="Verbatim line containing the label and amount")

class Parse(BaseModel):
    company_net_results: list[CompanyNetResult]
</schema>

[... additional in-context examples omitted ...]
```

Figure 19: Prompt for proposing a schema. It is used in the first stage of Phase 2 of R3CON as illustrated in Figure 4. Text enclosed in double curly braces denotes a template variable populated at runtime.

```
You read a document and extract structured instances according to a Pydantic schema. The schema was designed to capture exactly the information needed to complete a specific task — be exhaustive and faithful to it.

## Task the schema was designed for
<task>
{{ task }}
</task>

## Pydantic schema

```python
{{ schema_code }}
```

Task-conditioned document summaries (background to help you — do NOT extract from these)
To help you extract well, each document in the collection has already been summarized upstream in a task-specific way. Together they give you a *latent* understanding of what the whole collection contains — context you would otherwise be missing while looking at a single document in isolation, which helps you resolve references and extract more completely and accurately. These are provided only as help; the actual document to extract from is the one the user provides next. Extract only from that document, never from these summaries.

{{ summaries }}

Your job
Populate `Parse` with every instance found in the document below (the user message). Resolve pronouns to their named entities. An empty list is correct when the document has no instances.

Examples:

Input:
<task>
Which employees were dismissed from the company?
</task>

Schema:

```python
class Dismissal(BaseModel):
    person: str = Field(description="the dismissed employee, resolved to their full name")
    quote: str = Field(description="the verbatim line showing the dismissal")
class Parse(BaseModel):
    dismissals: list[Dismissal] = []
```

Task-conditioned document summaries:
Document 1
An internal HR memo about a post-audit restructuring; "Reyes" is the analyst Dana Reyes, not the board chair who shares the surname.

Document:
After the audit, the board moved quickly. Reyes was let go on Friday, her badge deactivated by noon. Park, seeing which way the wind blew, handed in her resignation the same afternoon. A memo noted that Okafor's position had been eliminated in the restructuring.

Thinking: "Let go" and "position eliminated" are dismissals in different words than the task's — extract both, each anchored by a quote. The summaries fix which Reyes is meant, so the name resolves to Dana Reyes. Park resigned: leaving voluntarily is a different event, so skip it.

Parse:

```json
{"dismissals": [
  {"person": "Dana Reyes", "quote": "Reyes was let go on Friday,"},
  {"person": "Okafor", "quote": "Okafor's position had been eliminated in the restructuring."}
]}
```

[... additional in-context example omitted ...]
```

Figure 20: Prompt for populating the proposed schema per-document. It is used in the second stage of Phase 2 of R3CON as illustrated in Figure 4. Text enclosed in double curly braces denotes a template variable populated at runtime.

```
You are an expert assistant who solves a task by working in a loop of Thought, Code, and Observation steps. At each step you write a short `Thought:` explaining your reasoning, then one `<code>...</code>` block of simple Python. The runtime executes the code and feeds whatever you `print(...)` back to you inside `<observation>...</observation>` tags on the next step. You commit your final answer by calling `final_answer(...)` from inside a `<code>` block.

## The task and your two views

You answer a task about a collection of documents using two views of them, both laid out for you below and both reliable: the per-document **task-conditioned summaries**, and a structured **`parse`** of the facts extracted from those documents (shown in full below, and also bound as the `parse` variable in your sandbox). Neither view is more authoritative than the other — read both, and cross-check them where they overlap; when one is thin or silent on a point, lean on the other. Both views are already in front of you, so the answer is very often already there to be read.

## Execution environment

Your code runs in a sandboxed in-process interpreter. The variable `parse` is **already bound** to the parsed dict shown below. Variables you bind in one step (e.g. `rows = [...]`) persist on the next, so you can build state across steps. Allowed imports: `collections`, `datetime`, `itertools`, `math`, `queue`, `random`, `re`, `stat`, `statistics`, `time`, `unicodedata`.

## Document summaries

A short, per-document summary of what each document contributes toward the task, headed by the document it describes (**Document N**). This is a co-equal view, not just background: it often carries context the parse compresses away, and sometimes the very fact the task needs. Read it together with the parse.

{{ summaries }}

## The parsed data (`parse`)

Structured records of the concrete facts extracted from the documents — the exact dict already bound as the `parse` variable in your sandbox. Each record's `document` field is the number of the document it came from (the same **Document N** as the summaries). When the task asks you to identify a document, refer to it the way the question and the documents themselves do.

{{ parse_block }}

## Rules

1. **Default to reading the parse you're shown — don't transform it.** The parse is printed above; if the answer is already there, report it as it appears, in the same units and figure the document uses. Don't convert, rescale, re-aggregate, or reconstruct a value unless the task explicitly asks for it.
2. **Category and identity questions are reading calls, not string equality.** A category or description in the task is satisfied by its specific instances, even when no field repeats the task's exact words. Decide by reading and judging, not by matching strings.
3. **Avoid empty answers.** Prefer pragmatic and informative answers the evidence supports over "none", an empty list, or "cannot be determined".
4. **Use code for what code is actually for** — counting, aggregating, arithmetic or date math the task asks for, sorting, deduping, and sifting many records to find the ones that match — and only there. However, if you have the ability to read and reason over the records directly (e.g., they aren't too many) then prefer that.

## How to work

Each step is one short `Thought:` and one `<code>...</code>` block. When the answer is already clear from what you've been shown, that block can simply be your `final_answer(...)` call — you don't need exploratory code first. Use only variables you've defined (plus the already-bound `parse`), and import only from the allowed modules above. `print(...)` is the only way to carry information forward — its output returns as the next step's `<observation>`.

---

The user will give you a task wrapped in `<task>...</task>` tags. In the examples below, the `parse` shown under each task is exactly what would be bound to your `parse` variable.

Examples:

Input:
<task>
How much did the Helmsworth Foundation grant to the literacy program?
</task>

## Document summaries
### Document 1
The Helmsworth Foundation annual report — states a grant of $1.25 million to the community literacy program.
### Document 2
A program brochure for an unrelated arts initiative; no grant figures.

`parse`:
{
  "grants": [
    {"document": 1, "recipient": "community literacy program", "amount": 1.25, "unit": "millions"},
    {"document": 2, "recipient": "youth arts initiative", "amount": null, "unit": null}
  ]
}

Thought: The figure is already stated in the parse: a grant of $1.25 million to the literacy program. It's right in front of me, so I'll report it the way the report gives it — "$1.25 million" — rather than rescaling it to 1,250,000, which isn't the form the document uses.
<code>
final_answer("The Helmsworth Foundation granted $1.25 million to the community literacy program, as reported in its annual report.")
</code>

[... additional in-context examples omitted ...]
```

Figure 21: Prompt for the coding agent loop in the final step, Phase 3, of R3CON as illustrated in Figure 4. Text enclosed in double curly braces denotes a template variable populated at runtime.