

We appreciate the reviewers' detailed feedback and have substantially revised the paper to address the main concerns around clarity, novelty positioning, and empirical validation. Below is a brief summary of the changes and how they respond to specific points.

## 1. Clarifying the technical novelty of HSEQ vs. prior unified / structural RAG

To address the concern that the unified interface “lacks novelty” and is not clearly positioned versus structural RAG:

We added a dedicated paragraph “**Structural and unified RAG interfaces**” in the Related Work section, grouping prior structural / graph-based / hierarchical RAG systems and explicitly contrasting them with HSEQ.

Our HSEQ has three technical differences:

1. **Reversible, modality-preserving representation:** HSEQ keeps modality-specific structure (text / table / KG) and alignment (offsets, coordinates, IDs) in the segment schema, instead of collapsing everything into a single opaque graph index.
  2. **Segment schema as first-class interface:** The unified segment format (id, level, parent, content, metadata) is used as the *only* interface to the controller, enabling structure-aware navigation and sufficiency-based stopping, rather than merely adding structure for re-ranking or expansion.
  3. **Formal properties and sufficiency-based control:** We provide a precise definition of faithful linearization and budgeted selection, plus a sufficiency-based iteration loop with guaranteed halting under fixed budgets, which is not present in prior structural RAG work.
- These distinctions are now clearly articulated in Related Work and referenced in the Method and Properties discussions.

## 2. Improving clarity and structure of the methodology

Section 4, Using HSEQ on LLMs are rewritten.

## 3. Efficiency trade-off: LLM-only vs. ToG vs. HSEQ

Reviewers requested analysis of the efficiency trade-off between LLM-only, graph-based methods (e.g., ToG), and HSEQ. We added analysis

## 4. Fairer comparison on TAT-QA and baseline consistency

One reviewer pointed out that TAT-LLM in our original Table 2 was reported with a different base model (LLaMA 2) compared to our HSEQ heads:

- We re-ran **TAT-LLM** using **Qwen3-4B-Instruct-2507** as the base model to better match the capacity of our open-source backbones.
- The updated **TAT-QA** numbers for the re-run are:
  - **Acc = 73.1, F1 = 81.0** for TAT-LLM (Qwen3-4B).

- Tables and text have been updated to reflect these new numbers, and we now clearly state which backbone is used for each baseline.
- HSEQ still outperforms this stronger TAT-LLM baseline variant on TAT-QA, but now under a more comparable model setting.

## **5. Example added on HSEQ-A**

Examples are added on HSEQ-A to support the design choices.

Overall, the revision aims to:

1. Clarify and sharpen the technical contribution of HSEQ relative to structural/unified RAG,
2. Improve the readability and coherence of the method section,
3. Provide a more careful and fair empirical comparison (including the re-run of TAT-LLM under a comparable backbone),
4. Explicitly analyze efficiency-accuracy trade-offs among LLM-only, ToG, and HSEQ, and
5. Strengthen examples to support the design choices.

We hope these changes address the reviewers' concerns and help the AC assess the paper more favorably.