Chronology-Driven RAG for Construction Meeting Minutes

Ioannis Kostis[0000−0002−1875−8587] and Pierre Schaus[0000−0002−3153−8941]

ICTEAM/INGI, UCLouvain, Belgium
{firstname.lastname}@uclouvain.be

In large-scale construction projects, key decisions evolve continuously and are primarily documented in meeting minutes. Tracking how these decisions change over time is critical for project management but typically requires tedious manual review of extensive archives. To address this challenge, we present a conversational system, developed as a case study in collaboration with a major Belgian construction company, that enables natural-language access to project documentation and returns semantically relevant, time-annotated answers reconstructing the chronology of decisions.

Our approach builds on a Retrieval-Augmented Generation (RAG) [6] framework that integrates semantic and temporal retrieval with Large Language Model (LLM) inference. This framework grounds responses in project-specific records, maintaining factual accuracy while revealing how decisions evolve. By fusing temporal awareness with semantic search, the system allows professionals to query large collections of meeting minutes and receive factual, time-ordered responses, thereby transforming unstructured archives into accessible chronological knowledge.

While prior research on RAG and meeting documentation systems [1, 2, 8] has advanced semantic retrieval and conversational access to project records, these approaches generally overlook the temporal evolution of information. Our work differs by explicitly modeling time as an independent retrieval dimension, enabling the reconstruction of decision histories and the chronological reasoning needed for long-duration projects.

The architecture extends the standard RAG pipeline (*indexing*, *retrieval*, and *generation* [3]) by introducing explicit time segmentation. Meeting minutes are parsed into timestamped textual passages that populate a Vector Database. For every user query, retrieval is performed independently over each chronologically segmented sub-index, ensuring that responses reflect only the knowledge available during the corresponding period of the project. Retrieved passages are then reranked for semantic relevance and passed to an LLM that synthesizes grounded responses ordered by time.

The framework integrates several key components: a hybrid retrieval mechanism combining semantic embeddings [5] with BM25 [7] keyword search to balance Precision and Recall; the aforementioned temporal sub-indexing strategy that mitigates contradictions between outdated and current decisions by segmenting the knowledge base over time; and a response synthesis module employing an LLM-as-a-judge [4] approach to merge redundant outputs and enhance overall readability. A lightweight graphical interface supports interactive

querying, displaying both the decision timeline and the retrieved source pages to facilitate transparent navigation of project history.

We evaluated the system using a real-world dataset compiled in collaboration with our industry partner. The dataset contains 60 anonymized French-language meeting minutes (around 3,900 words each) covering January 2022–June 2024, each recording stakeholder decisions with dated revisions. It is publicly available and accompanied by 8 expert-defined benchmark queries targeting decision-tracing tasks such as material approvals or regulatory comments. This dataset supports objective evaluation of retrieval performance and fosters reproducibility for future research on temporal information access in construction management.

Quantitative evaluation focused on retrieval and reranking quality using metrics such as Precision@5, Recall@5, and F1@5. Experiments show that temporal sub-indexing improves both speed and accuracy, with optimal batch sizes balancing efficiency and context granularity. Incorporating Reranking into the retrieval process raises Precision@5 from roughly 0.67 to 0.74 and F1@5 from 0.44 to 0.50. The system retrieves at least one relevant page for over 80% of queries within the top five results, demonstrating robust chronological alignment. These findings confirm that the proposed architecture effectively captures the temporal evolution of project decisions while remaining computationally practical for real-world use.

In summary, we introduce a chronology-aware RAG framework that enables the reconstruction of decision timelines from construction project documentation. By embedding temporal logic into retrieval and generation, the method produces interpretable, time-stamped answers that enhance transparency and traceability in project management. Looking ahead, we expect that our proposed approach could find applications in other domains, such as legal, medical, or historical archives, where chronology is integral to meaning. Further integration with temporal knowledge graphs and fine-tuned language models could enhance reasoning over evolving records. The open-source release of the dataset and code provides a foundation for advancing temporal retrieval and conversational knowledge systems across industries.

Acknowledgments. This work was conducted under the framework of the European Digital Innovation Hub project EDIH-CONNECT, co-funded by the European Union under the Digital Europe Programme (Grant Agreement No. 101083626). The case study was carried out in collaboration with BPC Group, whose input was instrumental to the research.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Baeuerle, S., Radyschevski, M., Pado, U.: Automeet: a proof-of-concept study of genai to automate meetings in automotive engineering. arXiv preprint arXiv:2507.16054 (2025)
2. Chen, G., Alsharef, A., Ovid, A., Albert, A., Jaselskis, E.: Meet2mitigate: An llm-powered framework for real-time issue identification and mitigation from construction meeting discourse. Advanced Engineering Informatics **64**, 103068 (2025)
3. Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Wang, H., Wang, H.: Retrieval-augmented generation for large language models: A survey. arXiv preprint arXiv:2312.10997 **2**(1) (2023)
4. Gu, J., Jiang, X., Shi, Z., Tan, H., Zhai, X., Xu, C., Li, W., Shen, Y., Ma, S., Liu, H., et al.: A survey on llm-as-a-judge. arXiv preprint arXiv:2411.15594 (2024)
5. Karpukhin, V., Oguz, B., Min, S., Lewis, P.S., Wu, L., Edunov, S., Chen, D., Yih, W.t.: Dense passage retrieval for open-domain question answering. In: EMNLP (1). pp. 6769–6781 (2020)
6. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.t., Rocktäschel, T., et al.: Retrieval-augmented generation for knowledge-intensive nlp tasks. Advances in neural information processing systems **33**, 9459–9474 (2020)
7. Robertson, S., Zaragoza, H., et al.: The probabilistic relevance framework: Bm25 and beyond. Foundations and Trends® in Information Retrieval **3**(4), 333–389 (2009)
8. Wu, C., Ding, W., Jin, Q., Jiang, J., Jiang, R., Xiao, Q., Liao, L., Li, X.: Retrieval augmented generation-driven information retrieval and question answering in construction management. Advanced Engineering Informatics **65**, 103158 (2025)