International Journal on Science and Technology (IJSAT)

E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

Memory Keeper: AI for Organizational Decision Traceability

Aryan Yadav[1], Chanchal Budhadeo[2], Sagar Zope[3], Ganesh Wadmare[4]

[1,2,3]Student, Dept. of AI & Data Science, K. J. Somaiya Institute of Technology, Mumbai, India
[4]Assistant Professor, Dept. of AI & Data Science, K. J. Somaiya Institute of Technology, Mumbai, India

Abstract

Critical business decisions do not happen in isolation in a particular location; rather, the operational environment in which such business decisions take place gets dispersed in messaging conversations, email threads, and meeting discussions. It may not be feasible to follow the rationale for a particular business decision taken months ago. In the end, the process becomes obscure, and there is a marked lack of traceability.

A hybrid solution named Memory Keeper is proposed to overcome the problem of data silos. By combining Knowledge Graphs (KG) with vector stores to build an improved Retrieval-Augmented Generation (RAG) model, the framework continuously feeds distributed communication flows into the model. It processes unstructured information to identify essential decisions and maps structural relationships between stakeholders, timelines, and topics into a queryable KG.

Although conventional semantic search approaches usually ignore the interdependencies of events, Memory Keeper can retrieve historical context with high accuracy by using graph-based reasoning to route vector embeddings. This approach makes it possible to generate traceable and explainable responses to user queries, thus making enterprise knowledge an active tool that supports rapid decision-making with well-documented evidence.

Keywords: Organizational Memory, Retrieval-Augmented Generation (RAG), Knowledge Graphs, Semantic Search, Enterprise Knowledge Management, Decision Traceability.

1. Introduction

A substantial amount of unstructured communication also comes from enterprises through different channels of their operational activities [1, 2]. Business decisions, reasons for the decisions, and task lists may also be present in this huge amount of information passing through the system. It is also observed that there may not be proper understanding of the situation, repeated work, and poor collaboration among people in the absence of a single tool that can effectively capture and hold all this information and knowledge.

The underlying implicit reasoning and complex cross-departmental dependencies that govern major organizational decisions rarely ever fall under the scope of traditional documentation approaches (meeting minutes, ticketing systems, or internal wikis, for example) [3, 13]. Severe operational continuity gaps are frequently created due to the loss of tacit knowledge when team members transition out of ever-evolving teams [14]. In this way, the institutional memory gradually erodes, making it extremely difficult for the present team to look back and attempt to establish the reasons why particular technical or strategic decisions were made in the first place.

New avenues for the automated capture, summarization, and contextualization of corporate dialogues have been paved by the recent developments in the field of Artificial Intelligence and its sub-stream of Natural Language Processing. The literature suggests that the underlying semantics of complex corporate dialogues could be parsed and the relevant information with regard to critical corporate decisions could be isolated [4]. The accurate retrieval of corporate decisions and justifications thereof along with relevant variables such as dates, project names, and people involved has been made possible [6, 7]. It is emphasized in the literature that by integrating such algorithms into Knowledge Graphs (KGs) and Retrieval-Augmented Generation (RAG) [5, 17], architectural frameworks have been enabled that transcend the traditional boundaries of data storage into the realm of active reasoning by machines.

In order to realize these concepts, it has been proposed that an intelligent and persistent organizational memory infrastructure be used to fully automate the life cycle of decision archiving and retrieval (this has been termed Memory Keeper). In the first instance, multimodal inputs from different modes of communication will be fed into the architecture using automated workflow pipelines (such as n8n and Gmail APIs) [8, 9]. Subsequently, entity and decision extraction will be carried out by applying LLMs, and the resultant mathematical embeddings will be stored in a vector database to allow for the execution of semantic search, while complex data relationships will be mapped using a temporal knowledge graph [15, 16]. By querying the conversational interface [11, 20], particular questions regarding the decision will be posed by staff members who want to know "why" or "when", and in turn, rich explanations will be provided by the system itself.

Throughout this manuscript, the fundamental architecture, workflow, and evaluation strategy for the Memory Keeper model are systematically outlined. This allows one to understand how data storage units in a dormant state can be utilized by artificial intelligence systems to generate flexible and self-updating memory networks [18, 19] (which in turn may be used to increase transparency within organizations, day-to-day accountability, and strategic alignment in long-term contexts).

2. Literature Survey

The history of development in organizational memory over the years has been well recorded and shows that there has been a move from abstract concepts to the development of the first information systems that included digital technology [12][13]. Fundamentally, early memory systems included human data entry and database systems [14], which often provided limited information.

For more than three decades, there has been intensive research into how institutions store and transfer knowledge. It is widely recognized that the first information systems were designed primarily to store information and make it accessible, but not to record the underlying context of why certain decisions were taken, what was being aimed at, and the subtleties of the reasoning behind decisions. In the early models of information systems, memory was considered to be embedded within the people, the culture of the workplace, and the physical objects surrounding them [1]. This was to underscore the importance of retaining the underlying reasoning behind decisions as a way of achieving organizational learning. As a matter of fact, conventional documentation such as meetings and emails was found to be inadequate as a way of retaining the complex interrelatedness of decisions.

As a direct response to such shortcomings, Organizational Memory Information Systems (OMIS) was investigated, where structured documentation and unstructured knowledge within the workplace were successfully integrated [2]. In these first frameworks, the strict archiving of final decisions, situational context, and associated timelines was highly emphasized. It is recognized, however, that such systems were largely static in nature due to the inherent requirement for manual updates and highly structured data input formats. This naturally limited scalability and real-world applicability due to an inherent lack of intelligence between sources.

In order to understand the architectural leap that has been taken with the advent of modern hybrid approaches, it is essential to first understand the fundamental limitations that were recognized with the early research on Organizational Memory. For example, the theoretical frameworks that were presented by early researchers on Organizational Memory emphasized that while it is possible to easily store individual data points, the "situational context," or the overall environmental factors that affect a decision, is almost always not stored during the data input process. Conventional OMIS architectures were highly dependent on highly structured, manual inputs, which required the employee to almost act as a database administrator, manually tagging their decisions. This, of course, has resulted in a level of friction within the system. As one can imagine, when placed in a stressful situation, it is inevitable that documentation procedures will not be followed, immediately resulting in a loss of Organizational Memory.

It has only been with the recent advancement of Natural Language Processing technology that a solution has been found to this problem. Recent studies have shown that Large Language Models can be used as an intelligent, background parser, successfully retrieving not only the explicit decisions made within communication logs but also the implicit decisions.

However, with this integration comes a new set of complexities, specifically in relation to trust. Where a machine learning system is being utilized in a manner that retrieves and summarizes important corporate history, stakeholders need a way to trust in its accuracy. This has led to one of the most rapid advances in Explainable AI (XAI) in relation to decision support systems [7, 11]. By being forced to show their work in relation to feature attributions, researchers have been able to bridge this gap between black-box neural networks and corporate governance. Furthermore, the matter is no longer confined to the domain of trust but is rapidly turning into a necessity for corporate governance in a wide range of industries. As the influence of insights from machine learning on the trajectory of businesses continues, the audit of accurate data provenance is a necessity, not a choice [20]. This is where attention weight visualization, as well as natural language justification matrices, can guarantee the maintenance of human oversight with the help of explainable artificial intelligence.

Significantly, the process of knowledge representation, retrieval, and subsequent reutilization has witnessed a paradigm shift due to the advancements in the field of artificial intelligence, especially in the area of natural language processing. Indeed, the ability of modern Large Language Models (LLMs) to process, summarize, and even reason about unstructured text has enabled the automatic extraction of decision-relevant insights from communication data in organizations; in other words, the extraction of relevant decision insights is enabled by the capabilities of LLMs, as discussed in [4, 10]. Furthermore, relevant decision history can be made accessible through query-based conversations by utilizing Retrieval-Augmented Generation (RAG) models that integrate LLMs. Indeed, the synergistic use of retrieval models and generative models, such as the one offered by LangChain [5], has ensured that contextual awareness in organizational memory is maintained.

At the same time, the data interconnectivity and semantic reasoning have been significantly improved with the emergence of Knowledge Graphs (KGs). In order to meet the requirements of dynamic knowledge flows, the concept of Self-Organizing Personal Knowledge Assistants (SOPKAs) has been proposed in the literature [6]. These knowledge assistants are based on a graph structure that is in a continuous state of change due to user interactions and information patterns. Complex relationships in the graphs can be identified in order to develop a robust interconnected organizational memory.
The importance of informal and implicit communication in the maintenance of knowledge continuity may be highlighted through the incorporation of Soft Organizational Memory in small and medium-sized enterprises [3]. It has been argued that informal communication processes play a critical role in the maintenance of organizational knowledge. In addition, the organizational memory management life cycle may be supported through the incorporation of automated technologies, for example, n8n. This may allow the ingestion, transformation, and synchronization of data across heterogeneous systems; thus, new knowledge may be automatically ingested into the central repository [8]. In this context, human intervention may be minimized significantly, leading to the self-updating and dynamic evolution of knowledge.

New ideas in Explainable AI (XAI) have been incorporated into the decision system of AI. As the role of summaries and recommendations generated by AI increases in today's organizations, it is necessary to understand the reasons behind a particular decision or conclusion made by the AI system [7, 11, 20]. This will create trust and accountability in the digital memory system.

3. Comparative Analysis

Nowadays, one can clearly see this shift towards less rigid, more AI-driven databases. If one looks back at older literature on Organizational Memory Information Systems (OMIS) from the 1980s and 1990s, one of the big issues was that these databases were quite structured, taking a lot of data inputs. Therefore, some of the nuances of the decision-making process were lost [1, 2]. There was also literature emphasizing the importance of 'informal chatter' in small and medium-sized enterprises [3], indicating that conventional databases were not adequate.

In order to overcome these structural issues, subsequent researchers suggested the integration of Large Language Models (LLMs) and traditional approaches based on the use of the Retrieval-Augmented Generation (RAG) pipeline [4, 5]. The revolution in the development of generative AI, as observed with the emergence of ChatGPT-based conversational interfaces, has demonstrated the capability of efficiently summarizing unstructured corporate dialogues in a fundamental manner [10]. It has been demonstrated that traditional vector-based retrieval approaches are extremely effective for the retrieval of isolated documents; however, significant challenges are encountered for multi-hop reasoning. A comprehensive synthesis of the developing methodologies and the technical contributions are provided in Table 1.

Table 1: Comparative Analysis of Organizational Memory Frameworks

System / Framework	**Model / Technique Used**	**Key Findings & Metrics Reported**	**Primary Use Case & Operating Conditions**
Organisational Memory Framework [1]	Cognitive & Structural Theory	Organizational memory defined as an entity embedded within individuals and culture.	Foundational theoretical baseline for institutional memory retention.
OMIS Model [2]	Information System Integration	Structured data storage enhanced, but semantic adaptability severely lacked.	Early physical attempt to digitize and systematize organizational memory.
Soft Organizational Memory [3]	Tacit Knowledge Capture	Informal and social knowledge highlighted as crucial for long-term memory retention.	Applied soft memory theory to modern digital enterprise ecosystems.
AI-Based Knowledge Retention [4]	LLMs & NLP Summarization	Automated extraction of decision insights from unstructured text enabled.	Demonstrated the baseline potential of AI for understanding organizational dialogues.
LangChain + RAG Stack [5]	Retrieval-Augmented Generation	Precision and contextual query capabilities significantly improved.	Combined standard retrieval mechanisms with generative AI for intelligent access.
Self-Organizing PKA [6]	Knowledge Graph & Cypher Queries	Self-evolving knowledge representation and contextual linking enabled.	Established dynamic topological linkages between people, processes, and policies.
XAI-Enhanced Decision Systems [7]	Explainable AI (SHAP, LIME)	Interpretability of underlying neural models successfully improved.	Enhanced overall transparency and traceability in AI-assisted

System / Framework	Model / Technique Used	Key Findings & Metrics Reported	Primary Use Case & Operating Conditions
			corporate decisions.
n8n Workflow Automation [8]	No-Code Data Integration	Multi-source data ingestion and synchronization fully automated.	Enabled seamless pipeline automation for real-time knowledge capture.
Gmail API Integration [9]	RESTful API Retrieval	Communication data efficiently extracted for downstream analysis.	Provided a reliable and secure method for harvesting historical organizational context.

Nevertheless, despite these architectural breakthroughs, the actual deployment of autonomous memory systems in a live enterprise setting is plagued by significant technical hurdles, in particular with regard to data heterogeneity. In fact, modern-day enterprise communication is extremely fragmented across a broad range of wildly diverse communication channels, from the asynchronous, text-dominant nature of email communication to the synchronous, audio-dominant nature of virtual meetings. While modern-day automated data pipelines, as discussed in [8, 9], are clearly effective in consuming such heterogeneous data streams, synchronizing the timestamp and performing entity resolution on the data remains a computationally costly operation. In fact, a project may be referred to by as many as three acronyms across Slack, Jira, and Zoom! As such, extremely sophisticated, context-aware entity resolution algorithms will be necessary to ensure that the knowledge graph does not fragment into a state that is effectively useless.

Furthermore, the temporal nature of corporate knowledge introduces a complex problem of "sunsetting" outdated knowledge. Indeed, as the organization itself changes and outdated projects are no longer relevant to modern operations, the autonomous knowledge system will need to recognize which relational entities are no longer relevant. However, while the knowledge system could potentially remove such entities, it will destroy the context of the knowledge. On the other hand, retaining outdated knowledge equally with modern knowledge will degrade query performance while introducing a massive amount of noise. Thus, developing dynamic algorithms that can smoothly sunset outdated knowledge without severing the topological connections between knowledge entities is a major research problem for the field of databases.

The self-updating nature of such a hybrid system introduces a profound security risk. Indeed, while the automated system is ingesting information from private communication channels to update the corporate knowledge graph, it is highly likely that it will ingest and globally expose sensitive personnel information or otherwise forbidden knowledge. Thus, developing robust Role-Based Access Control (RBAC) technologies that work directly at the vector retrieval and graph traversal layer to ensure that the AI only synthesizes knowledge using information that the querying employee is authorized to view is absolutely essential to the practical deployment of such a system.

If we compare the application of conventional RAG systems with systems that incorporate Knowledge Graphs, also denoted as Graph-RAG [6, 17], we can see that there are considerable improvements with regard to the accuracy of the provided facts and logical traceability. It has been systematically demonstrated that, by providing this topological mapping of the provided entities, stakeholders, and decisions, the Knowledge Graph actually provides a verifiable truth that can be used to heavily restrict the hallucination rates that are normally inherent with conventional generative systems. Finally, the application of these sophisticated systems in real-world enterprise environments has been made possible by the application of automated workflow systems [8, 9], which can be used to continuously ingest unstructured corporate data.

4. Open Challenges

Despite the recent advances in knowledge extraction, there are still a number of major technical challenges to overcome in creating self-managing organizational memory systems. The first and most significant challenge is multimodal corporate data. The technology has already overcome the challenge of automatic text extraction for common emails and chat systems. Nevertheless, reliable parsing and smooth integration of the noisy voice-to-text messages produced during virtual meetings [4, 8] require a great deal of improvement.

One of the challenges in hybrid RAG systems is the lack of robustness in deterministic query generation. Although semantic vector searches are very tolerant of ambiguous user queries, graph databases demand precise, perfectly formatted query languages. Direct translation of natural language queries into executable graph queries (Text-to-Cypher) frequently leads to syntax errors or empty sets. As was noted in recent studies of Large Language Models' performance on knowledge graphs [21], when the routing logic fails to identify the proper nodes because of slight syntax errors, the system is left with no choice but to rely entirely on the vector database, temporarily losing its hierarchical knowledge. Development of more robust and fault-tolerant translation interfaces between user queries and graph databases is a critical area for future improvement.

Another major problem is that there is a lot of computation required to update real-time temporal knowledge graphs [6]. Company knowledge is ever-changing as the company's teams are changing and projects are being completed. We need a system that can remove relational nodes gracefully without losing their entire history.

Aside from the mechanical difficulty in updating the graph, however, is the profound algorithmic difficulty in semantic conflict resolution. In any large corporation, unstructured communication is necessarily contradictory. Two different department heads may offer conflicting justifications for the same business decision from different email threads or meeting transcripts. When this contradictory information is ingested into the automated workflow [8, 9], the resulting Knowledge Graph may have a fractured and paradoxical node structure. Presently, most Large Language Models lack the deterministic logic to autonomously weigh the hierarchical status of different employees in the corporation or determine which statement represents the final, most recently updated corporate reality. Designing such a probabilistic weighting system that may assign "confidence scores" to these contradictory nodes in the graph, without any human intervention, is an important area for future database engineering.

Moreover, the latency induced by this hybrid routing remains an important bottleneck for any application that requires real-time performance. Although standard vector databases may perform similarity searches in milliseconds, the "Text-to-Cypher" query translation for Graph-RAG is necessarily computationally expensive. As discussed in recent evaluations of the performance of enterprise knowledge graphs [6, 17], when an LLM must first parse a user's natural language prompt, generate a syntactically perfect graph query, execute the query, and then synthesize the results into an intelligible answer, the overall response time may degrade significantly. To design organizational memory systems that function as seamless and real-time conversational interfaces, future researchers must significantly optimize the query execution layers of these models, potentially through the utilization of smaller and more fine-tuned routing models that avoid the latency induced by massive and generalized LLMs.

The issue of "cascading hallucinations" in hybrid RAG models represents a very severe threat to data integrity. As discussed in prior sections, while it is well understood that Knowledge Graphs severely inhibit standard generative hallucinations by grounding the AI in verifiable facts [5, 15], these models are highly susceptible to ingestion errors. If a parsing model mis-identifies an entity or a relationship during the data ingestion phase of this pipeline, this mistake is forever encoded into the topology of the final generated Knowledge Graph. This node will then be subsequently looked up as a “verifiable fact” and a historically incorrect explanation will be generated for it with confidence. The requirement for auditing such explanations to verify generated graph triples against original raw text embeddings is a necessity to avoid corrupting institutional memory.

Finally, as companies begin to integrate these autonomous systems, the absence of Explainable AI (XAI) systems [7, 11] is a major showstopper. Without the ability to explain how an AI system arrived at its conclusions, stakeholders simply won’t trust the AI system in a real-world corporate setting.

5. Future Directions

In order to move past the current shortcomings in automated organizational memory, future development needs to extend the paradigm of data ingestion well beyond the current boundaries of text-based communication. Scientists need to develop more complex multimodal processing pipelines that have the ability to semantically interpret complex audio and video sources, such as the voice-to-text transcripts created during virtual meetings [4, 8]. By seamlessly integrating these complex, unstructured interactions with the current chat and email logs, future systems will be able to create a comprehensive knowledge map of the enterprise.

Regarding database management, one of the most important areas of future research is the integration of lightweight dynamic algorithms into structured knowledge bases. As corporate structures, the scope of projects, and team dynamics are constantly in flux, it is essential that systems be able to dynamically restructure relational links in order to properly depreciate aging institutional knowledge [6]. Through the development of sound methods for dealing with this lifecycle issue, future versions can properly reflect the natural process of organizational memory without requiring full system recalculation.

Furthermore, as hybrid retrieval systems continue to mature as a technology and find their way into the enterprise, the development of strong infrastructures in Explainable AI (XAI) must be a key focus in the engineering of these systems moving forward. Future development efforts must focus on the creation of user interfaces that are friendly and capable of naturally expressing the complex mathematical routing logic behind the neural networks in simple, easy-to-understand linguistic terms [7, 11].

6. Conclusion

The basic transition of organizational memory from static repositories to dynamic, artificially intelligent systems has been analyzed in this review. It is argued that the traditional approaches to documentation are no longer adequate to capture the complex, tacit reasoning that goes into enterprise-level decision-making. With the integration of Large Language Models and Retrieval-Augmented Generation, the task of automated extraction and contextualization of corporate knowledge is successfully accomplished. Moreover, the important challenge of artificial intelligence hallucinations is effectively addressed by anchoring these generative models to structured Knowledge Graphs. As the automated data pipelines become more and more common in enterprises, a continuous and self-updating institutional memory is created. The future academic work needs to focus on improving multimodal data ingestion and improving the explainability of these complex neural networks. In the end, long-term strategic alignment and operational consistency are ensured when these sophisticated semantic frameworks are fully operationalized in the contemporary workplace.

References

1. J. P. Walsh, G. R. Ungson, "Organizational Memory", Academy of Management Review, 1991, 16 (1), 57-91.
2. A. Barros, C. Ferreira, M. Rodrigues, "Information Systems and Organizational Memory: A Conceptual Model", Procedia Computer Science, 2015, 64, 1174-1183.
3. A. Al-Edenat, "Soft Organizational Memory in Small and Medium Enterprises", Journal of Knowledge Management, 2022, 26 (3), 564-580.
4. OpenAI Research Group, "Advances in Large Language Models for Knowledge Extraction", AI Communications Journal, 2022.
5. LangChain Team, "Retrieval-Augmented Generation Framework for Contextual AI Systems", ArXiv Preprint arXiv:2307.09724, 2023.
6. F. Jilek, M. Schmitz, T. Ertl, "Towards Self-Organizing Personal Knowledge Assistants in Evolving Corporate Memories", Frontiers in Artificial Intelligence, 2023, 6.
7. S. Miller, L. Zhang, "Explainable Artificial Intelligence for Decision Transparency in Enterprise Systems", IEEE Access, 2023, 11, 14587-14603.
8. n8n Documentation Team, "Workflow Automation for Data Integration and Process Management", n8n Docs, 2023.
9. Google Developers, "Gmail API Overview – Automating Email Retrieval and Summarization", Google Cloud Documentation, 2024.
10. B. Mesko, "The ChatGPT (Generative Artificial Intelligence) Revolution", Journal of Medical Internet Research (PMC), 2023.
11. S. Ali, "Explainable Artificial Intelligence (XAI): What We Know and What We Don't", ScienceDirect, 2023.
12. G. Eldenberg, J. Pautzke, "How Can Organizational Memory Theories Contribute to Organizational Memory Systems?", ResearchGate / Information Systems, 2012.

13. M. Foroughi, "Organizational Memory Studies", Organization Studies, 2020, 41 (12), 1725-1748.
14. E. W. Stein, V. Zwass, "Actualizing Organizational Memory with Information Systems", Information Systems Research, 1995, 6 (2), 85-117.
15. J. Ghiran et al., "Knowledge Graphs as a Source of Trust for LLM-Powered Question Answering Systems", ScienceDirect / Information Discovery and Delivery, 2024.
16. M. Galkin, S. Auer, M. E. Vidal, S. Scerri, "Enterprise Knowledge Graphs: A Semantic Approach for Knowledge Management", ResearchGate, 2017.
17. L. Mariotti, V. Guidetti, F. Mandreoli, "Combining Large Language Models with Enterprise Knowledge Graphs", Frontiers in Artificial Intelligence, 2024.
18. S. Hameed, M. Soori, F. K. G. Jough, "AI-Based Decision Support Systems in Industry 4.0: A Review", Journal of Economy and Technology, 2024.
19. M. Caldarulo et al., "AI-Driven Decision Support Systems in Management: Enhancing Strategic Planning and Execution", ResearchGate, 2023.
20. G. Kostopoulos, G. Davrazos, S. Kotsiantis, "Explainable Artificial Intelligence-Based Decision Support Systems: A Recent Review", MDPI Electronics, 2024, 13 (14), 2842.
21. S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, X. Wu, "Unifying Large Language Models and Knowledge Graphs: A Roadmap", IEEE Transactions on Knowledge and Data Engineering, 2024.