From Knowledge Retrieval to Execution: Designing Executable Knowledge Systems

Sara Michelazzo, Parmeet Kaur and Saurabh Saxena
Thoughtworks, Chicago, USA

smichel@thoughtworks.com
parmeet.kaur@thoughtworks.com
saurabh.saxena@thoughtworks.com

Abstract: Generative AI is changing the role of knowledge in organisations. Traditional knowledge management (KM) systems have primarily supported storage, access and retrieval, assuming that knowledge is interpreted and applied by human users. In AI-enabled environments, however, organisational knowledge increasingly becomes a direct input into execution, shaping generated proposals, analyses, summaries, recommendations and other workflow outputs. This shift exposes a limitation of retrieval-oriented KM: fragmented, outdated or weakly governed knowledge can be amplified through AI-generated outputs, reducing consistency, reliability and trust. This paper introduces executable knowledge systems as a conceptual model for structuring organisational knowledge to support reliable human and AI-assisted execution. The term executable is used in a socio-technical sense. Knowledge does not necessarily become code, but is curated, validated and embedded into workflows so that it can guide outputs, decisions and actions. The paper distinguishes this concept from prior work on executable knowledge graphs and executable knowledge bases, which primarily focus on deterministic execution through rules, scripts or formalised representations. The paper further develops a framework of decay and compounding loops to explain how AI-enabled knowledge systems evolve over time. In decay loops, AI-generated outputs re-enter the knowledge environment without sufficient validation, allowing inconsistency and low-quality knowledge to accumulate. In compounding loops, curated knowledge assets are refined through governed feedback, domain ownership and controlled reuse, enabling improvements in reliability over time. The framework is informed by an exploratory case study within a global professional services organisation, where a curated knowledge environment was introduced to support AI-assisted workflows in the Retail, Consumer Products, Travel and Transportation domain. The evaluation compared outputs generated from a controlled, subject matter expert (SME)-validated knowledge dataset with outputs generated from an unconstrained organisational knowledge base. Findings indicate improved retrieval relevance and output quality when AI systems operate on validated knowledge assets. he paper contributes to KM research by reframing KM as a system design challenge for AI-enabled execution and by positioning governance, validation and feedback control as central mechanisms for reliable organisational knowledge use.

Keywords: Knowledge management, Generative AI, AI Reliability, Knowledge governance, Executable knowledge systems, AI-assisted workflows

1. Introduction

Organisations have invested in knowledge management (KM) for decades with the expectation that captured knowledge can be reused to improve efficiency, consistency and organisational learning. KM is commonly understood as encompassing knowledge creation, storage, retrieval, transfer and application (Alavi and Leidner, 2001). In practice, however, many KM initiatives have emphasised storage and retrieval through repositories, intranets, taxonomies and search tools. This has created a persistent gap between knowledge availability and knowledge application: organisations may possess large volumes of knowledge, yet employees still recreate work, rely on informal networks or struggle to apply existing knowledge in execution contexts.

The adoption of generative AI makes this gap more consequential. Recent KM research argues that generative AI reshapes core knowledge processes and requires new frameworks to manage productivity gains alongside risks to knowledge quality, expertise and organisational control (Storey, 2025). Studies of AI-enabled knowledge work similarly show that effective use of generative AI requires adaptable user control, transparent collaboration mechanisms and the ability to integrate organisational background knowledge with external information (Yun et al, 2025).

AI-enabled systems increasingly generate proposals, analyses, summaries, recommendations and other workflow outputs. Knowledge therefore becomes a direct operational input into execution workflows. When the underlying knowledge is fragmented, outdated or poorly governed, AI systems may amplify inconsistency at scale. This risk is consistent with research on large language model hallucination, which shows that generated outputs can be fluent and plausible while remaining unsupported, contradictory or unverifiable (Huang et al, 2025; Li et al, 2023).

This paper addresses this challenge by introducing executable knowledge systems: governed knowledge environments in which curated knowledge assets are designed to support reliable execution by both humans and AI systems. The term "executable" is used in an organisational and socio-technical sense. It does not imply that all knowledge becomes code; rather, it means that knowledge is structured as an operational input that shapes workflow outputs, decisions and actions.

The paper further clarifies how this concept relates to prior work on executable knowledge graphs and executable knowledge bases. Existing approaches, such as ExeKG, focus on representing knowledge so it can be translated into executable scripts or reusable workflow components for analytics and machine learning (Zheng et al, 2022a; Zheng et al, 2022b). Related work on executable clinical knowledge bases emphasises the formalisation of guideline rules for decision support (Quimbaya et al, 2014). By contrast, the model proposed here addresses generative AI environments in which reliability depends not only on formal representation, but also on governance, validation, ownership and feedback control.

The paper makes three contributions. First, it reframes KM as a system design challenge for AI-enabled execution rather than only a problem of knowledge access. Second, it defines executable knowledge systems and distinguishes them from deterministic executable-knowledge approaches. Third, it introduces decay and compounding loops as a framework for understanding how AI-enabled knowledge systems degrade or improve over time depending on governance.

2. Background and Literature

KM has traditionally been grounded in the premise that organisational knowledge can be captured, stored, shared and applied to improve performance. Foundational KM literature emphasises knowledge creation, transfer and reuse, while also recognising that knowledge creates value only when it is applied in organisational activity (Nonaka and Takeuchi, 1995; Davenport and Prusak, 1998; Alavi and Leidner, 2001). However, many KM systems have operationalised this agenda primarily through repositories, taxonomies, search and document access, contributing to a gap between knowledge availability and usability.

Information quality research helps explain this gap. Lee et al (2002) argue that information quality is multidimensional and must be evaluated in relation to use, including intrinsic, contextual, representational and accessibility dimensions. For AI-enabled work, this distinction is important because accessible knowledge may still be unsuitable for execution if it is outdated, duplicated, poorly contextualised or weakly governed.

Recent research on generative AI and knowledge work reinforces this limitation. Storey (2025) argues that generative AI changes the role of KM by altering how organisational and worker knowledge is accessed, applied and evaluated. Kirchner et al (2025) show that generative AI can support knowledge creation and sharing in software development, but expert review remains necessary to assess reliability, adapt outputs and maintain quality. Taken together, these findings suggest that AI-enabled KM requires mechanisms for validation, governance and contextual use.

The reliability challenge is also visible in research on large language model hallucination. Hallucination refers to outputs that are fluent but unsupported, contradictory or unverifiable. Huang et al (2025) show that hallucination remains a central challenge for LLM reliability, including in retrieval-augmented systems. Li et al (2023) similarly show that LLMs may fabricate unverifiable information and that external knowledge and structured reasoning steps can help reduce hallucination. In organisational contexts, hallucination should therefore be understood not only as a model-level problem, but also as a knowledge-system problem.

Research on retrieval-augmented generation (RAG) and knowledge graphs provides one response to this problem. RAG approaches seek to ground generated outputs in external knowledge sources, while knowledge graph-based approaches add structure and relationships to improve retrieval and reasoning. Zhu et al (2025) propose knowledge graph-guided RAG to preserve fact-level relationships between retrieved chunks and improve response coherence. Liu et al (2025) show in a clinical triage setting that knowledge graph-based RAG outperforms prompt-only and naïve RAG approaches, suggesting that structured domain knowledge can improve reliability in knowledge-intensive workflows.

This paper also builds on, but differs from, existing work on executable knowledge. Zheng et al (2022a) introduce ExeKG as an executable knowledge graph system for user-friendly data analytics, where semantic representations of machine-learning knowledge can be translated into reusable scripts. Zheng et al (2022b) extend this idea to an industrial machine-learning context through a Bosch welding-monitoring case. In healthcare, Quimbaya et al (2014) present an executable knowledge base for clinical practice guideline rules, illustrating how domain knowledge can be formalised for rule-based decision support.

These approaches are important foundations, but they address a different execution problem from the one examined in this paper. In executable knowledge graphs and executable clinical knowledge bases, execution is primarily deterministic: structured knowledge is translated into scripts, rules, queries or compliance mechanisms. In AI-enabled KM, execution is probabilistic and generative: systems interpret and recombine organisational knowledge to produce proposals, analyses, summaries and recommendations.

Taken together, these studies clarify both the foundation and the distinction of the present paper. Existing executable knowledge research shows that formalised knowledge can be operationalised through rules, scripts, queries and machine-learning workflow components (Zheng et al, 2022a; Zheng et al, 2022b; Quimbaya et al, 2014). The concept of executable knowledge systems proposed here extends this work to AI-enabled organisational settings, where decay and compounding loops describe how such systems evolve: without validation, generated outputs can degrade the knowledge environment; with governed feedback, they can improve curated knowledge assets.

3. Conceptual Framework

3.1 From Retrieval to Execution

The emergence of generative AI changes the role of knowledge in organisational work. In traditional KM environments, knowledge is retrieved, interpreted by humans and applied through manual execution. In AI-enabled environments, AI systems increasingly retrieve, interpret and recombine knowledge to generate outputs directly within workflows. Knowledge therefore becomes not only a resource for human interpretation, but an operational input into execution.

This paper defines executable knowledge systems as governed knowledge environments in which curated knowledge assets are structured to support reliable execution by humans and AI systems. The term executable is used in a socio-technical sense. It does not mean that all knowledge becomes code. Rather, it means that knowledge is organised so that it can shape workflow outputs, decisions and actions.

In executable knowledge systems, the unit of value shifts from static documents and repositories toward reusable knowledge assets embedded within workflows. These may include structured best practices, validated approaches to recurring problems, standardised decision frameworks and reusable workflow patterns. Such assets act as governed inputs that shape how AI systems generate outputs and support decision-making across contexts.

As knowledge becomes embedded in AI-enabled workflows, organisational control also shifts upstream. Rather than relying primarily on post-hoc review of generated outputs, organisations must govern the knowledge inputs that shape execution. This shift is summarised in Table 1.

Table 1: From retrieval-oriented KM to executable knowledge systems

Dimension	Traditional Knowledge Management	Executable Knowledge Systems
Primary goal	Improve access to knowledge	Enable reliable execution of work
Core paradigm	Retrieve, read and interpret	Generate, apply and refine
Unit of value	Documents and repositories	Governed knowledge assets and workflows
Quality control	Post-hoc review of outputs	Governance of knowledge inputs
Feedback model	Outputs re-enter knowledge base with limited validation	Only validated outputs update governed knowledge assets
Main failure mode	Fragmentation and low reuse	Inconsistency at scale
Success pattern	Improved access and discoverability	Improving reliability through governed reuse
Role of KM	Curate and organize content	Design and govern execution workflows
Measurement	Usage, access, reuse	Cycle time, consistency, output quality

3.2 Decay and Compounding Loops

The shift from retrieval to execution turns KM systems into feedback systems. In AI-enabled environments, knowledge is repeatedly used to generate outputs that may influence future knowledge assets and workflows. This creates two possible trajectories: decay loops and compounding loops.

A decay loop occurs when AI-generated outputs are reintroduced into the knowledge environment without sufficient validation or governance. In this situation, generated outputs become inputs for future execution regardless of quality. Over time, inconsistencies, duplications and inaccuracies become harder to distinguish from validated knowledge. The system may still produce outputs quickly, but reliability and trust decline.

A compounding loop occurs when knowledge is curated and governed before use in execution, and when generated outputs are reviewed before updating knowledge assets. In this model, domain experts remain responsible for validating and maintaining knowledge assets. Their role shifts from ad hoc contribution or repeated output creation toward continuous stewardship, refinement and improvement of reusable knowledge.

Figure 1 illustrates these contrasting dynamics. The decay loop is driven by uncontrolled feedback, while the compounding loop is driven by governed feedback, validation and ownership.

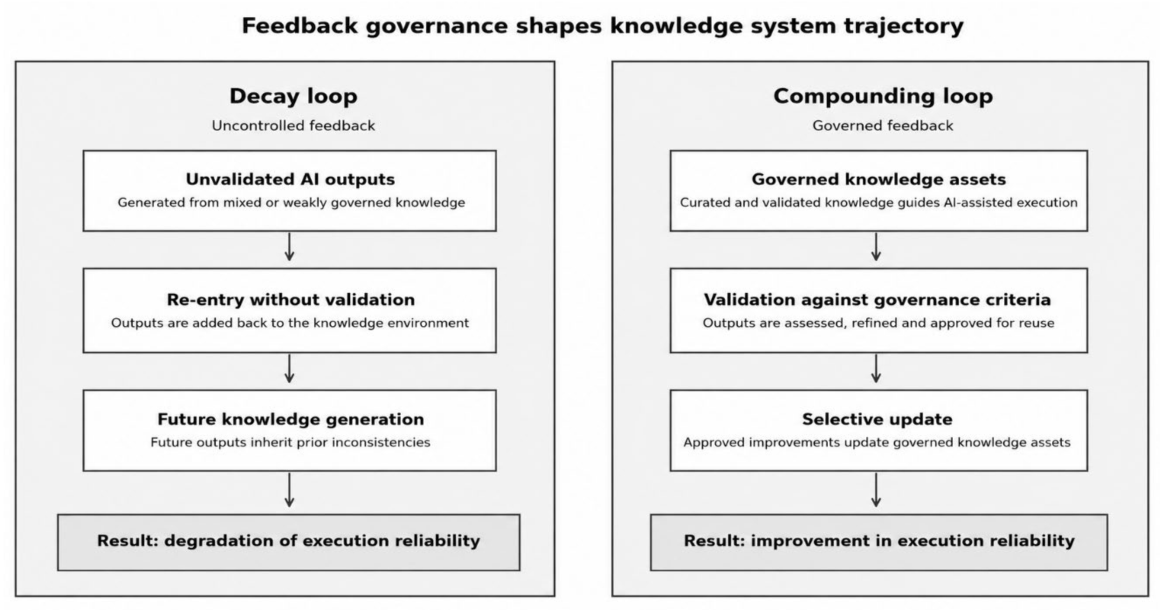

Figure 1: Decay and compounding loops in executable knowledge systems

In decay loops, signal quality declines as validated and unvalidated knowledge become mixed; provenance weakens because sources may remain traceable but lack clear authority; and trust declines as users encounter inconsistent outputs. In compounding loops, consistency improves because execution is guided by standardised inputs; reliability improves because outputs remain grounded in curated knowledge; and reuse accelerates because expert refinements become available for future workflows.

Table 2 summarises the distinction between the two trajectories.

Table 2: Decay and compounding loops

Feedback mechanism	Outputs re-enter the knowledge base without validation	Only validated outputs update governed knowledge assets
Knowledge evolution	Accumulation of inconsistency and duplication	Selective refinement and continuous improvement
Human contribution model	Ad hoc creation and interpretation	Validation, refinement and stewardship

User behavior	Increased reliance on individual judgment	Increased reliance on system-supported
System trajectory	Progressive degradation of knowledge	Progressive improvement in execution

The relationship between executable knowledge systems and decay and compounding loops is therefore direct. Executable knowledge systems are the socio-technical architecture through which knowledge is structured for AI-assisted execution. Decay and compounding loops describe how that architecture evolves under different feedback-governance conditions. Without validation, scale amplifies inconsistency. With governed feedback, scale amplifies validated knowledge, expert refinement and execution reliability. This interpretation is consistent with practitioner-oriented KM research emphasising that future-ready KM requires formal ownership, governance roles and knowledge embedded in the flow of work (APQC, 2025).

4. Case Study and Evaluation Design

4.1 Organizational Context and Intervention

The empirical component of this paper is an exploratory case study conducted within a global professional services organisation employing approximately 10,000 knowledge workers across multiple regions. The organisation delivers complex client engagements that require the reuse of institutional knowledge, including proposals, case studies, delivery patterns and industry insights.

The intervention focused on the Retail, Consumer Products, Travel and Transportation (RCTT) domain. This domain was selected because its work depends heavily on reusable industry knowledge, including retail e-commerce modernisation, marketplaces and omnichannel operations. The RCTT industry team partnered with the AI-First Knowledge Management team to establish a curated source of truth for industry knowledge assets designed for both human and AI-assisted workflows. Confluence was used as the primary platform for structuring and publishing validated knowledge assets.

Prior to the intervention, relevant knowledge was distributed across repositories, collaboration tools and personal storage systems. Although knowledge was abundant, it lacked consistency, validation and clear ownership. Employees were required to identify, interpret and adapt content independently, resulting in duplicated work, variable output quality and reliance on informal networks. These conditions reflected a retrieval-oriented KM model, where access was prioritised but usability within execution workflows was not systematically ensured.

The intervention introduced a controlled knowledge environment designed to improve the consistency, reuse and reliability of knowledge used in proposal development, delivery support and client-facing work. The evaluation focused specifically on RCTT workflows rather than the organisation as a whole.

4.2 Evaluation Design

The evaluation assessed whether constraining AI-assisted workflows to curated and validated knowledge assets improved retrieval relevance and output quality compared with an unconstrained organisational knowledge base. Two knowledge environments were compared:

- Controlled dataset: approximately 450 structured RCTT knowledge pages validated and approved by subject matter experts.
- Unconstrained knowledge base: the broader organisational knowledge environment containing more than one million documents.

Nine representative prompts were defined to simulate common knowledge-intensive tasks, including proposal development, industry analysis and client preparation. Each prompt was paired with an SME-approved reference response representing the expected high-quality answer. The reference responses were defined by an SME with more than 15 years of professional services experience across Retail, Consumer Products, Travel and Transportation.

For retrieval relevance, responses were evaluated using a pass/fail assessment. A response passed when it aligned with the SME reference in content, structure and relevance. Response accuracy was calculated as the proportion of responses aligned with SME-approved references. Resolution rate measured whether a response was generated for each prompt.

For output quality, three artefact types were evaluated: case studies, statements of work (SOWs) and client pitches. Outputs were assessed using a five-point scale across correctness, readiness for use, completeness and overall quality. Because the evaluation was conducted in an applied organisational setting, the design prioritised practical relevance over experimental control. The findings should therefore be interpreted as indicative rather than causal.

4.3 Evaluation Results

The controlled dataset produced substantially higher alignment with SME-approved reference responses than the unconstrained knowledge base. As shown in Table 3, both datasets achieved a 100% resolution rate. However, the controlled dataset achieved 89% accuracy against SME-approved references, compared with 44% for the unconstrained knowledge base.

Table 3: Retrieval relevance evaluation

Dataset	Total Prompts	Passed	Failed	Accuracy Score	Resolution Rate
Controlled dataset	9	8	1	89%	100%
Unconstrained dataset	9	4	5	44%	100%

The difference lies not in the ability to generate a response, but in the relevance and alignment of the response with domain expectations.

Output quality results showed a similar pattern. As shown in Table 4, scores are reported as controlled dataset / unconstrained knowledge base. The controlled dataset outperformed the unconstrained knowledge base across all evaluated artefact types. The largest gap appeared in SOW generation, where the unconstrained knowledge base produced outputs that required substantial correction before use.

Table 4: Output quality evaluation

Artefact type	Correctness	Readiness for use	Completeness	Overall quality
Case study	4 / 3	4 / 3	4 / 2	4 / 2.6
Statement of work	5 / 1	4 / 1	4 / 1	4.3 / 1
Client pitch	4 / 3	3.5 / 3	4 / 3	3.8 / 3

Across the evaluated artefacts, outputs generated from the controlled dataset showed stronger correctness, completeness and readiness for use. The SOW results suggest that tasks requiring precise delivery structure and validated organisational knowledge are particularly sensitive to knowledge governance. Overall, the findings suggest that constraining AI systems to governed knowledge inputs can improve retrieval relevance and output quality in AI-assisted workflows.

5. Operational Implementation and Observations

The controlled dataset evaluated in Section 4 was created through an AI-assisted operational implementation across discovery, prioritisation, classification, validation and publication. The implementation was designed to transform fragmented industry knowledge into a curated and reusable knowledge environment suitable for AI-assisted workflows. This shifted the focus of KM from storing knowledge toward governing how knowledge is identified, refined, validated and operationalised.

AI-assisted discovery processes identified approximately 4,500 RCTT-related resources distributed across team repositories, collaboration platforms and personal storage locations. AI-assisted prioritisation was then used to shortlist more than 1,040 resources for detailed review and curation. These figures are programme-level operational measures used to track the transformation of distributed content into governed knowledge assets.

AI-assisted classification and metadata enrichment organised resources across categories such as industry segments, value chains, offerings and reusable delivery patterns. SMEs then validated and refined the shortlisted assets to resolve inconsistencies, remove duplication and ensure alignment with organisational standards. This human-in-the-loop process was central to governance: AI accelerated discovery and structuring, while SMEs retained responsibility for determining which assets were sufficiently reliable for reuse.

Approximately 450 structured knowledge pages were published into the curated Confluence-based knowledge environment used as the controlled dataset for evaluation. AI-assisted workflows were constrained to operate on these validated assets rather than unrestricted organisational repositories.

The implementation also produced operational efficiency gains. Based on internal programme estimates, the delivery timeline was reduced from approximately 12 weeks to 4 weeks. Several classification and enrichment tasks required more than 90% less manual effort than the prior approach. These measures are indicative rather than experimental, but they suggest that AI-assisted curation can make governed KM more operationally feasible at scale.

The intervention also changed the role of domain experts. By reducing manual effort in discovery and structuring, SMEs were able to shift attention from execution-oriented activities toward validation, refinement and governance of reusable knowledge assets. This pattern is consistent with the compounding loop described in Section 3: expert effort is redirected from repeatedly producing or interpreting outputs toward improving the governed knowledge inputs that support future execution.

6. Discussion

The study suggests that the integration of generative AI into organisational knowledge systems exposes structural limitations in retrieval-oriented KM. When knowledge is primarily accessed and interpreted by humans, repositories, taxonomies and search tools may support reuse. However, when knowledge becomes a direct input into AI-assisted execution, weaknesses in structure, quality and governance can be amplified through generated outputs. The central issue is therefore not only whether knowledge can be found, but whether it is sufficiently governed to support reliable execution.

The evaluation provides indicative evidence for this argument. Across retrieval relevance and output quality assessments, outputs generated from the controlled dataset were more aligned with SME-approved references and more usable than outputs generated from the unconstrained knowledge base. This supports the paper's claim that AI reliability depends not only on model capability, but also on the knowledge environment in which AI systems operate.

The findings also clarify the relationship between executable knowledge systems and decay and compounding loops. Executable knowledge systems provide the socio-technical architecture through which knowledge is structured for AI-assisted execution. Decay and compounding loops describe how that architecture evolves over time. Without validation, generated outputs may re-enter the knowledge environment and amplify inconsistency. With governed feedback, outputs can be reviewed, refined and used to improve the knowledge assets that support future execution.

The case study further suggests that scalable knowledge governance may become more operationally feasible when supported by AI-assisted workflows. AI was useful in discovery, prioritisation, classification and metadata enrichment, while domain experts retained responsibility for validation and approval. Rather than removing human expertise, the intervention shifted expertise upstream, where it shaped the knowledge inputs that guide future workflows.

These findings extend prior executable knowledge work by moving from deterministic execution of formalised rules or scripts toward governed execution in generative AI environments. For practice, the implication is that organisations adopting AI need governance mechanisms that define ownership, validation criteria, feedback pathways and controlled updating of knowledge assets. Without these mechanisms, AI may accelerate production while degrading trust; with them, AI-assisted workflows can support compounding improvements in execution reliability.

7. Limitations

This study has several limitations. First, it is based on a single organisational context within a professional services environment, which may limit generalisability to other industries and organisational settings. The RCTT domain was selected because it depends heavily on reusable knowledge assets, but other domains may have different governance needs, knowledge structures or workflow characteristics.

Second, the empirical evidence is exploratory. The evaluation used nine prompts and relied on SME-approved reference responses and judgement-based quality assessment. This approach is appropriate for an applied organisational case study, but it does not provide the level of control associated with experimental methods. The findings should therefore be interpreted as indicative patterns rather than precise causal estimates.

Third, some operational measures, including reductions in delivery timeline and manual effort, were based on internal programme estimates and local benchmarking practices. These measures provide useful evidence of operational feasibility, but future research should apply more standardised measurement methods across workflows.

Fourth, the implementation remained transitional. Although curated knowledge assets, constrained execution and human validation were introduced, governance and feedback mechanisms were not yet fully standardised across all domains. The long-term sustainability of compounding loops therefore remains to be tested.

Finally, the framework of decay and compounding loops has not yet been validated through longitudinal or system-dynamics-based research. Future studies should examine whether these dynamics can be observed over time using measures such as knowledge asset reuse, validation latency, provenance quality, correction rates, output consistency and user trust.

8. Conclusion

Generative AI changes the role of knowledge in organisations. Traditional KM systems were designed primarily to improve access through storage, search and retrieval. However, when knowledge becomes a direct input into AI-assisted execution, access alone is insufficient. Knowledge must be structured, governed and continuously refined so that it can support reliable outputs and decisions.

This paper introduced executable knowledge systems as governed environments in which curated knowledge assets support reliable human and AI-assisted execution. It distinguished this concept from prior executable knowledge approaches by focusing on generative organisational workflows rather than deterministic execution of formalised rules, scripts or queries. The paper also proposed decay and compounding loops as a framework for explaining how AI-enabled knowledge systems evolve under different feedback-governance conditions.

The exploratory case study suggests that curated and validated knowledge inputs can improve retrieval relevance and output quality compared with unconstrained organisational knowledge sources. It also indicates that AI-assisted curation can make large-scale knowledge governance more operationally feasible when combined with SME validation and ownership.

The central contribution of the paper is to reframe KM as a system design challenge for the AI era. Reliable AI-assisted execution depends not only on model capability, but on the quality, governance and feedback mechanisms of the knowledge systems that support it.

Acknowledgements

The author acknowledges the contributions of the AI-First Knowledge Management team and the Research and Intelligence team at Thoughtworks. The author thanks Jessie Jie Xia and the Office of the CIO for their sponsorship of the knowledge management transformation described in this paper.

Ethics Declaration: This study is based on organisational observations conducted within a professional setting. It did not involve human subjects research requiring formal ethical approval. No personal or sensitive data was collected or analysed, and all information is presented in aggregated and non-identifiable form.

AI Declaration: AI tools were used to support the drafting, structuring and refinement of language in this paper. All AI-generated outputs were critically reviewed, edited and validated by the author. AI was not used to generate original research findings, core concepts or interpretations. The ideas, framework and conclusions presented in this paper are the result of the author's own analysis and professional experience.

References

Alavi, M. and Leidner, D.E. (2001) "Review: Knowledge Management and Knowledge Management Systems: Conceptual Foundations and Research Issues", MIS Quarterly, Vol 25, No. 1, pp 107–136.

APQC (2025) "Future Ready Knowledge Management: 10 Practices That Work", APQC, [online], available at: https://www.apqc.org/resource-library/resource-listing/future-ready-knowledge-management-10-practices-work

Davenport, T.H. and Prusak, L. (1998) Working Knowledge: How Organizations Manage What They Know, Harvard Business School Press, Boston.

Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H. et al. (2025) "A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions", ACM Computing Surveys, DOI: 10.1145/3703155.

Kirchner, K., Bolisani, E., Kassaneh, T.C., Scarso, E. and Taraghi, N. (2025) "Generative AI Meets Knowledge Management: Insights From Software Development Practices", Knowledge and Process Management, DOI: 10.1002/kpm.70004.

Lee, Y.W., Strong, D.M., Kahn, B.K. and Wang, R.Y. (2002) "AIMQ: A Methodology for Information Quality Assessment", Information & Management, Vol 40, No. 2, pp 133–146.

Li, J., Cheng, X., Zhao, W.X., Nie, J.Y. and Wen, J.R. (2023) "HaluEval: A Large-Scale Hallucination Evaluation Benchmark for Large Language Models", arXiv, DOI: 10.48550/arXiv.2305.11747.

Liu, S., Wright, A.P., McCoy, A.B., Huang, S.S., Steitz, B. and Wright, A. (2025) "Detecting Emergencies in Patient Portal Messages Using Large Language Models and Knowledge Graph-Based Retrieval-Augmented Generation", Journal of the American Medical Informatics Association, Vol 32, No. 6, pp 1032–1042.

Nonaka, I. and Takeuchi, H. (1995) The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation, Oxford University Press, New York.

Quimbaya, A.P., Muñoz, O., Londoño, D., Bohórquez, R., García, O.M., González, R.A. et al. (2014) "An Executable Knowledge Base for Clinical Practice Guideline Rules", Procedia Technology, Vol 16, pp 1446–1455.

Storey, V.C. (2025) "Knowledge Management in a World of Generative AI: Impact and Implications", ACM Transactions on Management Information Systems, DOI: 10.1145/3719209.

Yun, B., Feng, D., Chen, A.S., Nikzad, A. and Salehi, N. (2025) "Generative AI in Knowledge Work: Design Implications for Data Navigation and Decision-Making", Proceedings of the CHI Conference on Human Factors in Computing Systems, DOI: 10.1145/3706598.3713337.

Zheng, Z., Zhou, B., Zhou, D., Soylu, A. and Kharlamov, E. (2022a) "ExeKG: Executable Knowledge Graph System for User-Friendly Data Analytics", Proceedings of the 31st ACM International Conference on Information and Knowledge Management, pp 5064–5068, DOI: 10.1145/3511808.3557195.

Zheng, Z., Zhou, B., Zhou, D., Zheng, X., Cheng, G., Soylu, A. et al. (2022b) "Executable Knowledge Graphs for Machine Learning: A Bosch Case of Welding Monitoring", International Semantic Web Conference, Springer, Cham, pp 791–809.

Zhu, X., Xie, Y., Liu, Y., Li, Y. and Hu, W. (2025) "Knowledge Graph-Guided Retrieval Augmented Generation", arXiv, DOI: 10.48550/arXiv.2502.06864.