



International Journal of Knowledge Management and Leadership

Journal homepage: <https://trpub.org/journal.php?id=obXzcXzc8r>

TRPub

IJKML

International Journal of
Knowledge
Management and
Leadership

Vol. 1 Issue 1 2026

Perspective

Open Access

When Enterprise Chat Assistants Become the Organisational Memory, Who Curates It

Tuan Salgado^{1*}, Dongwoo Kowalski², Erik Nakamura³

¹Department of Business, Levantine University Centre for Business, Kenya

²Department of Business, Orinoco Academy of Business, Croatia

³Department of Business, Douro Institute of Business, Taiwan

Article Info

Volume 1, Issue 1, 2026

[doi: 10.62715/ijkml.2026.v1.i1.pp27-32](https://doi.org/10.62715/ijkml.2026.v1.i1.pp27-32)

Received: 20 October 2025

Accepted: 23 December 2025

Published: 19 February 2026

Abstract

Organisations are connecting conversational assistants built on large language models to their internal documents, message archives and ticketing systems, and employees are beginning to ask the assistant where they once asked a colleague or searched an intranet. This article argues that the change is larger than a new interface to old stores. When the assistant becomes the ordinary route to what the organisation knows, it becomes, in effect, the organisational memory, and the decisions that shape its answers become decisions about what the organisation remembers. Those decisions include what is admitted to the corpus, what is withdrawn when it is superseded, how contradictions between sources are settled and who checks that answers are right. In most deployments they are made by default, by whoever configured the connectors. The article sets out why conversational retrieval makes curation both more necessary and less visible than it was for document repositories, considers the objections that the technology will curate itself or that curation cannot scale, and proposes that responsibility be divided among content owners, a stewardship function and the users themselves, with records management and knowledge management standards as the starting point. The argument is conceptual and draws on published research; no deployment was studied.

Keywords: *Organisational memory, Enterprise chat assistants, Large language models, Knowledge curation, Retrieval-augmented generation, Knowledge governance*

© 2026 Tuan Salgado et al. This is an open access article under the CC BY license (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

1. Introduction

For three decades the standard account of organisational memory has been that of Walsh and Ungson (1991), who located it in a set of retention bins: individuals, culture, the transformations embodied in work procedures, structures, the physical ecology of the workplace and archives outside the organisation. Information systems were added to the picture soon afterwards, and the knowledge management movement of the following years invested heavily in repositories, intranets and directories of expertise. Throughout, the systems were stores. A person searched, received a list of documents, opened some of them and judged what to make of each.

* Corresponding author: Tuan Salgado, Department of Business, Levantine University Centre for Business, Kenya. E-mail: tuan.salgado@levantine.edu

© 2026 Tuan Salgado et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

The enterprise chat assistant changes the last step. An employee types a question and receives an answer composed for the occasion from whatever the system has retrieved. Early evidence suggests the arrangement can be valuable. [Brynjolfsson et al. \(2025\)](#), studying the introduction of a generative assistant among customer support agents, found gains in productivity concentrated among the less experienced, and interpreted them as the spread of practices that the best agents had developed. That is a description of a memory at work. The question this article raises is who is looking after it.


2. Memory as a Process, Not a Store

The retention-bin model was criticised on a point that has become newly relevant. [Ackerman and Halverson \(2000\)](#), after observing how a telephone helpline answered its callers, argued that organisational memory is not one thing in one place. It is a set of memories of different kinds, held in people, records and tools, that are brought together in the act of answering, and each time a piece of information moves from one holder to another it is stripped of some context and must be supplied with context again by the person who uses it. On this account remembering is work, and much of the work is the judgement by which a competent member decides that a given record is current, applies to the case and can be trusted.

Curation, in the sense used here, is the organised form of that judgement. Records managers have long practised it. The international standard on records management asks that records have authenticity, reliability, integrity and usability, and that their retention and disposal be decided and not left to chance ([ISO, 2016](#)). The standard on knowledge management systems speaks of retaining knowledge and of handling knowledge that is outdated or invalid ([ISO, 2018](#)). [Markus \(2001\)](#) showed how much effort is needed to make a record usable by someone other than its author. None of these concerns is new. What is new is that an assistant now performs part of the judgement itself, and performs it out of sight.

3. What Conversational Retrieval Changes

Most enterprise assistants follow the pattern that [Lewis et al. \(2020\)](#) named retrieval-augmented generation. The question is used to fetch passages from an index of the organisation's content, and a language model writes an answer conditioned on those passages. Figure 1 shows the layers involved. Each layer makes selections that the user does not see: which sources were admitted to the index, which passages were ranked highest, and how the model reconciled or ignored what they said.



Three consequences follow. The first is the loss of provenance cues. A search result arrives with a title, an author, a date and a location, from which an experienced employee infers a good deal: that the page

belongs to a team known to be careless, or that the policy predates a reorganisation. A composed answer blends its sources into a single voice, and even where citations are offered the path of least effort is to read the answer and stop. The second is fluency. Language models produce text that is confident and well formed whether or not it is supported, and the tendency to generate unsupported content is extensively documented (Ji et al., 2023). Bender et al. (2021) warned that readers take coherent text to be the product of understanding and extend to it a trust it has not earned.

The third consequence concerns the user's vigilance. Research on automation established long ago that people come to rely on aids that are usually right and cease to check them, a pattern Parasuraman and Riley (1997) called misuse and that later writers have termed automation bias. Sparrow et al. (2011) found that people who expect to have access to information later remember where to find it and not what it says, extending to the internet the transactive memory that Wegner (1987) had described among people. An employee who treats the assistant as the place where answers are kept will have less reason to remember the answer, the source or the colleague who knew. Dell'Acqua et al. (2023), in a field experiment with consultants, found that performance improved on tasks within the capability of the model and worsened on a task beyond it, where those given the model were less likely to reach a correct solution than those working without it. The boundary between the two kinds of task was not apparent to the people using the tool.

4. The Curation Gap

Taken together, these changes mean that the quality of what the organisation remembers now depends on a set of tasks that the deployment project rarely assigns. Table 1 lists them. In practice the corpus is defined by which systems had connectors available, exclusions are driven by access permissions and not by accuracy, superseded documents remain indexed beside their replacements, and nobody is charged with noticing that the assistant gives two teams different answers to the same question.

Table 1: Curation tasks raised by an enterprise chat assistant, why each matters and where responsibility might sit		
Curation task	Why it matters	Suggested owner
Admission to the corpus	Casual chat and drafts acquire the standing of record once indexed	Stewardship function with content owners
Authority ranking	An approved policy should outweigh a remark in a message thread	Content owners, by designation
Retirement	Superseded material continues to be retrieved unless withdrawn	Content owners, on a review cycle
Resolving contradictions	The model will blend or pick between conflicting sources unpredictably	Stewardship function
Checking answers	Fluent errors pass unnoticed without sampling and user reports	Stewards and users jointly
Recording gaps	Questions the corpus cannot answer show what was never written down	Stewardship function
Guarding against feedback	Generated text saved as documents re-enters the corpus as a source	Platform owner

Two of these tasks deserve comment. Admission matters because indexing erases a distinction that organisations used to maintain without effort. A remark in a project channel and a signed procedure were different kinds of thing, kept in different places and read with different expectations. Once both are passages in an index they compete on similarity to the question. The feedback problem is newer. Employees paste assistant output into reports and wiki pages, and those pages are indexed in turn. Shumailov et al. (2024) showed that models trained repeatedly on generated data degrade. An enterprise

index is not a training set, and the analogy is loose, but the direction of the concern is the same: a memory that increasingly consults its own earlier answers drifts from its sources.

There is also the matter of what was never in the corpus. The retention bins of Walsh and Ungson included individuals and culture, and the knowledge held there is by definition not in the documents. An assistant can make written knowledge so convenient that the habit of asking a person weakens, and with it the occasions on which unwritten knowledge was passed on. [de Holan and Phillips \(2004\)](#) described how organisations forget what they fail to exercise. A memory that answers every question it can, and gives no sign of the questions it cannot, makes that kind of forgetting hard to notice.

5. Objections

It may be objected that the technology will solve the problem it creates. Retrieval can be tuned to prefer recent and authoritative sources, and models can be instructed to cite and to decline. These measures help, and they presuppose curation: a system can prefer authoritative sources only if someone has said which sources are authoritative, and it can flag a superseded policy only if supersession is recorded. The tools make curatorial decisions cheaper to apply. They do not make them. There is a further difficulty in asking a system to audit itself: the assistant has no access to the facts against which an organisational record would have to be checked, only to other records.

A second objection is that intranet search never had curators and organisations managed. The reply is that search left judgement with the reader, who saw the sources, whereas the assistant takes over part of that judgement, and that the neglected intranet was in any case a familiar failure of knowledge management, not a success. A third objection is cost. No organisation can review everything its staff write. This is true, and it is the reason for curating by class and not by item: deciding which repositories count as record, which are admitted as context and which are excluded, and concentrating human review on the small body of material that carries authority.

A last objection holds that curation centralises control over what the organisation is allowed to know. The risk is real. [Orlikowski \(1992\)](#), studying the introduction of groupware, showed that a collaborative technology is shaped by the reward structures and assumptions of the place that adopts it, and an assistant curated solely by a central office will reflect that office's view. The answer is to distribute the work and make it visible, and not to leave it undone.

6. Implications for Practice and Research

A workable division of responsibility has three parts. Content owners, meaning the functions that issue policies, standards and procedures, designate their authoritative documents and keep them current, as records management already requires. A stewardship function, which may be a small team drawn from knowledge management, records and information technology, decides admission by class, samples answers, follows up contradictions and reports the questions that went unanswered. The last of these duties is easily overlooked and may be the most useful, since a log of questions the assistant could not answer is a direct inventory of what the organisation knows only through its people. Users are given the means and the expectation to check sources and to flag wrong answers, which makes their vigilance part of the system. Regulation points the same way. The risk management framework of the National Institute of Standards and Technology treats governance as a function that runs through the whole life of an artificial intelligence system ([NIST, 2023](#)), and the European Union's Artificial Intelligence Act requires organisations that deploy such systems to ensure a sufficient level of literacy among the staff who use them, and requires human oversight of systems it classes as high-risk ([European Union, 2024](#)).

For researchers the assistant offers an unusual opportunity. For the first time the questions that members of an organisation ask of its memory are recorded, together with the answers they receive. With appropriate safeguards, such records could show what employees need to know and cannot find, how answers change as the corpus changes and whether reliance on the assistant alters whom people consult. Studies comparing curated and uncurated deployments would test the claim made here, which rests on argument from adjacent research and not on observation of the systems in question.

7. Conclusion

An organisation that routes its questions through a chat assistant has acquired a memory with new properties: it answers in one voice, hides its sources unless pressed, never admits that a matter was simply not written down, and is trusted in proportion to its fluency. Such a memory needs keepers. The tasks are recognisable from records and knowledge management and can be shared among content owners, stewards and users. What cannot continue is the present arrangement in many deployments, in which the memory of the organisation is shaped by the defaults of a connector and nobody has been asked to mind it.

Conflict of interest

To be declared by the author(s) before publication.

Funding

To be declared by the author(s) before publication.

References

- Ackerman, M.S. and Halverson, C.A. (2000). [Reexamining organizational memory](#). *Communications of the ACM*, 43(1): 58-64. <https://doi.org/10.1145/323830.323845>
- Bender, E.M., Gebru, T., McMillan-Major, A. and Shmitchell, S. (2021). [On the dangers of stochastic parrots: Can language models be too big?](#) *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610-623. <https://doi.org/10.1145/3442188.3445922>
- Brynjolfsson, E., Li, D. and Raymond, L.R. (2025). [Generative AI at work](#). *The Quarterly Journal of Economics*, 140(2): 889-942.
- de Holan, P.M. and Phillips, N. (2004). [Remembrance of things past? The dynamics of organizational forgetting](#). *Management Science*, 50(11): 1603-1613.
- Dell'Acqua, F., McFowland, E., Mollick, E.R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraye, L., Candelon, F. and Lakhani, K.R. (2023). [Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality](#). Harvard Business School Working Paper 24-013.
- European Union (2024). [Regulation \(EU\) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence \(Artificial Intelligence Act\)](#). Official Journal of the European Union, L series, 12 July 2024.
- ISO (2016). [ISO 15489-1:2016 Information and documentation: Records management. Part 1: Concepts and principles](#). International Organization for Standardization, Geneva.
- ISO (2018). [ISO 30401:2018 Knowledge management systems: Requirements](#). International Organization for Standardization, Geneva.
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y.J., Madotto, A. and Fung, P. (2023). [Survey of hallucination in natural language generation](#). *ACM Computing Surveys*, 55(12): 1-38. <https://doi.org/10.1145/3571730>
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S. and Kiela, D. (2020). [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). *Advances in Neural Information Processing Systems*, 33: 9459-9474.
- Markus, M.L. (2001). [Toward a theory of knowledge reuse: Types of knowledge reuse situations and factors in reuse success](#). *Journal of Management Information Systems*, 18(1): 57-93.
- NIST (2023). [Artificial Intelligence Risk Management Framework \(AI RMF 1.0\)](#). NIST AI 100-1. National Institute of Standards and Technology, Gaithersburg, MD. <https://doi.org/10.6028/NIST.AI.100-1>
- Orlikowski, W.J. (1992). [Learning from Notes: Organizational issues in groupware implementation](#). *Proceedings of the 1992 ACM Conference on Computer-Supported Cooperative Work*, 362-369.

<https://doi.org/10.1145/143457.143549>

Parasuraman, R. and Riley, V. (1997). [Humans and automation: Use, misuse, disuse, abuse](#). *Human Factors*, 39(2): 230-253. <https://doi.org/10.1518/001872097778543886>

Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R. and Gal, Y. (2024). [AI models collapse when trained on recursively generated data](#). *Nature*, 631: 755-759.

<https://doi.org/10.1038/s41586-024-07566-y>

Sparrow, B., Liu, J. and Wegner, D.M. (2011). [Google effects on memory: Cognitive consequences of having information at our fingertips](#). *Science*, 333(6043): 776-778.

<https://doi.org/10.1126/science.1207745>

Walsh, J.P. and Ungson, G.R. (1991). [Organizational memory](#). *Academy of Management Review*, 16(1): 57-91. <https://doi.org/10.5465/amr.1991.4278992>

Wegner, D.M. (1987). [Transactive memory: A contemporary analysis of the group mind](#). In: Mullen, B. and Goethals, G.R. (eds), *Theories of Group Behavior*. Springer, New York, 185-208.