

From Text to Knowledge Gems: Design Principles and a Hybrid AI Approach for Sustainable Knowledge Curation

Lucia Cvitanic¹, Heidelinde Hobel¹, Verena Parzer¹, Denis Havlik¹ and Ghazal Etminan¹

¹*Austrian Institute of Technology (AIT)*

Abstract

This paper introduces Knowledge Gems, a class of compact, semantically structured artifacts, and presents a hybrid AI approach designed to support expert knowledge curation by semi-automating Gem creation from heterogeneous documentation sources. It formulates five core design principles for sustainable knowledge systems, addressing key challenges in knowledge structuring, semantic interoperability, scalability, and long-term reuse. Building on these principles, the paper proposes a modular agent-based pipeline that separates retrieval, generation, and human-in-the-loop validation and refinement of generated content. Initial experiments and expert observations suggest that, during early-stage knowledge curation, knowledge graphs can provide substantial practical value through provenance tracking, schema alignment, and systematic validation, even before more advanced semantic reasoning capabilities become feasible. While the proposed framework shows promise as a first-draft generation system, the results reveal persistent challenges in long-form synthesis, source alignment, and automated quality assurance. These findings highlight the continued importance of human oversight and further validation to ensure accuracy, reliability, and traceability. This work outlines an architectural blueprint for community-driven knowledge engineering in complex, interdisciplinary environments and contributes early design knowledge on the application of hybrid AI approaches for sustainable knowledge systems.

Keywords

Knowledge Graphs, Generative AI, Retrieval-Augmented Generation, Hybrid AI, Knowledge Engineering, Knowledge Curation, Agent-based Systems, Provenance-aware AI, Human-in-the-loop AI

1. Introduction

Large collections of heterogeneous documents often contain valuable insights, yet these remain difficult to discover and reuse due to fragmentation across formats, contexts, and levels of abstraction [1, 2]. Even when knowledge is extracted, long-term maintainability remains a challenge: many knowledge systems become static once projects end, leading to outdated content, rigid schemas, and increasing maintenance effort [3]. This creates a need for approaches that support continuous, structured, and evolvable knowledge curation.

Semantic technologies and knowledge graphs provide a principled foundation for structured and interoperable knowledge representation [4, 5], but require substantial effort for schema design, population, and ongoing maintenance [2]. Large language models (LLMs) offer scalable extraction and generation from unstructured text [6], yet lack guarantees regarding structure, consistency, grounding, and traceability [7, 8]. Hybrid approaches such as RAG and graph-augmented retrieval aim to improve grounding [9], but remain primarily retrieval-centric and do not fully address schema-driven organization or long-term structured knowledge persistence. Taken together, these existing methods address important components of the problem space, but do not fully address the challenge of creating, validating, and sustaining structured knowledge over time. This gap motivates the work presented in this paper.

The work presented here is carried out within the broader NEB Junction initiative [10, 11], an EU project that provides a rich, real-world environment with diverse documentation, evolving requirements, and multiple stakeholder groups¹. This setting offers both the complexity and practical constraints needed to explore how hybrid AI approaches can support sustainable knowledge curation in practice. Against this background, this paper presents the first phase of our investigation into hybrid

Posters, Demos, Blue Sky, and Tutorials at SEMANTiCS 2026, September 2026, Ghent, Belgium



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://nebjunction.eu>

AI-supported knowledge curation. Drawing on experiences gained while creating knowledge artefacts from heterogeneous project documentation, we derive five design principles for maintainable hybrid AI knowledge systems. These principles, referenced throughout the paper, are: (1) structured intermediate representations, (2) incremental knowledge graph integration, (3) modular decomposition of retrieval and generation, (4) validation-driven workflows, and (5) support for long-term schema and content evolution.

2. Knowledge Gems and the Challenge

We operationalize these five design principles through the use of Knowledge Gems: compact, structured and validated knowledge artefacts that package practical insights, solutions, methods, or experiences from project documentation into an accessible form for reuse by stakeholders. Knowledge Gems are designed to support both human consumption and machine processing, while preserving links to the underlying evidence and sources to enable traceability and verification.

In the context of the NEB Junction initiative, Knowledge Gems package practical, validated insights from New European Bauhaus projects and are intended to help stakeholders learn from existing work, get inspired, and take action. They may capture different types of reusable knowledge, including solutions, tools, methodologies, case studies, and lessons learned.

To support discovery and reuse across different stakeholder needs and contexts, Gems are expressed using a consistent structure and enriched with semantic metadata and provenance information. Figure 1 schematically depicts a concrete Knowledge Gem derived from one of the NEB-STAR demonstrators² and highlights its provenance structure by expanding one of its branches.




Figure 1: Structural view of a Knowledge Gem: Expanded provenance links illustrate how generated Gem elements can be traced back to previously extracted information in the knowledge graph and source document chunks, supporting validation, traceability, and incremental knowledge base evolution.

The example illustrates a simplified, initial version of the intended structure and purpose of Knowledge Gems. The goal of the Gem concept is threefold: (i) enable stakeholders to consume information expressed in familiar language and domain terminology; (ii) support semantic search, comparison, and aggregation through a consistent schema; and (iii) ensure traceability by preserving explicit links to extracted information and original sources for verification and deeper exploration.

A related challenge is to efficiently create and maintain such structured knowledge artefacts from large and heterogeneous collections of project documentation. Relevant information is distributed across

²<https://nebstar.eu/>

documents that differ in format, structure, terminology, and level of detail, while the resulting Gems need to remain consistent, traceable, and adaptable as project knowledge and stakeholder needs evolve. Manual curation can provide the required quality and contextual understanding, but becomes difficult to scale and maintain over time. This creates a need for approaches that can support the semi-automatic creation and continuous evolution of Knowledge Gems while retaining human oversight and reliable links to source evidence.

3. From Initial Experiments to a Hybrid Approach

The initial phase of the work focused on identifying the structural and quality requirements for Knowledge Gems. NEB innovation managers manually curated a small set of high-quality Gems and developed a detailed content template defining structure, tone, and vocabulary. These exemplars established editorial norms and served as few-shot examples for large language models. Prompt engineering was used to encode these guidelines into the generation process, enabling rapid experimentation while keeping costs low and avoiding premature commitment to evolving schemas.

Initial AI-supported Gem generation experiments were conducted using an AI-expert system platform, originally developed in the MAIA project [12]. The platform provided a straightforward retrieval-augmented generation (RAG) pipeline with components for summarization, question answering, and semantic annotation. While this baseline enabled early prototyping, applying it in the NEB context revealed several intertwined challenges: (i) heterogeneous and fragmented source materials, (ii) the complexity of transforming documents into structured knowledge, (iii) evolving requirements and schemas, and (iv) the need to support non-expert users.

Addressing these challenges required moving beyond a retrieval-centric RAG pipeline, towards a workflow that explicitly separates retrieval, structured generation, validation, and human review. This motivated the development of a hybrid AI workflow that combines retrieval, knowledge graphs, constrained LLM generation, and validation.

Figure 2 presents the resulting agent-based pipeline. Rather than relying on end-to-end LLM generation, the workflow separates retrieval, generation, and human validation into distinct stages. This modular design improves transparency, allows individual steps to be repeated independently, and facilitates incremental refinement.

In the first stage, project deliverables are ingested, chunked, embedded, and stored in a hybrid knowledge base combining vector search with lightweight graph structures. The graph captures entities and relationships, while the vector index supports semantic retrieval. The pipeline uses the Mistral API as the LLM backend and Neo4j as a combined vector- and graph-based knowledge base, enabling unified semantic and structural retrieval.

A retrieval agent then assembles context for downstream processing. Its role is deliberately constrained: rather than generating new interpretations, it identifies relevant document passages, structured graph entities, previously curated Knowledge Gems, and task-specific instructions derived from the Gem schema. By assembling this material into a structured prompt, the agent helps reduce unsupported generation and keeps subsequent generation anchored in identifiable source evidence.

A generation agent transforms the retrieved context into a structured Knowledge Gem aligned with the predefined schema. This includes producing concise narrative text, extracting or inferring semantic metadata, and populating the output model in a form suitable for ingestion into the knowledge base. Generated Gems undergo syntactic and semantic validation to verify completeness and schema compliance. Validation failures trigger selective re-execution of the affected step rather than restarting the entire workflow.

Human experts remain integral to the process. They review generated Gems, refine prompts, and evolve the schema as project requirements change. Human oversight ensures correctness, maintains editorial quality, and guides the gradual expansion of automation as the knowledge base grows.

Three operational considerations guided the implementation: minimal viable automation, ensuring that only stable and well-defined tasks are automated; explicit grounding and traceability, requiring that




Figure 2: Agent-based pipeline for semi-automatic Knowledge Gem creation.

all generated content be linked to source evidence; and human control and validation, maintaining expert oversight during early project phases and supporting schema evolution over time. These considerations also provide a foundation for the design principles discussed in the following section.

4. Design Principles Emerging from Hybrid AI Knowledge Curation

As we began experimenting with semi-automatic creation of Knowledge Gems, one pattern became clear almost immediately: the system could reliably produce short, structured elements such as titles, keywords, and summaries, but struggled with longer, more interpretive fields. Narratives, lessons learned, and impact descriptions often arrived incomplete or overly generic. This experience highlighted the importance of *structured intermediate representations* (Principle 1). Introducing additional structuring steps, particularly topic extraction prior to generation, provided clearer semantic anchors, improved consistency, and reduced variability in generated outputs.

One of the main insights emerging from the experiments was that knowledge graphs provide immediate value through structure, validation, provenance, and schema alignment, while graph-based reasoning becomes more relevant as the knowledge base matures. Although we initially expected graph-based reasoning to play a central role, early benefits emerged primarily from keeping the system organized and grounded. The experience reinforced the value of *incremental knowledge graph integration* (Principle 2): rather than building a complete graph upfront, gradually enriching the graph with validated entities and relationships proved both practical and beneficial. Sparse population limited grounding capabilities, but even a lightweight graph improved consistency in fields tied to explicit entities such as authors, topics, and project names.

As the system evolved, the modular pipeline design proved essential. Separating retrieval, generation, validation and structuring made the workflow easier to understand, debug, and extend. When a step failed, it was possible to rerun only the affected component rather than restarting the entire workflow. This supported iterative refinement and adaptation to changing project requirements, highlighting the value of *modular decomposition of retrieval and generation* (Principle 3).

Evaluation surfaced recurring challenges that shaped the next phase of development. Source-

alignment errors were among the most critical: occasionally, generated content was coherent and factually plausible but grounded in the wrong source material. For example, in Figure 1, a tool from another project could be associated with the current Knowledge Gem due to the retrieval of semantically similar but contextually different document fragments. Long-form synthesis remained less stable than extraction-oriented tasks despite strict prompting, highlighting the remaining limits of constraint-driven generation for interpretive content. These observations reinforced the importance of *validation-driven workflows* (Principle 4). Structured outputs, provenance links, schema checks, and iterative correction mechanisms became essential for prompt refinement, error detection, and schema evolution. At this stage, the pipeline functioned best as a first-draft generator—accelerating content creation while still requiring human oversight for correctness, attribution, and grounding.

Finally, the project environment itself demonstrated the importance of designing for change. New stakeholder requirements, additional document sources, and evolving information needs repeatedly required adjustments to both the Gem structure and the underlying knowledge model. This experience underscored the need to support *long-term schema and content evolution* (Principle 5). Lightweight schemas, human-centered templates, and a hybrid symbolic–subsymbolic architecture provided sufficient flexibility to accommodate change while preserving consistency and traceability.

Taken together, these experiences led to the five design principles introduced in Section 1. While the current system remains human-supervised, these principles provide a foundation for more scalable and sustainable AI-supported knowledge curation.

5. Limitations and Outlook

Initial experiments indicate that semi-automatic Knowledge Gem creation is feasible when combined with lightweight semantic structures, modular orchestration, and constraint-driven prompting. At the same time, the current results highlight important limitations, particularly in long-form synthesis, source alignment, and automated quality assurance. The findings and design principles presented in this paper are based on initial experiments and qualitative expert observations and have not yet been systematically validated. Several directions for further research therefore emerge.

The first direction concerns the continued evolution of the Knowledge Gem schema itself. Early iterations revealed where additional structure improves consistency and where flexibility is needed to accommodate different types of knowledge. Future work will refine the schema to better support grounded generation, provenance, and long-term maintainability. In parallel, the agent-based pipeline can be extended with specialized components for tasks such as entity normalization, cross-document alignment, and context validation, with the aim of reducing manual correction efforts while preserving human oversight.

A second direction is systematic quantitative evaluation. Future experiments will introduce quantitative measures for accuracy, completeness, grounding quality, and user effort. Such evaluation will allow us to assess the effectiveness of individual design choices, identify where automation provides the greatest benefit, and better understand the balance between human oversight and machine assistance.

The role of the knowledge graph also warrants further investigation. As the graph structure and validation processes mature, it may provide stronger support for semantic retrieval, disambiguation, and structural guidance. Future work will examine whether greater graph coverage improves grounding and generation quality, and when graph-based reasoning becomes beneficial.

Finally, the long-term applicability of the approach needs to be tested beyond the current project context. Its use of lightweight schemas, modular components, and explicit validation is intended to support adaptation to different domains and community-driven environments, but its scalability and transferability remain to be demonstrated. Future research therefore, needs to investigate how the approach performs across larger and more diverse document collections, changing schemas, and stakeholder groups with different levels of expertise.

Overall, the next phase of work will focus on moving from a promising, human-supervised first-draft generation system towards a more systematically evaluated and increasingly reliable approach to

sustainable knowledge curation.

Acknowledgements and Declaration on Generative AI

The work leading to this paper has received funding from the European Union's Horizon Europe research and innovation programme under the NEB Junction project, grant agreement No. 101215634, and its precursor, the MAIA project, grant agreement No. 101056935.

Special thanks go to Clémentine Coujard and Karine Laffont-Eloire (Dowel Innovation), who served as innovation managers and contributed to the design of the Knowledge Gems and the preparation of the exemplary Gems used in the case study.

During the preparation of this work, the authors used Microsoft Copilot in order to: check grammar and spelling, paraphrase and reword. After using this service, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] L. M. Augusto, From symbols to knowledge systems: A. newell and h. a. simon's contribution to symbolic ai, *AI & Society* 2 (2021) 29–62. URL: <https://philpapers.org/archive/AUGFST-2.pdf>.
- [2] K. M. Wiig, Knowledge management: An emerging discipline rooted in a long history, *Knowledge Horizons: The Present and the Promise of Knowledge Management* (2000).
- [3] H. Hobel, J. Heurix, A. Anjomshoaa, E. Weippl, Towards security-enhanced and privacy-preserving mashup compositions, in: L. J. Janczewski, H. B. Wolfe, S. Sheno (Eds.), *Security and Privacy Protection in Information Processing Systems*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 286–299.
- [4] B. Motik, P. F. Patel-Schneider, B. Parsia, C. Bock, A. Fokoue, P. Haase, R. Hoekstra, I. Horrocks, A. Ruttenberg, U. Sattler, M. Smith, *OWL 2 Web Ontology Language: Structural Specification and Functional-Style Syntax (Second Edition)*, W3C Recommendation, World Wide Web Consortium (W3C), 2012. URL: <https://www.w3.org/TR/owl2-syntax/>, accessed: 2025-11-20.
- [5] R. Widuri Sri MP, Semantic approaches in knowledge management: A systematic literature review, *ELite Journal : International Journal of Education, Language and Literature* 5 (2025) 34–40. URL: <https://journal.unesa.ac.id/index.php/elite/article/view/39839>.
- [6] T. B. Brown, B. Mann, N. Ryder, et al., Language models are few-shot learners, *Advances in Neural Information Processing Systems* 33 (2020) 1877–1901. URL: <https://arxiv.org/abs/2005.14165>.
- [7] Z. Ji, et al., Survey of hallucination in large language models, *arXiv preprint* (2023). URL: <https://arxiv.org/abs/2305.13534>.
- [8] C. Sharma, Retrieval-augmented generation: A comprehensive survey of architectures, enhancements, and robustness, in: *Frontiers in Computer Science (arXiv)*, volume abs/2506.00054, 2025. doi:10.48550/arXiv.2506.00054.
- [9] J. Ren, A comprehensive survey of knowledge graph-augmented generation (graph-rag) for trustworthy large language models, *Advances in Engineering Innovation* 17 (2026).
- [10] European Commission, New european bauhaus: About the initiative, https://new-european-bauhaus.europa.eu/about/about-initiative_en, 2025. NEB at a glance and a timeline of main activities.
- [11] European Commission, Neb junction: The new european bauhaus hub for results & impact, <https://cordis.europa.eu/project/id/101215634>, 2025. Grant Agreement ID: 101215634, Funded under Horizon Europe, Climate, Energy and Mobility.
- [12] D. Havlik, A. Chettakattu, V. Karagiannis, G. Etminan, P. Pale, Harnessing genai for climate change knowledge management: Use cases, system design, and experimental evaluation, *Research Square preprint v1* (2025). URL: <https://www.researchsquare.com/article/rs-6344237/v1>, preprint.