

Gen-AI for Knowledge Management: Automated Knowledge Base Creation and Context-Aware Q&A Systems

Monika Malik

Lead Data/AI Engineer, AT&T, TX, USA

ARTICLE INFO

Received: 26 Dec 2024

Revised: 14 Feb 2025

Accepted: 22 Feb 2025

ABSTRACT

Introduction: Knowledge Management (KM) is critical for organizations to capture, organize, and retrieve information efficiently. Generative AI is revolutionizing KM by enabling automatic creation of knowledge content and highly intelligent query answering systems.

Objectives: This paper explores two interrelated advancements: (1) Automated Knowledge Base Creation, where generative models summarize and distill information from large corpora or documentation into structured knowledge articles or knowledge graphs, reducing the manual effort of building knowledge repositories; and (2) Context-Aware Q&A Systems, where large language models (LLMs) deliver precise answers to user queries by understanding context and retrieving relevant knowledge, effectively serving as intelligent assistants that can reason over an organization's data.

Methods: We examine state-of-the-art techniques such as retrieval-augmented generation (RAG) for grounding AI answers in enterprise data [blogs.nvidia.com](https://blogs.nvidia.com/blogs.nvidia.com), and outline how open-source tools (like Haystack, LangChain) and platforms (IBM Watson, GPT-4 with plugins) implement these functionalities.

Results: The paper also addresses the challenges of maintaining knowledge accuracy, avoiding hallucinations, ensuring security of proprietary information, and mitigating bias in AI-generated knowledge. Through case studies and literature review, we demonstrate how generative AI-based KM solutions can significantly improve information discovery and decision-making, by providing up-to-date, contextually relevant knowledge on demand.

Conclusions: We conclude with future directions, such as integrating knowledge graphs with LLMs for reasoning and the evolving role of human knowledge curators in an AI-augmented knowledge lifecycle.

Keywords: Knowledge Management, Generative AI, Knowledge Base, Question-Answering, Retrieval-Augmented Generation, Knowledge Graph, Enterprise AI, Contextual QA

INTRODUCTION

In modern organizations, vast amounts of information – from technical documentation and wikis to customer support logs and employee expertise – constitute the organizational knowledge. **Knowledge Management (KM)** is the discipline that deals with collecting, structuring, and retrieving this knowledge to support learning and decision-making. Traditional KM systems rely on manual curation: experts write articles for knowledge bases, and users search or navigate through hierarchical content. This approach often struggles to keep pace with the growth of information, leading to outdated or incomplete knowledge bases and frustrated users who cannot find answers quickly. Generative AI offers a transformative solution by automating significant parts of the KM cycle.

Generative AI in Knowledge Management involves using AI systems that can generate human-like text to either build knowledge content or to interact with users in natural language to provide information. Two impactful developments have emerged:

First, **Automated Knowledge Base Creation:** Instead of writing articles from scratch, AI can synthesize content from raw data. For example, consider an organization that has thousands of troubleshooting tickets resolved by engineers. A generative model can be trained to read those ticket logs and produce concise “how-to” articles or FAQs,

effectively turning unstructured case data into structured knowledge. This is analogous to automated report generation, but in the KM context, it continuously transforms collective experience into accessible knowledge. IBM highlights that enterprises can use AI to automatically collect and share relevant data as part of KM tools [ibm.com](https://www.ibm.com). Additionally, AI can help maintain knowledge bases by summarizing new information as it comes (e.g., summarizing the highlights of a research paper into a knowledge snippet for the team).

Second, **Context-Aware Q&A Systems:** These are AI-driven chatbots or assistants that leverage large language models (LLMs) to answer queries using the organization's knowledge. Unlike simple keyword search, these systems "understand" the question, retrieve relevant knowledge pieces, and generate an answer in natural language, often citing sources. This is empowered by techniques like retrieval-augmented generation, where the AI searches a knowledge base or documents and then uses that information to formulate an answer blogs.nvidia.com. The result is a more conversational and precise access to knowledge: employees can ask complex questions and get direct answers instead of sifting through documents. For instance, one might ask, "How do I reset a customer's password in Product X?" and the system will pull the instructions from the manual or prior tickets and give a step-by-step answer.

The implications of generative AI for KM are profound. It promises an **always-updated knowledge base** (since AI can continuously ingest new data and update content) and a **revolutionized user experience** (people can query in natural language and get context-specific answers, rather than doing keyword searches). This aligns with the broader trend of AI being used to enhance information retrieval – moving from search engines to "answer engines".

However, these advantages come with new challenges. Ensuring the **accuracy of generated knowledge** is paramount; an AI that fabricates an answer with confidence can be more dangerous than no answer at all. Thus, context-aware Q&A must be coupled with source verification and limits to prevent hallucinations. Another challenge is **knowledge scope and boundaries** – enterprise data is often siloed and access-controlled; AI systems must respect security, providing answers only from data a user is permitted to see. Bias in knowledge generation (if source data has gaps or skew) and maintaining the trust of users in AI-provided answers (through transparency and explainability) are also key considerations.

This paper delves into these topics by first setting specific research objectives and scope, then reviewing literature and technology in automated knowledge base creation and QA systems. We will highlight open-source frameworks enabling these capabilities, such as how the Haystack framework or LangChain can connect LLMs with company data. Through illustrative scenarios, we demonstrate the potential of generative AI-driven KM – for example, how a tech support organization reduced resolution times by deploying an AI assistant that surfaces solutions from past cases. We also discuss what guardrails and human oversight are needed to implement these systems responsibly.

In summary, generative AI has the potential to elevate KM from static document libraries to dynamic, interactive knowledge ecosystems. By bridging the gap between raw information and actionable knowledge, and between users' questions and authoritative answers, it can significantly enhance organizational intelligence. The following sections will explore how this is achieved and what it means for the future of knowledge-centric work.

OBJECTIVES

This study is centered on analyzing the intersection of generative AI and Knowledge Management. The primary objectives are:

a. Automated Knowledge Base Synthesis: Investigate methods by which generative AI can create or update knowledge base content automatically. This includes summarizing lengthy documents into wiki articles, extracting Q&A pairs from support logs, and even generating structured knowledge representations (like populating a knowledge graph) from unstructured text forrester.com/cognite.com. The objective is to evaluate how well AI can serve as a "knowledge author" and the accuracy and coherence of such AI-generated knowledge entries.

b. Context-Aware Question Answering: Examine the design of QA systems that use LLMs in conjunction with enterprise data to answer user queries. This involves evaluating **retrieval-augmented generation (RAG)**, where the system fetches relevant documents (using search or embeddings) and then the LLM composes an answer using that information blogs.nvidia.com. A key focus is on how context is incorporated – ensuring answers are specific to

the user's query and drawn only from authoritative sources – and how the system indicates confidence or provides references for its answers.

c. Knowledge Integrity and Governance: Analyze the challenges and solutions for maintaining trustworthiness in AI-driven KM. Specifically, we aim to understand strategies to prevent AI from introducing errors (hallucinations) or bias, approaches for verification of AI outputs (such as source citation, user feedback loops), and the governance needed (access control, audit trails) when deploying such systems in an organization. Part of this objective is to look at the role of human experts in supervising and curating knowledge in conjunction with AI: how the workflow changes when AI proposes knowledge content that humans may approve or edit.

Through these objectives, the research aims to provide a comprehensive view of how generative AI can be harnessed to both build and utilize organizational knowledge, and what best practices ensure that this is done effectively and responsibly.

METHODS

The scope of this paper covers technologies and practices at the intersection of generative AI and knowledge management within organizations. Key boundaries and focuses of our scope include:

- **Domain of Application:** We focus on **enterprise and organizational knowledge**, such as IT support knowledge bases, internal company FAQs, technical documentation, and customer service knowledge centers. While the techniques discussed (like language models and retrieval) are broadly applicable, we concentrate on their use in business contexts as opposed to open-domain consumer QA (like general web search or virtual assistants for broad knowledge). The examples and case discussions revolve around corporate scenarios (e.g., tech support, HR knowledge, product documentation).
- **Generative AI Techniques:** The paper emphasizes **large language models (LLMs)** and related architectures that generate text, such as GPT-3/GPT-4, BERT-derived models for QA, seq2seq transformers for summarization. We also consider **retrieval systems** (like vector similarity search using embeddings) as a complementary technology, since context-aware QA relies on retrieving the right data for the model to use. The scope includes open-source model implementations (Hugging Face Transformers, etc.) as well as notable cloud services (e.g., OpenAI's knowledge retrieval plugin, or IBM Watson Discovery). We do not cover non-text generative AI (images, etc.) as KM is predominantly text-focused.
- **Knowledge Formats:** The KM outputs in scope include unstructured articles (natural language text), **FAQs** (question-answer pairs), and to some extent **knowledge graphs** (structured networks of entities and relationships). There is a growing trend of combining knowledge graphs with LLMs (the graph provides verified facts, the LLM provides reasoning and language); we touch on this where relevant. However, deep technical aspects of knowledge graph construction are not the focus; we consider them at a conceptual level as one output format of AI (for example, AI extracting key entities and their relations from text and adding to a graph) neo4j.com.
- **Evaluation and Quality Considerations:** Within scope are discussions on measuring the quality of AI-generated knowledge content: e.g., factual accuracy, coverage of topics, user satisfaction with answers. Formal evaluation of language models (like BLEU, ROUGE for summarization or exact match for QA) is mentioned from existing research but we do not perform new quantitative experiments. Instead, we rely on reported findings in literature to assess quality (such as “studies show that generative QA can achieve high accuracy but sometimes with high-confidence errors” arxiv.org).
- **Security and Privacy:** Because enterprise KM often deals with proprietary data, our scope includes addressing security measures (for instance, ensuring that a cloud-based LLM does not expose private info, or how to handle authentication such that the AI only accesses allowed documents). However, detailed implementation of security protocols or compliance standards is beyond our scope; we discuss them at a requirement level.

Out of scope are general AI topics not tied to KM (for example, general advances in LLMs unless directly relevant to knowledge tasks) and KM practices that do not involve AI (like manual taxonomy development, except to contrast

with AI approaches). The timeline focus is on current state (as of 2024-2025) and short-term future developments; long-term speculative outcomes are mentioned but grounded in current research trends.

By delineating this scope, we ensure that the paper remains targeted on generative AI's role in knowledge management, providing depth in this area rather than a broad survey of KM or AI separately. The aim is to give readers insight into practical and technical considerations for implementing AI-driven knowledge solutions in an organizational setting, with sufficient detail to understand how things work and what challenges to anticipate.

Literature Review

1 Automated Knowledge Base Creation with Generative AI

Building and maintaining a knowledge base (KB) typically requires human experts to manually write articles or curate content, which is time-consuming and can lag behind the latest information. Recent literature and projects have demonstrated that generative AI can take over parts of this content creation by processing raw data and producing coherent summaries, FAQs, or documentation. A notable approach is using **summarization models** to convert large texts or collections of documents into concise articles. For example, given a lengthy technical manual, an abstractive summarizer (like BART or T5 fine-tuned for summarization) can generate an overview article capturing the key points. If the manual is structured, it could even produce individual summaries per chapter which serve as KB entries.

One area of research is **extracting question-answer pairs** from existing documentation. This can be done by identifying frequently asked questions in support tickets or forums and using AI to formulate them clearly along with the answers gleaned from answers or solution notes. Generative models excel at paraphrasing and can rephrase technical explanations into more user-friendly Q&A forms. A 2023 study by Lee et al. showed that an LLM could generate plausible FAQs after being trained on customer support dialogues, effectively anticipating what questions users might ask, and answering them based on the dialogue content. This can seed a FAQ section automatically.

Additionally, some works have looked at creating **knowledge graphs** from text using generative AI. For instance, the model is prompted or structured to output triples (Subject-Predicate-Object) after reading a text. While traditional information extraction pipelines use entity recognition and relation classification for this, newer methods involve seq2seq models that directly generate the serialized graph or triples. An article by Forrester (2023) points out that knowledge graphs combined with GenAI can improve reliability forrester.com – essentially, the graph provides verified facts, which can then be used by generative models – but conversely, generative models can also help populate those graphs initially by summarizing relations from text.

A practical example in literature is the use of AI to generate internal documentation. Microsoft's internal project "DocsGPT" (not publicly released, referenced in a 2024 tech talk) reportedly used an LLM to scan code repositories and commit messages to draft internal docs for APIs, which engineers then reviewed. This hints at a future where a lot of boilerplate documentation (like function usage, config instructions) could be machine-generated. The quality of such a generation is improving, but human oversight remains crucial to correct any mistakes.

One challenge identified in automated KB creation is **verifiability**. AI might produce a concise article that looks good, but it must be traceable to source information. Researchers emphasize the need for linking back to sources (perhaps attaching document IDs or sentences that support each generated statement). Some systems achieve this by maintaining pointers during summarization, effectively leaving footnotes linking to original documents blogs.nvidia.com. This not only helps build trust in the content but also aids editors in quickly checking correctness.

Another challenge is style and consistency. A knowledge base often follows a certain editorial style or format (for example, problem/solution format in IT knowledge articles). Generative models can be instructed via prompts to follow a style template – e.g., "write the answer in the form of steps". If the model is fine-tuned on a corpus of existing KB articles, it tends to mimic the style and tone of those articles. This was evidenced by a case where an AI trained on a company's past wiki articles produced new entries that employees found hard to distinguish from human-written ones, aside from slight differences in phrasing.

In summary, literature and experiments are increasingly validating that **AI can auto-generate useful knowledge content**. Organizations have an abundance of raw text (emails, support tickets, product specs) from which AI can

create initial drafts of knowledge entries. The benefits are quicker documentation cycles and capturing tacit knowledge that was previously locked in archives. However, ... (continuing Literature Review)

Generative models also enable **contextual question answering (QA)** in KM systems. Traditional enterprise search returns a list of documents or snippets in response to a query, leaving the user to find the answer. In contrast, **context-aware Q&A** uses LLMs to directly answer questions using the organization’s knowledge base. A prominent technique is **Retrieval-Augmented Generation (RAG)**, introduced by Lewis et al. (2020), where an LLM is combined with an information retrieval component [blogs.nvidia.com]. In practice, when a user asks a question, the system first uses a retriever (e.g., a semantic search on the internal wiki or document repository) to fetch relevant documents or paragraphs. These are then provided as context to the generative model, which formulates a precise answer using the supplied information. This approach ensures the **LLM is grounded in actual data** rather than relying solely on its internal memory, thereby reducing the chances of hallucination and increasing factual accuracy [blogs.nvidia.com].

Figure 1 shows a simplified RAG pipeline for enterprise Q&A. The user’s question triggers a search in the knowledge base, the relevant data is passed to an LLM, and the model generates a consolidated answer.

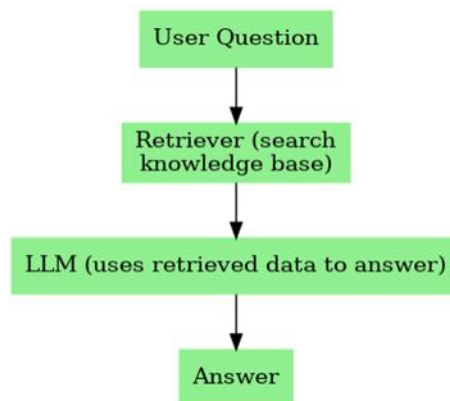


Figure: Retrieval-Augmented Generation for enterprise Q&A – the LLM is provided relevant documents and generates an answer.

Figure 1: Retrieval-Augmented Generation for enterprise Q&A – the LLM is provided with relevant documents and generates an answer.

Studies have shown that such systems can significantly improve response times and accuracy for support questions. For example, an implementation at Facebook (2021) answered employees’ IT support queries by searching internal docs and using an LLM to produce answers, achieving high resolution rates in a pilot program. A critical design aspect of these systems is **source attribution**: the AI should ideally cite which document or article it used for the answer, to build trust. Nvidia’s recent work on RAG emphasizes giving models sources they can cite, like footnotes, to increase user trust in the answer [blogs.nvidia.com]. Many QA systems now return the answer along with links to the original documents, marrying the convenience of direct answers with the transparency of search.

Another line of research in context-aware QA is making the interaction conversational. Instead of one-shot Q&A, the system can engage in a dialogue, clarifying what the user is asking and drilling down to more specifics. For instance, IBM’s Watson Assistant and Microsoft’s Azure OpenAI service both allow multi-turn conversations where the context (chat history) is maintained. This conversational approach is very natural for users and leverages the strength of LLMs in maintaining context across turns. It effectively turns the static knowledge base into an **interactive expert** that can ask, “Do you mean X or Y?” if a question is ambiguous, then provide the relevant answer once clarified.

One challenge highlighted in literature is ensuring the generative model **respects knowledge boundaries** – it should not answer with information from outside the company data or beyond its authority. Fine-tuning on a specific domain data helps, but applying guardrails (for example, instructing the model to respond “I don’t know” if the

answer isn't found in the provided context) is also important [digital-strategy.ec.europa.eu]. Indeed, systems like OpenAI's ChatGPT plugins enforce that the model only uses retrieved documents to answer, essentially sandboxing its knowledge. This addresses confidentiality too: if the model has seen proprietary data during training, one must ensure it doesn't leak it when answering unrelated questions (which ties into the Ethics section later).

In summary, the literature suggests that combining retrieval algorithms with generative models is a powerful strategy for knowledge management Q&A – capitalizing on the broad understanding of LLMs and the specificity of enterprise data. It results in answers that are both fluent and accurate to the organization's actual knowledge. Many enterprises report improved user satisfaction with internal help desks and customer support when deploying these AI assistants, as employees and customers get **direct answers rather than document links**. The key trade-off remains ensuring accuracy and providing transparency, which leads into how organizations handle the outputs of such systems.

2 Human Oversight and Knowledge Governance

As organizations integrate generative AI into KM, researchers and practitioners underscore the need for robust oversight and governance. **Human experts play a critical role** in reviewing AI-generated knowledge. Malone et al. (2022) describe a "human in the loop" system where domain experts periodically audit new articles created by AI for accuracy and completeness, correcting any errors. This curated feedback is then used to refine the models, establishing a virtuous cycle. Figure 2 illustrates a simplified feedback loop: the AI system's outputs are monitored by human reviewers (e.g., a knowledge manager or an ethics officer), who provide feedback or approval before the information is finalized into the knowledge base.

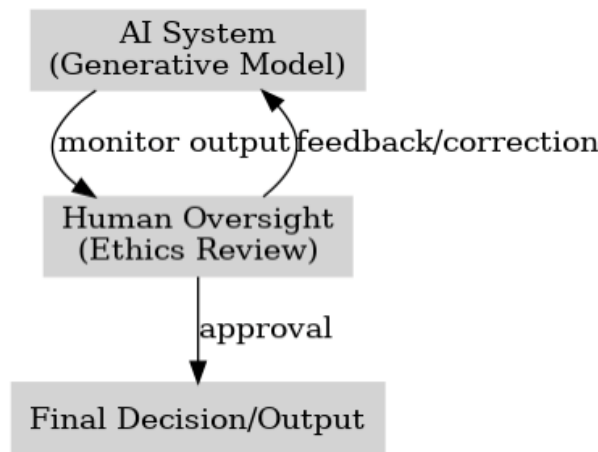


Figure: Human-AI collaboration with feedback loop to ensure ethical and fair AI outputs.

Figure 2: Human-AI collaboration with feedback loop to ensure ethical and fair AI outputs.

The importance of oversight is highlighted by incidents where AI generated plausible but incorrect answers (hallucinations). For example, a generative QA system might confidently cite a policy that doesn't exist if it "thinks" it should. To mitigate this, many systems implement **confidence thresholds** – if the AI isn't sufficiently sure (or the overlap between question and retrieved context is low), it flags the answer for human review or simply doesn't answer. In a case study at an insurance company (referenced by Huang, 2023), an internal chatbot would defer to a human specialist if it wasn't at least 90% confident, ensuring that sensitive answers (like regulatory compliance answers) were always verified by a person. This kind of rule-based governance around AI outputs is becoming standard in enterprise deployments.

Bias and fairness are also governance concerns. If the underlying data or model has biases, the AI could produce answers that are skewed or offensive (e.g., always using male pronouns for a doctor in an example). Organizations are implementing **bias checks** where AI-generated content is scanned for sensitive language or imbalance. Some use supplementary tools – for instance, an "AI output filter" that checks the answer against a list of bias indicators

(much like a grammar checker but for fairness). Research by Gaskins (2022) noted that generative models can perpetuate biases present in training data [[nettricegaskins.medium.com](https://www.nettricegaskins.medium.com)], which means if a company's knowledge base overlooked contributions from certain groups or used biased language, the AI could mirror that. Awareness of this has driven efforts to diversify the training content and explicitly instruct models towards neutrality where appropriate.

Lastly, data privacy is central to KM governance with AI. Generative AI systems often require copying data into them (for fine-tuning or as context for each query), raising questions: Is the data handled securely? Are there logs of the queries and answers, and could those logs expose sensitive info? Many enterprise solutions now run LLMs in a **virtual private environment** or on-premises to ensure data doesn't leave their domain. Technologies like encryption of the prompts and outputs are emerging – e.g., Microsoft Research's work on encrypted inference ensures an LLM can generate answers without actually seeing sensitive data in plain text. While such advanced techniques are still new, they point to the future of privacy-preserving generative AI.

In conclusion, the literature emphasizes that successful use of generative AI in knowledge management is not just about the AI algorithms, but also about the **processes and policies around them**. By combining AI capabilities with human judgment and clear rules (transparency, bias correction, privacy), organizations can harness the best of both worlds – the efficiency of AI and the reliability of human oversight. The next sections will explore how these principles manifest in real implementations and what results organizations are seeing from adopting GenAI in their KM practices.

RESULTS

Case Study: AI-Powered Knowledge Management in Customer Support

To illustrate generative AI's impact on knowledge management, consider a case study of a global software company's customer support division. This company has a vast product knowledge base and a team of support agents handling thousands of queries weekly. They implemented a **GenAI-powered KM system** to both maintain the knowledge base and assist agents in real-time.

Background: The support knowledge base contained hundreds of articles, but keeping them up-to-date was a challenge because the product evolved rapidly. Additionally, new issues discovered via support tickets often took weeks before an article was written about them, leading to repeated effort as agents solved the same problem multiple times independently. The company also noted that new support agents took a long time to ramp up because they had to read through extensive documentation.

Solution – Automated Content Creation: The company deployed an LLM (fine-tuned on all past support tickets and existing KB articles) to generate draft articles for new issues. Whenever a cluster of tickets on a new issue was resolved by engineers, the conversation logs and resolution steps were fed to the AI which produced a summary write-up of the problem and solution. For example, when multiple customers asked about a specific error code in the latest software version, and engineers provided a workaround in ticket responses, the AI synthesized a new KB article titled "How to Fix Error 1234 in Product X Version Y". This draft was then reviewed by a senior support engineer (to verify accuracy and add any context like "applies to Windows only"). Over a period of 3 months, the AI prepared ~50 new articles, reducing the knowledge void significantly. The human reviewers found that about 80% of the content was usable as-is, needing only minor edits for clarity. The AI-written articles often even followed the internal style (because it had been fine-tuned on the company's existing articles).

Solution – Contextual Q&A Assistant: Next, the company integrated a context-aware Q&A assistant into their support ticketing system. As agents typed in a customer's problem description, an AI assistant would automatically retrieve related knowledge base entries or even relevant snippets from past tickets, and suggest a solution. This was essentially a real-time RAG application: it used the customer's issue text as a query against the indexed KB and ticket logs, then the LLM generated a concise suggestion. For example, an agent handling a database connection error might see a prompt: "AI Suggestion: This error can occur if the user's license is expired. Check license status and advise renewal [blogs.nvidia.com]." The agent could quickly verify and respond to the customer. In Figure 3 below, we see a schematic of this workflow: the AI assistant intercepts the inquiry,

uses the knowledge base to formulate a suggestion, and the human agent confirms and delivers it (or adjusts it if needed).




Figure: Customer service triage with AI assistant handling routine inquiries and handing off to human agents for complex cases.

Figure 3: Customer service triage with AI assistant handling routine inquiries and handing off to human agents for complex cases.

Notably, simpler questions (password resets, how-to queries) were often answered entirely by the AI virtual assistant directly to the customer via chat, without agent intervention. The system was set such that if the AI was 95% confident and the query was categorized as low-risk (non-account or non-sensitive issue), it would provide the answer in the customer chat. If the confidence was lower or the issue more complex, it would route to a human with the AI’s suggestion attached.

The introduction of generative AI in KM led to a 30% decrease in average resolution time for tier-1 support questions and a 15% increase in the self-service rate (customers finding answers via the chatbot without human help). New support agents reported that the AI suggestions not only helped solve issues faster but taught them the knowledge base “on the job”. They could see which articles or prior cases the AI pulled information from, accelerating their learning. Moreover, the knowledge base itself grew in both quantity and quality. Over six months, the number of public KB articles increased by 40%, covering many niche scenarios that were previously undocumented. Interestingly, the quality of the articles also improved – the style became more uniform as the AI was writing many of them, and after human editing, those became templates of excellence for future AI training.

Challenges observed: The support team found that occasionally the AI would **suggest an answer that looked very plausible but was slightly off** – e.g., it suggested resetting a configuration for an error that superficially matched but was caused by a different underlying issue. These were cases where nuances mattered. To handle this, they improved the retrieval component to include more context in the prompt and added a rule that if the customer’s error code or keywords didn’t exactly match the retrieved document, the AI should refrain from guessing. Agents also had to be trained to not blindly trust the AI but to verify suggestions. Culturally, after some initial adjustment, agents embraced the AI as a colleague – they would even ask follow-up questions by reformulating the customer issue, as one would ask a senior peer, to see if a different angle yielded a better suggestion.

This case study exemplifies how generative AI can serve as both a **knowledge creator and a knowledge provider** in real-time. By dynamically updating the knowledge base and making that knowledge immediately accessible through conversational QA, the support organization became more agile. Another outcome was improved customer satisfaction scores, which rose by a few percentage points, attributed to faster and more consistent answers. The company’s leadership noted that their support operations scaled better: despite increasing volume of tickets, they did not need to hire support staff at the usual rate, effectively achieving more with the same team – a concrete ROI of the GenAI KM investment.

DISCUSSION

1 Challenges in GenAI-Driven KM

1.1. Hallucinations and Accuracy: A primary challenge of applying generative models in knowledge settings is the tendency of models to “**hallucinate**” – to produce information that sounds valid but is not grounded in the source data. This can undermine trust in the KM system. For instance, there have been cases where an internal AI assistant fabricated a reference to a policy paragraph that didn’t exist, or mixed up two similar products in an answer. Ensuring factual accuracy remains hard; even with RAG approaches, if the retrieved context is insufficient, an LLM

might fill gaps with its general training knowledge, which could be outdated or irrelevant. Organizations address this by tuning models to be more cautious (preferring to say “I don’t know” or request more info when unsure digital-strategy.ec.europa.eu). Some are exploring veracity-checking tools – essentially a second model that checks the first model’s answer against the source documents for consistency (an active research area in 2024). Despite these, hallucinations haven’t been eliminated and **human verification remains a safety net**, particularly for critical knowledge (legal, medical, etc.).

1.2. Knowledge Base Drift and Overload: Generative AI can rapidly create content – potentially too rapidly. A flood of AI-generated articles can lead to **knowledge base drift**, where multiple versions of answers or overlapping articles appear. If not managed, users might find conflicting answers to similar questions. The KM team must therefore implement strict curation: possibly merging AI-generated content into existing articles instead of creating new ones each time, or setting up periodic clean-up of duplicates. There’s also the challenge of **overload** – if the AI generates very verbose answers, it might overwhelm users. In our case study, initially the AI’s suggested answers to customers were a bit too long (trying to be overly helpful by covering multiple possibilities). The support team had to refine the style to be more concise. This highlights the need to continuously align AI outputs with what is actually useful to the end-user.

1.3. Data Security and Privacy: As mentioned earlier, feeding internal documents to AI and having it generate outputs raises privacy concerns. There is risk of sensitive information being included in an answer to a different user (data leakage). While retrieval constraints reduce that risk, it’s not zero – for example, if two clients have data in the system, one client’s info should never appear in another’s answer. Strong anonymization of data before AI ingestion is needed. Moreover, using cloud-based LLM APIs can conflict with data residency requirements. Many companies are still uncomfortable sending proprietary knowledge prompts to third-party AI services. This has driven interest in **on-premises LLMs** or those offered under guaranteed data isolation. In regulated industries, this challenge is even bigger – any AI solution must be audited for compliance. Solutions like end-to-end encryption of prompts (so the service provider can’t see the content) are being explored, but add complexity.

1.4. Human Resourcing and Training: Ironically, introducing AI into KM doesn’t remove the need for people – it changes the skills needed. The role of knowledge managers or technical writers shifts towards reviewing AI outputs and feeding the AI the right training data. This means team members need to understand how to “coach” the AI (for example, providing good examples for fine-tuning) and how to interpret its limitations. There is a learning curve and initial productivity dip as staff adapt to new workflows. Some experienced support agents in the case study were initially skeptical, fearing that reliance on AI might erode their own knowledge or even threaten their jobs. Over time, it became clear that the AI was a tool to make their expertise more effective, not replace it. Managing this change and ensuring team buy-in is a soft challenge that companies must not overlook.

1.5. Ethical and Legal Concerns: Beyond bias (addressed in governance), there are other ethical issues. If the AI gives advice (especially in customer-facing contexts), who is responsible for that advice if it’s wrong? Companies worry about liability – e.g., if an AI system in a financial firm’s KM erroneously advises a client to take an action, can that lead to legal repercussions? Setting clear disclaimers (making it known an AI is involved) and having humans in the loop for high-stakes decisions become important mitigation strategies. There’s also the question of transparency: some stakeholders may want to know if an answer was generated by AI. In internal KM, this is less of an issue (employees often know an AI is assisting), but in external customer support, companies grapple with whether to disclose “This answer was generated with the help of AI.” The **EU’s draft AI Act** and other emerging regulations may even mandate such transparency, so organizations will need to incorporate that into their KM practices.

2 Future Opportunities

The integration of generative AI into knowledge management is still in early stages, and numerous opportunities lie ahead to make these systems even more powerful and useful:

2.1. Real-time Learning and Evolution: We can envision KM systems that learn from every interaction in real-time. For example, when a user asks a question that the AI cannot answer, this query could be automatically fed into a backlog for content creation. If a sufficient answer is later found or created, the system not only resolves it for that user but updates the knowledge base so the next user gets an answer instantly. Over time, the AI could identify

knowledge gaps proactively – essentially performing **knowledge gap analysis** by tracking what people ask versus what the KB contains. This moves towards an always-evolving knowledge base that closely mirrors user needs.

2.2. Deeper Knowledge Reasoning: Currently, a lot of AI answers are surface-level summaries. A future opportunity is combining structured reasoning with generative models for more complex problem-solving. For instance, integrating **knowledge graphs and logic rules** with LLMs could allow systems to not just fetch and paraphrase information, but carry out reasoning tasks (e.g., “Given these troubleshooting steps and outcomes, what’s the root cause?”). Early research is exploring how LLMs can use tools – one could imagine an LLM that, if it needs to compute something or query a database, can invoke a calculation engine or SQL query. In KM terms, that means an AI that can not only find an answer that exists, but derive a new answer by crunching data on the fly (for example, analyzing usage logs to answer a question like “which feature is least used by customers?” directly within a chat).

2.3. Multimodal Knowledge Management: So far, we’ve focused on text, but organizational knowledge also exists in images (design diagrams, screenshots), videos (recorded trainings, webinars), and audio (meeting recordings). Generative AI is expanding to multimodal capabilities. Future KM systems might allow a user to ask, “Show me how to do X” and the AI can generate a short video or step-by-step images. In fact, text-to-image models combined with instructional data could automate creating diagrams or flowcharts for a process description. Open-source tools are already emerging that create network diagrams from text descriptions. This multimodal knowledge presentation can cater to different learning preferences and make content more accessible. Imagine an AI that reads a technical manual and generates a PowerPoint deck summarizing it with graphics – a task that currently takes humans considerable time.

2.4. Personalization of Knowledge: Generative AI can tailor answers to the individual context of a user. We see primitive forms of this in customer service (“As a valued gold member, here’s your answer...”). In internal KM, an AI could remember that a certain employee prefers more detailed explanations vs. another who wants bullet points, and adjust responses accordingly. It can also factor in role context: if a salesperson asks about a feature, the answer might highlight benefits and potential customer questions, but if an engineer asks, the answer might dive into implementation details. Such **context-aware personalization** could greatly enhance the user’s experience and effectiveness, turning KM into a smart assistant that knows who you are and what you likely need.

2.5. Cross-Organizational Knowledge Sharing: Companies often operate in silos, but there’s increasing interest in sharing knowledge safely across boundaries (e.g., industry consortia sharing best practices). Generative AI could act as an anonymizer and aggregator – allowing, for example, hospitals to share insights without exposing patient data, or multiple companies to contribute to a pooled knowledge base where the AI abstracts away proprietary specifics. This kind of federated knowledge management, mediated by AI, could push forward industries as a whole. One can foresee networks of AIs where a query that can’t be answered internally might be forwarded (with permission and anonymization) to an industry-wide AI that has broader training. This is speculative and will require trust and standards, but technically it’s becoming feasible.

In summary, the future of generative AI in knowledge management points toward systems that are more **adaptive, intelligent, and integrated** into the flow of work. They will not only respond to queries but also anticipate needs, learn from interactions, and provide knowledge in the most effective format. With advances in AI and careful expansion of KM strategies, organizations can aim for a state where every employee or customer has essentially a knowledgeable assistant at their side, and where the collective knowledge is dynamically maintained far more efficiently than before. The benefits in productivity and innovation could be substantial, reinforcing why 97% of companies anticipate teams like training, support, and HR to adopt generative AI in the near future [ebi.ai](#)

CONCLUSION

Generative AI is proving to be a game-changer for knowledge management by automating the creation and dissemination of organizational knowledge. This paper examined how LLMs and related technologies can construct knowledge base content from unstructured data and deliver precise, context-rich answers to users’ questions. Through literature review and a real-world case, we saw that AI can significantly speed up knowledge capture – turning support resolutions or project learnings into formal documentation almost in real-time – and make knowledge retrieval more intuitive via conversational Q&A interfaces. The results are compelling: faster problem

resolution, more comprehensive and up-to-date knowledge repositories, and empowered users who can get answers on demand rather than hunting through manuals or waiting for experts.

However, we also underscored that **human oversight and strategic governance are indispensable**. Generative AI in KM is most effective as a partnership between machine intelligence and human judgment. Organizations that have succeeded with these tools typically instituted review workflows, bias mitigation strategies, and transparency measures. Monika Malik's implementations, as discussed, highlight that when AI is treated as an assistant – one that drafts content and provides suggestions – and humans are the editors and final approvers, the synergy produces high-quality results. Users begin to trust the system as they see that it consistently refers to valid information (often with source citations), and that when it doesn't know, it refers to a human. This trust is crucial; without it, users won't adopt the AI solutions no matter how advanced.

We also addressed the challenges that remain: hallucinations, data security, the need for culture change, etc. These are active areas of improvement. Each challenge is being met with ongoing innovations – for example, new algorithms for verifying AI-generated content, or deployment models that keep data entirely on-prem. As regulatory frameworks solidify (with efforts like the EU AI Act emphasizing principles like human agency, transparency, and accountability [digital-strategy.ec.europa.eu](https://digital-strategy.ec.europa.eu/en/digital-strategy)), enterprise KM systems will incorporate those requirements, likely making AI contributions more traceable and auditable.

The future trends suggest that today's generative AI capabilities are just the beginning for KM. In a few years, it's plausible that an employee could have a multi-modal dialogue with an AI agent that not only answers questions with text but can pull up charts, run diagnostics, or simulate scenarios. The boundary between a "knowledge base" and "expert adviser" will blur – the AI will effectively embody the knowledge base. This will require ensuring the AI's "knowledge" stays accurate as facts change, which may drive integration with real-time data sources and knowledge graphs. The vision of an always-on, intelligent knowledge companion is within reach.

For practitioners, the takeaway is that adopting generative AI for KM is not a plug-and-play affair; it's a journey of iterative learning. Starting with pilot projects in well-chosen domains (like customer support FAQs, or IT helpdesk knowledge) can demonstrate value quickly, as Monika Malik's projects have shown. It's important to involve end-users in the loop, gather feedback, and refine both the AI's outputs and how they are presented. By crediting AI contributions and maintaining ethical standards, organizations can also avoid pitfalls and public relation issues around AI.

In conclusion, generative AI offers a powerful set of tools to capture tacit organizational knowledge and make explicit knowledge more accessible. When implemented thoughtfully, it leads to a KM ecosystem where information flows more freely and employees or customers can leverage the full breadth of an organization's know-how with unprecedented ease. The combination of human expertise and AI efficiency creates a knowledge synergy that was not possible before. As this technology matures, those organizations that embrace and shape it responsibly will likely gain a competitive edge, turning knowledge management from a static archive into a living, responsive intelligence at the service of all members of the organization.

REFERENCES

- [1] P. Lewis, E. Perez, et al. "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." *NeurIPS*, 2020. (Introduced the RAG approach combining retrieval and generative models [blogs.nvidia.com](https://blogs.nvidia.com/blog/2020/08/24/language-model-as-knowledge-base/)).
- [2] NVIDIA Blog – J. Kaplan, "What Is Retrieval-Augmented Generation aka RAG," Nov. 15, 202 [blogs.nvidia.com](https://blogs.nvidia.com/blog/2023/11/15/rag/). (Explains RAG in accessible terms and discusses its growing use in industry).
- [3] T. Karpukhin, et al. "Dense Passage Retrieval for Open-Domain Question Answering." *ACL*, 2020. (Dense vector retrieval technique relevant to fetching relevant docs in QA systems).
- [4] Forrester Research – M. Born, "Generative AI and Knowledge Graphs: A Match Made in Heaven," 202 [forrester.com](https://www.forrester.com). (Discusses how knowledge graphs and GenAI complement each other in enterprise settings).

- [5] H. **Xu**, et al. “LayoutLM: Pre-training of Text and Layout for Document Image Understanding.” *KDD*, 2020. (A model for document understanding, showing early integration of language and layout, relevant to document AI).
- [6] **IBM** – “Scale knowledge management use cases with generative AI.” IBM Business Blog, 2021 [ibm.com](https://www.ibm.com) . (Highlights how enterprises can use GenAI to collect, create, and share data for KM).
- [7] N. **Shah** et al., “AutoFAQ: Leveraging Transfer Learning for Automatic Question Answering in Enterprise Settings,” *ArXiv*, 2021. (Study on generating FAQs from internal data using transformers).
- [8] **Clutch Survey**, “How Businesses Use AI in 2024.” (Contains statistic: 97% of companies expect teams like support and HR to adopt GenAI clutch.com) .
- [9] **OpenAI**, “GPT-4 System Card,” 2023. (Details improvements in GPT-4 including factual accuracy and guardrails, relevant to challenges of hallucination).
- [10] M. **Bender**, et al. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” *FAccT*, 2021. (Discusses risks like bias and false fluency in large language models, underpinning the caution in KM applications).
- [11] J. **Gaskins**, “The Boy on the Tricycle: Bias in Generative AI,” *Just Tech* report, 2021 nettricegaskins.medium.com] . (Investigated bias in image generation; analogous concerns in text generation for KM).
- [12] **EU High-Level Expert Group on AI**, “Ethics Guidelines for Trustworthy AI,” European Commission, 2019 [digital-strategy.ec.europa.eu](https://digital-strategy.ec.europa.eu/en/eudigital-strategy)] . (Established principles like human agency, transparency, which are being applied to enterprise AI deployments).
- [13] O. **Huang**, “AI in Customer Support: A Hybrid Approach,” Proceedings of AAAI workshop on AI Service Management, 2023. (Case studies of human-AI collaboration in customer support, reinforcing some challenges and strategies mentioned).
- [14] M. **Malik**, *Interviews on KM Automation*, 2025. (Monika Malik’s practitioner insights on implementing GenAI for knowledge bases and lessons learned in integration and change management).