INSTITUTO
UNIVERSITÁRIO
DE LISBOA

AI-Powered Chatbots for Technical Knowledge Management

Paulo Alexandre Abrantes Moreira

Master in Digital Technologies for Business

Supervisor:
PhD. Rúben Filipe de Sousa Pereira, Assistant Professor,
ISCTE-IUL

Co-Supervisor:
PhD. João Pedro Beato Antão, Invited Assistant Professor,
ISCTE-IUL

December 2025

TECNOLOGIAS
E ARQUITETURA

Departamento de Ciências e Tecnologias da Informação

AI-Powered Chatbots for Technical Knowledge Management

Paulo Alexandre Abrantes Moreira

Master in Digital Technologies for Business

Supervisor:
PhD. Rúben Filipe de Sousa Pereira, Assistant Professor,
ISCTE-IUL

Co-Supervisor:
PhD. João Pedro Beato Antão, Invited Assistant Professor,
ISCTE-IUL

December 2025

To Rosa, for encouraging me to start this Master's degree
and supporting me throughout the entire journey.

Acknowledgements

I would like to express my sincere appreciation to my supervisors, Professor Rúben Pereira and Professor João Pedro Antão, for their availability, guidance, support and constructive feedback throughout the development of this work.

I would also like to thank all the teachers of the Master's programme in Digital Technologies for Business for the knowledge shared during the course, as well as my colleagues and project group members for the collaboration, positive discussions, and mutual support during this academic journey.

My gratitude is also extended to Tekuchi and its team for their collaboration, particularly to Ana Lino and Nuno Victor for their essential support in organising and preparing data that was critical to this research, as well as for their participation in the testing and evaluation processes. I would also like to thank Rui Madeira, Rafael Pereira, Lara Duarte, and Mark Harris for their availability to participate in interviews and evaluation activities, which were fundamental to grounding this research in a real organisational context.

On a personal note, I am deeply grateful to my partner Rosa for her constant encouragement, motivation, and patience throughout this process, and to my family for their continued support.

Finally, I would like to acknowledge all those who have directly or indirectly been a positive influence on my personal and professional path. I would also like to express my appreciation to the individuals and communities who have contributed to the development of Artificial Intelligence in a responsible and constructive manner, particularly within open-source ecosystems, whose collaborative efforts continue to advance innovation and shared knowledge.

Resumo

No atual contexto de transformação digital, a Gestão do Conhecimento é um ativo estratégico crítico para a eficácia e competitividade organizacional. Muitas organizações continuam a enfrentar desafios persistentes relacionados com silos de informação, transferência ineficiente de conhecimento e os custos operacionais associados a questões técnicas repetitivas. Esta tese aborda estes desafios no contexto de uma organização tecnológica internacional, explorando de que forma a Inteligência Artificial pode ajudar as organizações a ultrapassar a documentação estática e a responder de forma mais eficaz às suas necessidades em constante evolução.

A investigação adota a metodologia de Design Science Research para conceber, desenvolver e avaliar o TekBot, um agente conversacional baseado em Inteligência Artificial, integrado nas ferramentas de comunicação diária da organização. A solução proposta visa melhorar o acesso ao conhecimento técnico, reduzir a dependência de indivíduos-chave e permitir a captação eficiente e intuitiva do conhecimento organizacional. Ao longo do projeto, o sistema evoluiu de uma abordagem inicial baseada na recuperação de informação através de embeddings não estruturados para uma solução mais robusta, capaz de processar consultas através da integração de dados estruturados e raciocínio baseado em SQL.

Os resultados demonstram que a integração de Inteligência Artificial Generativa nas plataformas de comunicação corporativa melhora significativamente a eficiência operacional, ao reduzir os tempos de pesquisa, minimizar as interrupções aos especialistas na matéria e facilitar o armazenamento da informação. Este trabalho evidência o valor prático e científico dos agentes conversacionais baseados em IA como facilitadores de uma gestão do conhecimento eficaz em ambientes organizacionais digitalmente transformados.

Palavras-chaves: Inteligência Artificial, Chatbots, Gestão de conhecimento

Abstract

In the current context of digital transformation, Knowledge Management is a critical strategic asset for organisational effectiveness and competitiveness. Many organisations continue to face persistent challenges related to information silos, inefficient knowledge transfer, and the operational costs associated with repetitive technical queries. This thesis addresses these challenges within an international technology organisation by exploring how Artificial Intelligence can help organisations move beyond static documentation to better address their evolving needs.

The research adopts a Design Science Research methodology to design, develop, and evaluate TekBot, an AI-powered conversational agent integrated into the organisation's daily communication tools. The proposed solution aims to improve access to technical knowledge, reduce dependence on key individuals, and enable the efficient and user-friendly capture of organisational knowledge. Throughout the project, the system evolved from an initial retrieval-based approach using unstructured embeddings into a more robust assistant capable of handling queries through structured data integration and SQL-based reasoning.

The results demonstrate that integrating Generative AI into corporate communication platforms significantly enhances operational efficiency by reducing search times, minimising interruptions to subject-matter experts and facilitates information storage. This work highlights the practical and scientific value of AI-driven conversational agents as enablers of effective knowledge management in digitally transformed organisational environments.

Keywords: Artificial Intelligence, Chatbots, Knowledge Management

Table of Contents

Acknowledgements ..... *iii*

Resumo ..... *v*

Abstract ..... *vii*

List of Tables ..... *xi*

List of Figures ..... *xiii*

List of Acronyms ..... *xv*

Introduction *1*

Literature Review *3*

2.1. Research Methodology **3**

2.2. Results **4**

Research Methodology *11*

3.1. Problem identification **12**

3.2. Objectives of the solution **13**

3.3. Evaluation process **13**

TekBot Prototype *16*

4.1. First DSR Iteration **16**

4.1.1. Design and Development 16

4.1.2. Demonstration and Evaluation 23

4.2. Second DSR Iteration **24**

4.2.1 Design and Development 25

4.2.2. Demonstration and Evaluation 27

4.3. Third DSR Iteration **29**

4.3.1. Design and Development 29

4.3.2. Demonstration and Evaluation 31

4.4. Fourth DSR Iteration **32**

4.4.1. Design and Development 33

4.4.2. Demonstration and evaluation 35

4.5. Fifth DSR Iteration **36**

4.5.1. Design and Development 37

4.5.2. Demonstration and evaluation 39

Conclusions *41*

5.1. Contributions **41**

5.2. Limitations **42**

5.3. Future work **42**

References *44*

List of Tables

Table 1 – Scientific publications selected for the literature review 5
Table 2 – Literature summary 6
Table 3 – DSRM mapping 12
Table 4 – DSR iterations 14
Table 5 – Interview questionnaire 15
Table 6 – First iteration: main components and technology 21
Table 7 – First iteration: interview results 23
Table 8 – Second iteration: improvement suggestions 25
Table 9 – AI Model comparison 27
Table 10 – Second iteration: interview results 28
Table 11 – Third iteration : improvement suggestions 29
Table 12 – Third iteration: interview results 32
Table 13 – Fourth iteration: improvement suggestions 32
Table 14 – Performance comparison between Notion-based and SQL-based retrieval 34
Table 15 – Fourth iteration: interview results 35
Table 16 – Fifth iteration: improvement suggestions 36
Table 17 – Message Classifier analysis 38
Table 18 – Fifth iteration: main components and technology updated 38
Table 19 – Fifth iteration: interview results 39

List of Figures

Figure 1 – Flow diagram for systematic review (PRISMA) 4
Figure 2 – DSRM diagram 11
Figure 3 – n8n workflow 16
Figure 4 – Message classification 18
Figure 5 – Data storing 19
Figure 6 – Question answering 20
Figure 7 – Data embedding 20
Figure 8 – Example of question 22
Figure 9 – Pinecone VectorStore integration for data embedding 26
Figure 10 – Pinecone VectorStore integration for AI agent tool 26
Figure 11 – "Other" AI agent subsystem 30
Figure 12 – Example of conversational flow between user and TekBot 31
Figure 13 – AI Agent with SQL-based retrieval 33
Figure 14 – Example of list generation and comparison 34
Figure 15 – Save request interaction 37

List of Acronyms

AI	Artificial Intelligence
API	Application Programming Interface
DB	Database
DSRM	Design Science Research Methodology
FAISS	Facebook AI Similarity Search
GenAI	Generative Artificial Intelligence
IT	Information Technology
KM	Knowledge Management
LLM	Large Language Model
NLP	Natural Language Processing
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses
Q&A	Questions and Answers
RAG	Retrieval-Augmented Generation
SQL	Structured Query Language
UI	User Interface
UX	User Experience

CHAPTER 1

Introduction

Knowledge management (KM) plays an important role in driving innovation and performance in teams [1], [2]. The sharing of knowledge, the capture of tacit and hard-to-transfer expertise, and the continuous updating of information repositories contribute to making teams more effective and adaptable. KM plays a key role in bridging the gap between remote work and strong team performance. Such practices support more efficient information sharing, foster innovation, and improve overall team performance [2].

Simultaneously, organisations face numerous challenges in maintaining accurate and current technical documentation. These include cultural resistance to documentation practices [3], a loss of focus caused by fast-paced development cycles, and a tendency to measure productivity solely by deliverables in the form of software or code [4]. In addition, informal and verbal exchanges are frequently left undocumented in the appropriate platforms, and the complexity of managing multiple tools and processes further complicates efforts to maintain proper documentation [4].

In recent years, particularly following the impact of the COVID-19 pandemic, there has been a significant shift in work dynamics, prompting many organisations to adopt distributed models. In these models, employees collaborate across different geographical locations, time zones, and cultural contexts [5], [6]. This transformation has introduced new operational and communicational challenges such as asynchronous coordination, the need to access new information from different parts of the organisation [7], reduced face-to-face interactions and language barriers [8]. As a result, tools that ensure rapid, clear and efficient access to information, without relying on the simultaneous presence of team members, have become even more necessary [5].

Technology can play a key role in accelerating and enhancing knowledge transfer [1]. In particular, recent advances in artificial intelligence (AI) and conversational agents (chatbots) [9] have opened new possibilities for improving access to knowledge [10]. The evolution of chatbots has been marked by a significant transformation, moving from rule-based systems to advanced conversational agents powered by AI [11]. These developments have enabled chatbots to become strategic tools across industries, supporting customer engagement, automating operations, and reducing operational costs [12]. Today, AI-powered chatbots are actively contributing to digital transformation in sectors such as banking, e-commerce, tourism, and customer service [8], [9].

The use of AI-powered chatbots as virtual assistants represents a promising approach to improving technical KM within organisations [10]. These systems offer several advantages, such as 24/7 availability, real-time responses without the need for human intervention, reduced search time [13], and the ability to interact using natural language across multiple languages [14].

This study aims to develop an AI-powered chatbot prototype to improve access to internal technical knowledge and support the inclusion of new documentation within an organisation. To achieve this, the Design Science Research Methodology (DSRM)[15] is followed, covering all phases from design to implementation and evaluation, ensuring a structured and rigorous research approach.

This document is structured into five chapters. The first chapter presents the introduction, outlining the general context of the project and detailing its objectives. Chapter 2 is dedicated to the literature review, which examines the current state of the art in the field and presents key theories, approaches, and findings based on previous relevant studies. Chapter 3 presents the research methodology used. Chapter 4, titled “TekBot Prototype”, describes the multiple iterations of design, development and evaluation of the functional prototype. Finally, Chapter 5 offers a summary of the project’s key outcomes, limitations and reflects on potential future developments, identifying research opportunities that could be explored in this specific area.

CHAPTER 2

Literature Review

2.1. Research Methodology

This chapter was developed following the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) guidelines, to document how studies were identified, selected, assessed, and included, thereby ensuring that the review is transparent, comprehensive, and reproducible.

The scientific databases used in this research were Scopus [16], IEEE Xplore [17] and Web of Science [18] as they offer extensive coverage of both scientific and technical literature. The search terms were defined based on three dimensions: technology, purpose and scope. Using these dimensions as a foundation, the following search expression was constructed and subsequently applied across the databases:

("Chatbot" OR "conversational agent" OR "RAG" OR "Retrieval Augmented-Generation") AND
("knowledge*" OR "documentation" OR "information") AND
("technical team" OR "technical support" OR "technical assistance" OR "technical work")

Given the rapid technological developments in this field, and to ensure the relevance and currency of the literature, a temporal filter was applied, restricting the search to the last five years, specifically from 2020 to May 2025. To facilitate readability, only documents written in English or Portuguese were included.

The initial selection of results was carried out through the reading and analysis of titles and abstracts to ensure their relevance to the topic. Following this step, the availability of each document for full-text access was verified. Subsequently, a thorough reading of the selected documents was performed.

During the reading phase, several support tools were used, including:

- Zotero [19], for saving and referencing the collected records.
- Microsoft Excel [20], for organising the information and performing statistical analysis.
- ChatGPT [21] for translation and summarisation.
- SciSpace [22] for analysis and summarisation.

All AI translations and summaries were reviewed and validated against the original sources.

2.2. Results

During the selection process, and following the initial screening, the records to be included in this literature review were identified. Figure 1 illustrates and describes the article selection process following the PRISMA guidelines.

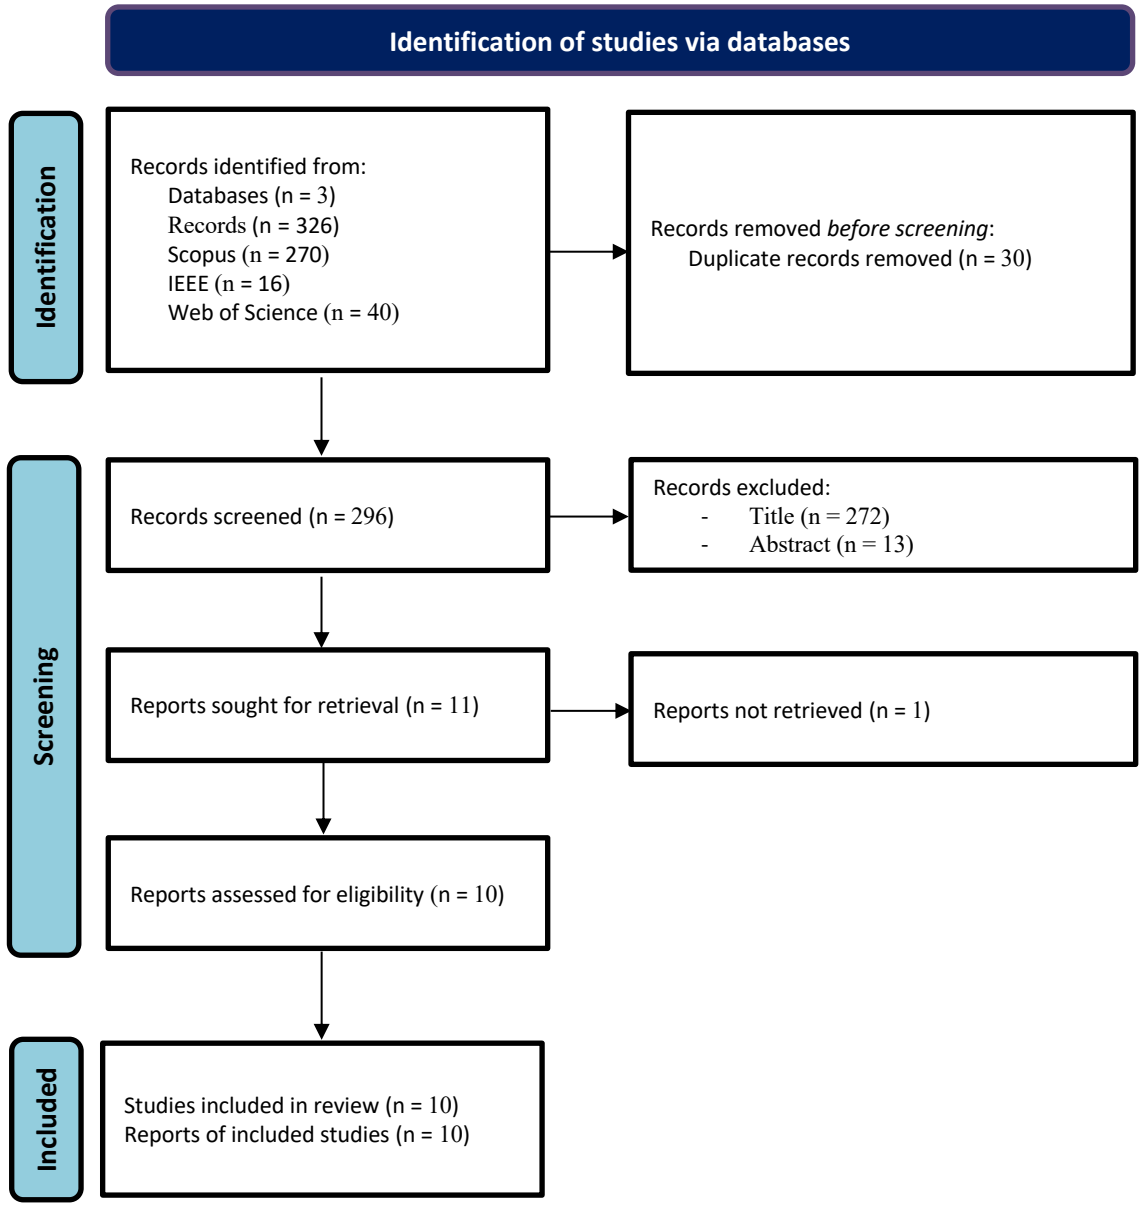


Figure 1 – Flow diagram for systematic review (PRISMA)

As a result of this process, the documents presented in Table 1 were identified.

Table 1 – Scientific publications selected for the literature review

TITLE	AUTHORS	DATE
Developing and Optimizing a Model for Creating a Learning Chatbot Using the Example of Technical Support [13]	Tishchenko, A.A.; Tishchenko, P.A.; Shestakov, T.A.	2021
(Re)Designing IT Support: How Embedded and Conversational AI Can Augment Technical Support Work [23]	Poser, M.; Bittner, E.A.C.	2021
Chatbot Based on Conversational Artificial Intelligence for the Attention of an IT Service Desk in Florida [24]	Torres, J.O.; Barraza, V.X.; Castaneda, P.S.	2024
Development of an RAG-Based LLM Chatbot for Enhancing Technical Support Service [9]	Lee, H.-C.; Hung, K.; Man, G.M.-T.; Ho, R.; Leung, M.	2024
Development of Helpdesk Chatbot for Incident Classification and Resolution using NPL and Deep Learning [25]	Tipe-Palomino, D.; Auccahuasi, W.	2024
Generative AI in Customer Support Services: A Framework for Augmenting the Routines of Frontline Service Employees [26]	Reinhard, P.; Li, M.M.; Peters, C.; Leimeister, J.M.	2024
LLM Support for Real-Time Technical Assistance [27]	Scotti, V.; Carman, M.J.	2024
Multi-Language Conversational Agent for Tech Support: Design and Implementation [14]	J. Benita; R. M. R. Kumar; E. E; G. M. Reddy; N. R. K	2024
Retrieval-Augmented Generation with Knowledge Graphs for Customer Service Question Answering [28]	Xu, ZT; Cruz, MJ; Guevara, M; Wang, T; Deshpande, M; Wang, XF; Li, Z;	2024
TechAssist: A RAG-Based Chatbot for Accessing Technical Information from StackOverflow [29]	Vasaniya, R.; Visodiya, M.; Patel, A.K.	2025

The initial selection targeted studies published from 2020 to 2025 and shows a significant concentration of publications in 2024. This clustering reflects the recent advances in generative AI (GenAI) technologies and the broader momentum of innovation and research interest currently shaping the field.

To support the understanding of each document's main contribution, Table 2 was created to summarise their primary objectives and the technological approaches adopted.

Table 2 – Literature summary

TITLE	OBJECTIVE	C.I. (a)	RAG	AN (b)	AD (c)	TECHNOLOGY	MODEL
Developing and Optimizing a Model for Creating a Learning Chatbot Using the Example of Technical Support, 2021 [13]	Chatbot aimed at automating support tasks to alleviate the workload on existing technical support systems	YES	NO	YES	NO	Python Seq2seq	N/A
(Re)Designing IT Support: How Embedded and Conversational AI Can Augment Technical Support Work, 2021 [23]	Proposes a redesign of the IT support process integrating AI conversational agents	NO	-	N/A	N/A	N/A	N/A
Chatbot Based on Conversational Artificial Intelligence for the Attention of an IT Service Desk in Florida, 2024 [24]	Implementation of a chatbot using Dialogflow for IT service management	YES	YES	YES	NO	Dialogflow Conversational AI Platform Integration with Zoho Desk for tickets Integraton with WhatsApp	N/A
Development of an RAG-Based LLM Chatbot for Enhancing Technical Support Service, 2024 [9]	Developed a domain-specific and RAG-based LLM chatbot to enhance the software house technical support of an IT consultant in Canada	YES	YES	YES	NO	Flowise AI Pinecone vector database (Indexing Stage) Ollama Google Cloud, Bigquery and Google Cloud	Llama2 7B
Development of Helpdesk Chatbot for Incident Classification and Resolution using NPL and Deep Learning, 2024 [25]	Developed a Chatbot to assist with help desk processes using NLP and Deep Learning techniques	YES	NO	YES	NO	PYTORCH POSTGRESQL/Amazon S3 - Database	N/A
Generative AI in Customer Support Services: A Framework for Augmenting the Routines of Frontline Service Employees, 2024 [26]	Conceptualisation of AI augmentation in customer support services, specifically focusing on how generative AI (GenAI) can enhance the routines of frontline service employees	NO	-	YES	NO	N/A	N/A
LLM Support for Real-Time Technical Assistance, 2024 [27]	Demo web application utilizing Large Language Models (LLMs) to enhance user support	YES	YES	YES	NO	PyTerrier for lexical search Sentence Transformer for semantic search	N/A
Multi-Language Conversational Agent for Tech Support: Design and Implementation, 2024 [14]	Multilingual chatbot that helps users who speak different languages communicate more effectively and improves customer service across global markets	YES	NO	YES	NO	NLP	mBART BERT
Retrieval-Augmented Generation with Knowledge Graphs for Customer Service Question Answering, 2024 [28]	LLM-based customer service question answering system that integrates RAG with a knowledge graph	YES	YES	YES	NO	RAG with Knowledge Graphs	N/A
TechAssist: A RAG-Based Chatbot for Accessing Technical Information from StackOverflow, 2025 [29]	Specialized question-answering chatbot designed to streamline the retrieval of technical information.	YES	YES	YES	NO	Gradio (for web interface) FAISS (Facebook AI Similarity Search)	Mistral 7B

(a) C.I. - Chatbot Implementation (b) AN/AD – Ability to answer questions (c) Ability to add new knowledge into the system using the chatbot

The selected literature reveals a clear distinction between studies that focus primarily on the technical development and implementation of chatbot solutions [9], [13], [14], [24], [25], [27] [28], [29], and those that examine the organisational impact of AI-powered chatbots, particularly how business processes can be adapted or transformed[23], [26].

The literature highlights several advantages in using chatbots to support employees and business operations. These include improvements in customer service and problem resolution, reduction in response times [14], facilitate user interaction via natural language enhancing communication [14], [23], automation of repetitive tasks, reduction of workload for technical support teams [24] and the ability of chatbots to respond to user queries when no specialised personnel are available that is particularly important for ensuring operational continuity and enhancing employee autonomy [25].

The design and implementation of AI-powered chatbots in the studies reviewed offered several advantages and also presented some challenges. In terms of functionality, all the chatbots discussed were capable of responding to questions based on a predefined dataset.

A wide range of technological approaches can be seen in Table 2 with different development platforms, natural language processing (NLP) techniques, database structures, and integration with external services. The technological diversity across the reviewed studies show the diversity of available tools and the experimental nature of chatbot deployment in the KM field.

Some studies opted for widely accessible frameworks that streamline chatbot development. As examples, Torres et al [24] used Google´s Dialogflow [30] to create a conversational support agent tailored for information technology (IT) service desk management. Vasaniya et al [29] utilised Gradio for its easy web deployment and FAISS (Facebook AI Similarity Search) to support semantic search, ensuring relevance in retrieving technical knowledge from large unstructured data sources. And Lee et al [9] incorporated Pinecone as a vector store for document indexing, combined with Flowise AI and Ollama, enabling the chatbot to perform semantic matching between user queries and relevant technical documentation and proposing a flexible and scalable solution.

Early studies, like Tishchenko et al [13] in 2021 chatbot implementation used classical models like Seq2Seq in Python. Recent studies have adopted Retrieval-Augmented Generation (RAG) integrated with Large Language Models (LLMs) to improve data accuracy [9], [13], bring higher contextual relevance [9], reduce hallucinations by ensuring that the generated responses are based on relevant information [28], access updated data and adaptability of chatbot-based technical support systems.

Most of the recent works on this review use RAG-based architectures. Lee et al., Xu et al [25], and Vasaniya et al. [26] combine the LLMs generative capabilities with external knowledge retrieval, allowing the chatbots to access indexed content in real time before producing a response. Lee et al [9] presented significant improvements when comparing their RAG chatbot with a LLM model, showing the value of this approach. Despite these benefits, Vasaniya et al [29] addressed the challenge of maintaining and updating the external knowledge base.

Data preparation and curation is vital in model training. It is essential to clean and structure the data to ensure its relevance and quality. This includes removing duplicates [13], correcting typos [9], filtering irrelevant information, and tokenising text [13].

The use of vector databases and semantic indexing engines, such as Pinecone, FAISS, is a central element in RAG implementations. These technologies enable the retrieval of the most relevant documents or knowledge chunks before feeding them into the language model for answer generation.

The study of Xu et al [28] explores the integration of knowledge graphs (KGs) with RAG frameworks to enhance retrieval accuracy by preserving the structure and relations within the data. This approach helps in reducing hallucinations by ensuring that the generated responses are based on well-structured and relevant information.

The study conducted by Torres et al. [24], which involved the development of a chatbot solution for an IT support team in Florida (USA), reported positive outcomes, with a 20% improvement in general user satisfaction regarding issue resolution times, and a 12% increase in incident response efficiency among advanced users. In terms of financial costs, the same study also demonstrated an estimated cost reduction of 42.26% through task automation and decreased hiring needs. These advantages, particularly the reduction in manual workload and recruitment demands, align with the current shortage of qualified professionals in technical fields, as discussed by Lee et al. [7].

As chatbot interactions occur primarily through written communication, Benita et al. [14] emphasise the advantages of using multilingual chatbot systems such as the capability of handling technical support requests in multiple languages and ensuring accurate translations while preserving the nuances of technical language. This enables users to receive accurate support in their native or preferred language, enhancing global accessibility. Despite these benefits, chatbots may encounter challenges when dealing with rare languages, complex expressions, or highly specialised technical terms, which can affect their effectiveness and reliability in providing expert support.

One of the emerging questions presented within this technological progression of generative AI is how professional collaboration between humans and AI tools, such as chatbots, will evolve. Based on both theoretical and empirical investigations, Reinhard et al. [26] identify different levels of human and automated participation in task execution. In improvement routines, humans retain an active role, able to accept or reject AI-generated suggestions, such as automated responses or recommendations in collaborative writing tasks. In execution-only routines, human involvement shifts towards data curation and bot supervision. The study's main conclusion is that AI does not replace human agents, but rather complements them, enhancing their strategic involvement and the value of their contributions. This correlates with the conclusions of Poser et al [23] in their study about redesigning IT Support, that present a hybrid work process that introduces the substitution and augmentation of human tasks to address identified work-related issues and improve service delivery using AI.

This literature review reveals a recent and concentrated academic interest in the application of conversational agents within the domain of technical KM and support services. The limited number of available publications, most of which have emerged since 2024, demonstrates the novelty of the field and reflects its current phase of rapid evolution and innovation. The literature shows a gap in the use of AI-powered systems as tools to support the creation and addition of new knowledge within organisational systems. This study aims to address this gap by developing a chatbot prototype that facilitates access to internal technical knowledge while also enabling the addition of new information to support a more dynamic and evolving knowledge base.

CHAPTER 3

Research Methodology

A DSRM was used to develop the AI-powered chatbot prototype, as it provides a structured and rigorous approach for the design, development, and evaluation of technological artefacts (Figure 2).

The methodology starts with the identification of the problem and the underlying research motivation, followed by the definition of clear objectives for the solution. The subsequent phases involve the design and development of the prototype, its demonstration to relevant organization stakeholders, and the evaluation of its performance. If improvements are identified, the process can be repeated through additional iterations, in this case five iterations were completed.

The final phase consists of the communication of results, where the discoveries and contributions of the research are shared with the academic and professional community.

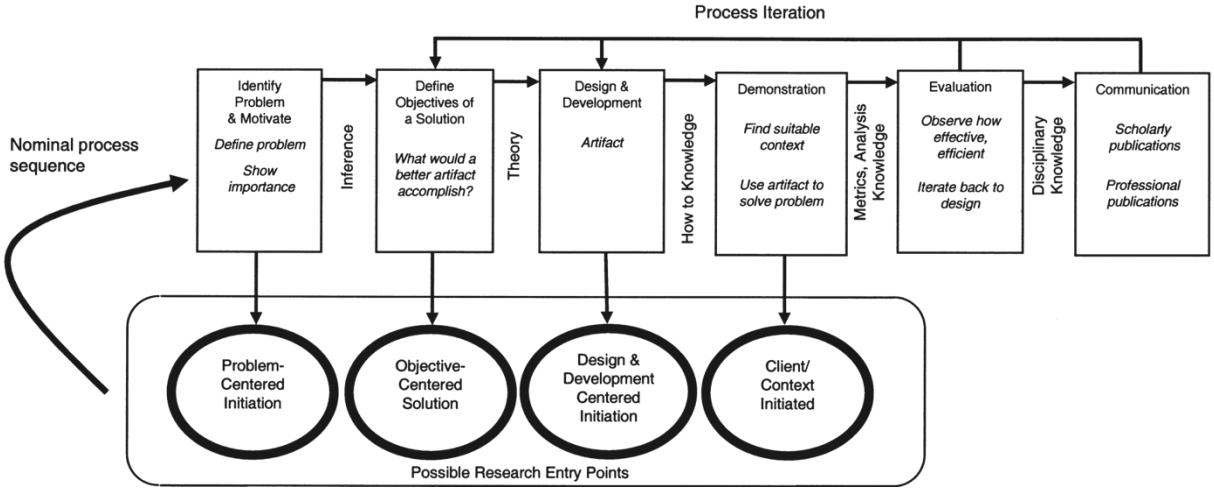


Figure 2 – DSRM diagram

A defining characteristic of DSRM is its iterative nature: each cycle of design, development, demonstration, and evaluation integrates feedback to refine and improve the artefact. This loop is particularly valuable in the case of AI-powered chatbots, where success depends on technical implementation and on usability, accuracy, user experience (UX), and alignment with organisational needs. Through this process, the prototype evolves progressively towards a more effective, reliable, and sustainable solution.

Table 3 presents a structured correspondence between the six phases of the DSRM and the respective sections of the document where each phase is addressed.

Table 3 – DSRM mapping

DSRM PHASE	DOCUMENT SECTIONS
1. Problem Identification and Motivation	Chapter 3 Section 3.1
2. Definition of Objectives for a Solution	Chapter 3 Section 3.2
3. Design and Development	Chapter 4 Sections 4.1.1, 4.2.1, 4.3.1, 4.4.1, 4.5.1
4. Demonstration	Chapter 4 Sections 4.1.2, 4.2.2, 4.3.2, 4.4.2, 4.5.2
5. Evaluation	Chapter 4 Sections 4.1.2, 4.2.2, 4.3.2, 4.4.2, 4.5.2
6. Communication	Chapter 5

3.1. Problem identification

The challenges associated with technical knowledge management are identified in the context of Tekuchi, an international company operating across Portugal, the United Kingdom, Hong Kong, China, Japan, and Australia. Its operations depend heavily on collaboration between distributed teams with distinct responsibilities, ranging from software development to the installation and configuration of systems in technically demanding environments, such as setting up sales offices and short-term events.

Although the company maintains internal communication channels and technical documentation, quick and reliable access to the needed information is not always guaranteed. Time pressure, particularly during the installation and configuration of software, hardware, computer networks and other IT tasks setting up events, makes it difficult to conduct time-consuming searches on existing internal KM platforms or to wait for responses from colleagues who may be unavailable due to time zone differences or competing project demands. In such contexts, the success of the intervention may depend on the ability to obtain rapid and useful answers to specific technical questions.

Another challenge lies in the recurrence of repetitive questions, frequently directed at the same individuals within teams. This not only causes interruptions and reduces efficiency but also places an unnecessary burden on key staff members. An automated mechanism capable of handling recurring technical queries could significantly reduce interruptions, improve efficiency, and promote knowledge reuse.

A centralised and intelligent access point to technical information also plays a critical role in onboarding new employees. A well-structured chatbot would enable newcomers to learn at their own pace, clarify basic questions independently, familiarise themselves with procedures, and build confidence without being overly dependent on more experienced colleagues. Finally, the ability to capture and store newly emerging information in a quick way would strengthen the documentation of evolving subjects and ensure that technical knowledge is continuously updated and preserved.

The findings of the literature review reinforce the significance of these challenges. Prior studies consistently highlight the difficulty of ensuring effective KM in distributed and time-constrained environments, where conventional repositories and communication tools fall short in providing timely and contextually relevant support.

The review also identified the recurring problem of information silos, duplication of effort, and over-reliance on key experts, all of which are observable in the Tekuchi context. The literature suggests that AI-powered conversational agents, particularly when combined with RAG and continuous knowledge capture mechanisms, can effectively address these limitations. By addressing the specific needs of Tekuchi while aligning with broader research insights, the proposed prototype aims to demonstrate both organisational relevance and scientific contribution.

3.2. Objectives of the solution

This solution aims to develop an AI-powered chatbot prototype that improves two main topics: ensuring access to internal technical knowledge and support the inclusion of new documentation within an organisation. Access to technical information should be provided in a simple, rapid, and effective manner, preferably by integrating with communication tools already in use within the company. The solution must also follow a strategy of sustainable growth of the knowledge base, while being designed with the principle that the system should allow future technological evolution. Regarding adding new knowledge, the prototype should make it possible to capture basic information that can later be validated, refined, and incorporated into the system. This ensures that knowledge is not only made accessible but also continuously improved and maintained as a reliable organisational resource.

3.3. Evaluation process

The group includes professionals with diverse technical and managerial roles, ensuring a broad and representative perspective on the system's usability and effectiveness. Table 4 summarises the profiles of the participants involved in the iterative evaluation cycles.

All participants possess significant professional experience within the company which ensures a good understanding of the data used for testing and internal knowledge management practices.

Regarding communication tools, all participants had several years of experience using Slack, confirming that the UI selected for TekBot was aligned with users' daily workflows. In terms of prior familiarity with AI-based assistants, most participants reported daily use of AI chatbots, with only one participant indicating occasional use. This ensured that all users were sufficiently familiar with conversational AI, allowing the evaluation to focus on the system's functionality and integration rather than on basic tool adoption.

Table 4 – DSR iterations

ITERATION	ROLE	COMPANY EXPERIENCE	SLACK EXPERIENCE	AI CHATBOT USE
1	Senior Interactive Project Manager	10 Years	5 Years	Daily
2	Senior Interactive Designer	15 Years	5 Years	Daily
3	Senior Interactive Developer	3 Years	3 Years	Daily
4	Technical Project Manager – Hardware and Exhibitions	12 Years	5 Years	Occasionally
5	Senior Interactive Developer	12 Years	5 Years	Daily

During the testing sessions, each participant was asked to perform a series of predefined tasks, including submitting technical questions and requesting the insertion of new information into the system.

Their interactions were observed to assess how easily they were able to complete these tasks and how accurate and helpful the bot's responses were. In addition, Table 5 questionnaire was used to collect feedback on various aspects of the experience, including ease of use, accuracy of answers, and the likelihood of using the chatbot in their daily work. The combination of direct observation and structured feedback provided valuable insights into TekBot's performance and helped identify areas for improvement in future iterations.

The test was conducted remotely via a Zoom call, during which the participant shared their screen. The user interacted with TekBot using its own Slack software and account.

Two main tasks were defined for this test: asking TekBot a range of technical questions and submitting requests for the chatbot to store new information. Before beginning, the user was briefed on the type of data currently available in the system, the Slack channel to be used, and the need to mention TekBot on all messages. The interviewees were then asked to pose several questions to the system and to make a few data storage requests. Following these interactions, a short interview was conducted based on the follow-up questions presented in Table 5.

Table 5 – Interview questionnaire

ID	QUESTION	POSSIBLE ANSWERS
Q1	How easy was it to interact with TekBot when asking questions?	Very hard/Hard /Neutral /Easy /Very Easy
Q2	How easy was it to register new information?	Very hard/Hard /Neutral /Easy /Very Easy
Q3	Were the answers accurate?	Very Satisfied/Satisfied/Neutral/ Dissatisfied/Very Dissatisfied
Q4	Were the answers useful?	Very Satisfied/Satisfied/Neutral/ Dissatisfied/Very Dissatisfied
Q5	How satisfied are you with the response time?	Very Satisfied/Satisfied/Neutral/ Dissatisfied/Very Dissatisfied
Q6	Would you use TekBot as part of your daily work routine?	Yes/No
Q7	In which situations do you see yourself using TekBot?	Open Answer
Q8	Do you have any suggestions for improving TekBot?	Open Answer

CHAPTER 4

TekBot Prototype

4.1. First DSR Iteration

4.1.1. Design and Development

Considering the defined objectives and the understanding of the organisational problem at Tekuchi, this AI powered chatbot solution, TekBot, was structured around a few fundamental components, which are presented on the diagram below (Figure 3). These include the communication between the user and the system through tools already adopted by the company, the decision-making process to determine whether a user request represents a query to be answered or a contribution of new knowledge, the addition of new information into the system, the ability to provide accessible and well-structured responses in natural language, and finally the storage of new knowledge whenever requested by the user.

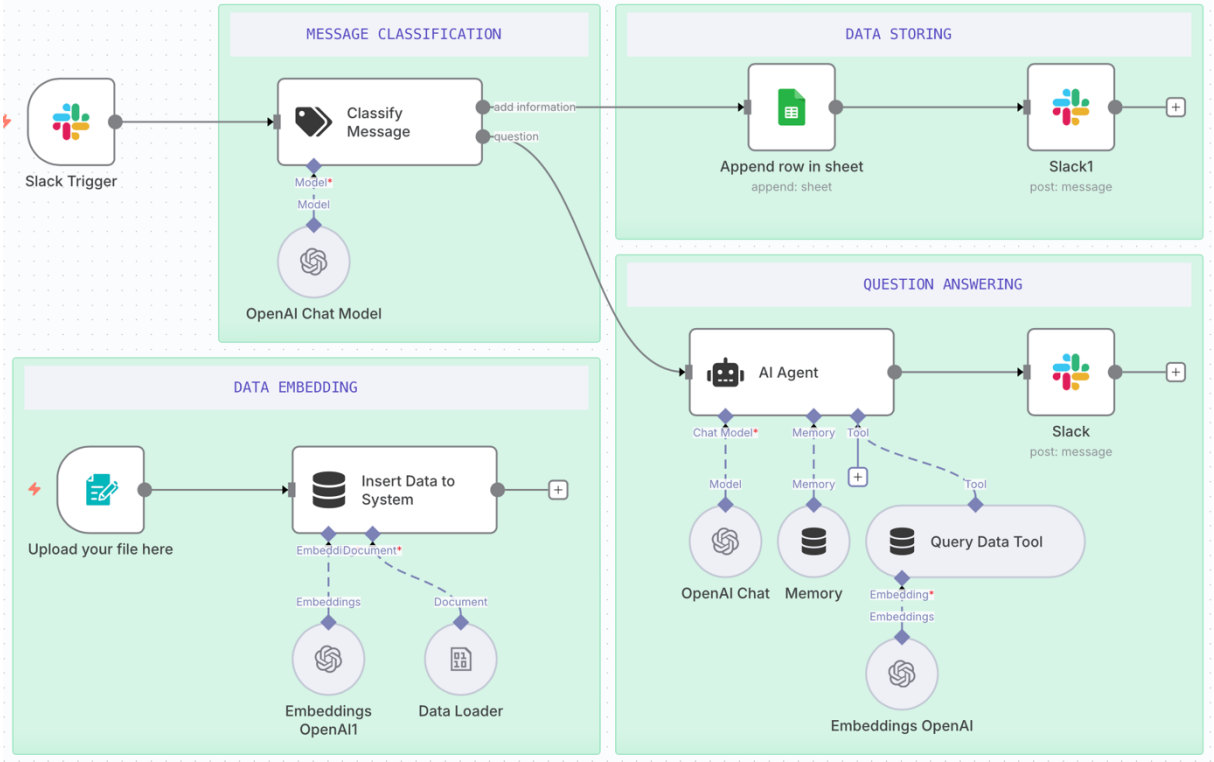


Figure 3 – n8n workflow

To respond to the above needs, the selected solution for the design and development of the artefact was the n8n platform, selected for its versatility, fast prototyping features and capacity to integrate a wide range of services in a seamless manner.

n8n is a workflow automation system that enables the orchestration of complex processes using a modular visual node-based approach that doesn't require advanced knowledge in programming languages. Its architecture supports connections to a large set of Application Programming Interfaces (API), databases, third-party applications allowing flexibility and scaling. The platform can be installed and self-hosted or operated as a cloud-based service; for this artefact, the latter option was chosen to ensure simplicity, ease of maintenance, and a faster development turnaround.

As required by the scope of this project, n8n facilitated the integration of the different components required for the TekBot Prototype and the third-party applications, cloud services and AI models. This ability to evolve iteratively aligns well with the DSRM, as it simplifies the design and development of the artefact and allows continuous improvement in response to evaluation and feedback. The modularity of the platform also allows a foundation of future enhancements, allowing new integrations and functionalities to be incorporated.

Following an analysis of the most used tools within the company, Slack was chosen as the user interface (UI). Slack is used by all employees and provides the necessary functionalities for exchanging information across a wide variety of internal topics, supporting both daily operations and communication between teams, departments, and offices. It is therefore considered an essential element of the company's workflow.

The integration of the chatbot prototype into Slack brings several advantages: it ensures ease of access, minimises the time required for training or adaptation to this new way of accessing knowledge, and reduces the need for additional software or websites. A further significant advantage is that Slack has been extensively tested, and its availability is guaranteed across all the geographical areas in which the company operates, including China, where many platforms and services face considerable restrictions.

The integration between n8n and Slack is possible and requires a set of security steps to ensure the communication between both systems. The availability of the chatbot to the company users is ensured through a virtual user named TekBot.

To interact with the prototype, employees use a public channel within Slack company workspace, where they can either post a question or submit a request to add new information by mentioning the TekBot user. The system receives the message, processes it, and replies directly in the same channel using natural language. The response may either provide the requested information or confirm that a new knowledge item has been successfully stored.

When a message is addressed to the system, its content is first assessed to determine its intent. This validation is performed through AI, using an integration with the GPT-4.1 nano model from OpenAI. A predefined prompt is applied, which guides the model to decide whether the message represents a request to add information to the system or a question to be answered using the existing knowledge base (Figure 4).

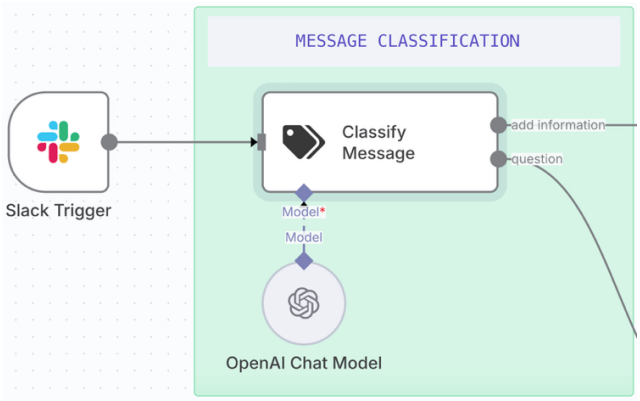


Figure 4 – Message classification

If the message received is identified as a request to store new information, the system records the content in a cloud-based solution (Figure 5). In this iteration, Google Sheets was selected to record the information, as it is already widely used within the company's daily operations and it can be integrated with the n8n platform. The organisation has the necessary user accounts, folder structures, and security procedures in place, making it a reliable and accessible tool for all employees. It also benefits from an established access and sharing framework that aligns with the company's internal governance. To integrate both systems it was necessary to setup all the necessary permissions on the Google platform.

A dedicated spreadsheet was created to capture and register the information submitted through TekBot. Each entry stores the name of the user who submitted the request, the text of the message, and the date and time when the request was made. This setup enables regular reviews of the collected data, which can then be validated, organised, and prepared for structured inclusion. In this way, the information can be added into the TekBot system and made available through other existing platforms, ensuring consistency and integration across the company's KM system. When new information is successfully added to the Google Spreadsheet, the system automatically posts a confirmation message in Slack, mentioning the user and indicating that the information has been recorded successfully.

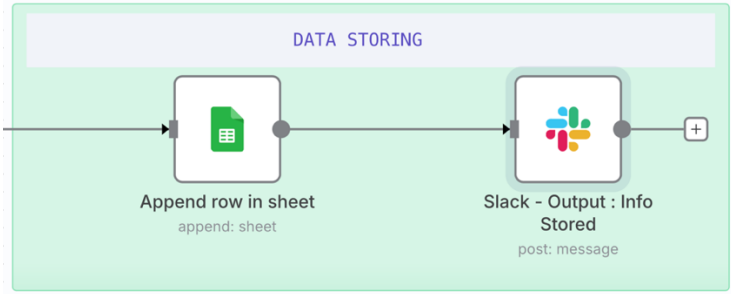


Figure 5 – Data storing

For TekBot to be able to respond effectively to user queries, several key elements are required. These include access to an AI model capable of generating natural language answers, a memory mechanism to retain previous interactions, and an indexed collection of company data to serve as a knowledge base (Figure 6). For this artefact, ChatGPT-4.1 Nano was chosen to provide answers to user questions. The model was configured through a predefined system prompt, which establishes the necessary context and provides clear instructions to ensure that responses are delivered with accuracy, clarity, and a communication style aligned with the organisational context.

In addition to the integration with the AI model, TekBot relies on a memory component to ensure consistency and contextual relevance in its responses. The memory enables the system to retain previous interactions, which allows follow-up questions to be contextual and supports a more natural conversational flow.

The Query Data Tool makes use of OpenAI Embeddings to transform both stored documents and user queries into vector representations, enabling semantic similarity search. This allows the system to retrieve the most relevant organisational knowledge items and provide accurate and contextually appropriate answers. To minimise the risk of hallucinations or contradictory information, and to ensure that outputs remain focused exclusively on organisational data, the system prompt explicitly instructs the model not to access or rely on external online sources.

These elements ensure better quality in the responses and alignment with the company's technical knowledge and practices. The system prompt used is:

"You are an AI assistant designed to provide accurate answers based solely on the technical information contained in the company's internal documents.

- Do not retrieve or rely on any external online sources.

- Do not use general model knowledge that is not supported by the provided documents.

- If the requested information is not available in the documents, reply with: "I don't have information about that topic."

- Ensure that all responses are clear, concise, and written in a professional tone"

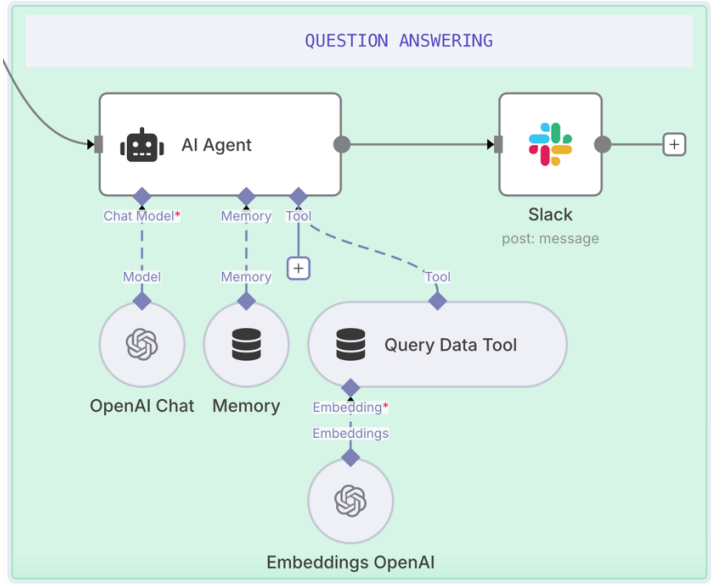


Figure 6 – Question answering

Figure 7 illustrates the data embedding workflow implemented to integrate information into the TekBot system. The system allows uploading files containing relevant data. Once uploaded, the document passes through a Data Loader, which prepares the content for processing.

The prepared text is processed using the OpenAI Embeddings model, which converts the textual content into vector representations (embeddings). The embeddings allow the system to perform similarity searches and retrieve contextually relevant information during user queries.

The document data and its corresponding embeddings are inserted into the system's vector database through the Insert Data to System module ensuring that the data is searchable and accessible within the TekBot knowledge base.

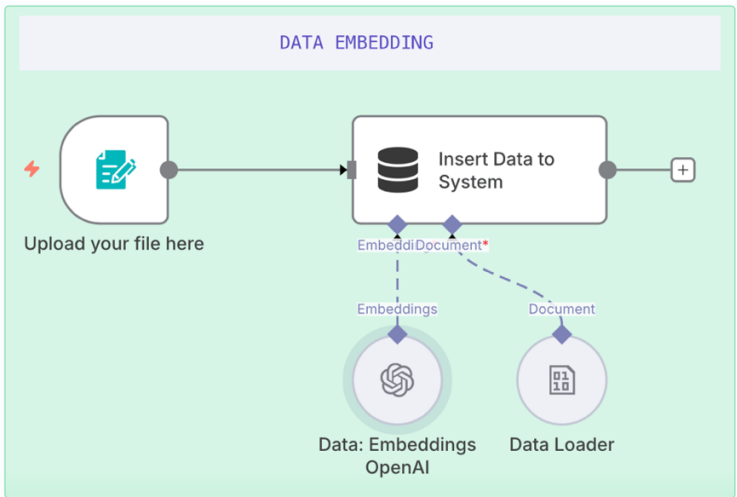


Figure 7 – Data embedding

The use of n8n platform ensured that the different components could be integrated into the system and communicate according to their specific needs. Table 6 provides a summary of the components above described and the technologies used for each functionality.

Table 6 – First iteration: main components and technology

COMPONENT	DESCRIPTION	TECHNOLOGY
Communication Interface	The channel through which users interact with TekBot. In this project, Slack is used as the primary interface for submitting questions and adding new information.	Slack
AI Model	The AI model responsible for interpreting user messages, determining intent, and generating responses.	GPT-4.1 Nano
Knowledge Repository	An indexed database that stores uploaded company data and supports information retrieval through semantic similarity search.	OpenAI Embeddings
Memory Module	Component that retains previous interactions for conversational context	n8n Memory Node
Embeddings	Component that transforms text data and user queries into vectors to enable semantic matching and RAG	OpenAI Embeddings
New Information Storage	Used to store new knowledge contributed by users before validation, preparation and integration into the system.	Google Sheets (Cloud-based)

After defining the components and technologies that compose the TekBot Prototype, the next step is to demonstrate how the system will be used in the company operational contexts. The following examples help to understand how the prototype supports both information access and knowledge growth within the organisation.

In a typical scenario, an employee, who is planning the installation of an application, posts a message to a designated Slack channel. The message is received by the system via n8n, where it is analysed to determine whether it represents a question or a request to add new information.

Figure 8 presents an example where the user prompt is classified as a question, and the AI agent accesses the vector store to retrieve the most relevant internal documents and uses the GPT-4.1 Nano model to generate a clear and context-aware response. If there is information available to answer the query, it is then returned to the user within the same Slack channel. If there is no information available to reply to that question the system posts a message stating that there is no information available about that topic.

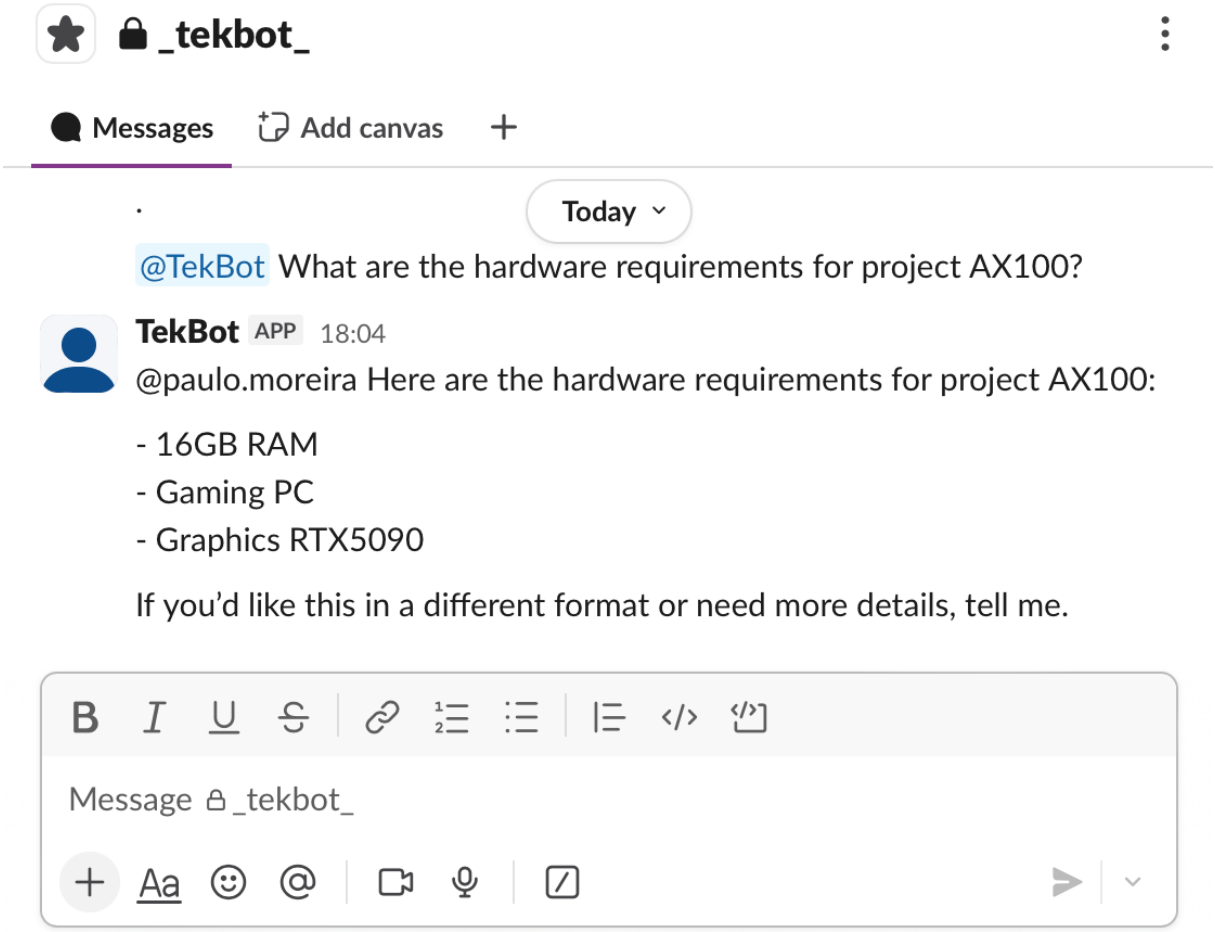


Figure 8 – Example of question

In another common situation, a user may ask TekBot to store new information, such as a recently updated procedure or configuration guideline. In this case, the system recognises the intent to add knowledge, validates the input, and records it into Google Sheets. This process ensures that new insights are captured efficiently and can later be reviewed and integrated into the main knowledge repository, maintaining the continuous growth and accuracy of the company's knowledge base.

4.1.2. Demonstration and Evaluation

For this first iteration, the testing was conducted with a senior project manager, that has been actively involved in contributing to and consulting the organisation's KM system, has more than five years of experience using Slack, and employs AI Tools like ChatGPT daily.

The results of this first test, presented in Table 7, were overall positive, demonstrating that the TekBot Prototype is intuitive and easy to use.

Table 7 – First iteration: interview results

ID	QUESTION	ANSWER
Q1	How easy was it to interact with TekBot when asking questions?	Very Easy
Q2	How easy was it to register new information?	Very Easy
Q3	Were the answers useful?	Satisfied
Q4	Were the answers accurate?	Neutral
Q5	How satisfied are you with the response time?	Satisfied
Q6	Would you use TekBot as part of your daily work routine?	Yes
Q7	In which situations do you see yourself using TekBot?	Would use TekBot regularly, primarily to search for information without leaving Slack.
Q8	Do you have any suggestions for improving TekBot?	Answers should be more accurate. Chatbot should indicate clearly when no relevant information is available and avoid giving incorrect or incomplete answers. Provide more factual responses. Support for data comparison and list generation. Broader knowledgebase would be more useful. Personable and polite communication style. Proactive suggestions when no data are found, could improve the user experience.

The participant rated the interaction experience as very easy both when asking questions and when adding new information. The usefulness of the answers was considered satisfactory in terms of utility, but the accuracy was only rated as neutral, indicating there was still room for improvement in the precision of the information retrieved. The response speed was evaluated positively.

When asked about general use, the participant confirmed that she would use TekBot regularly, primarily to search for information without leaving Slack, emphasising that the integration within the company's main communication tool was a major advantage. She suggested that TekBot would be more useful if it could also handle other data operations, such as comparing items, providing lists or totals based on specific variables or conditions.

In terms of suggestions for improvements, the participant recommended that the chatbot should indicate clearly when no relevant information is available and avoid giving incorrect or incomplete answers. And, suggested that the bot should provide more factual responses and that a broader knowledge base would be more useful.

Additionally, she mentioned that a more personable and polite communication style, as well as proactive suggestions when no data are found, could improve the user experience.

4.2. Second DSR Iteration

After analysing the interview results and suggestions, the improvements outlined in Table 8 were defined to guide the refinement of the artefact's second iteration. During the design and development phase, it was decided to focus on the items suggested by the interviewee I1, I2, I3, and I6, which mainly relate to improving the accuracy and quality of the chatbot's responses. The ability to compare data was left for future iterations (I4). Additionally, item I8 was addressed to ensure that the data added to the system remains persistent and is not lost after a period or following server restarts.

Due to time constraints and to maintain focus on improving the system, it was decided to postpone the implementation of proactive suggestions when no data are found (I7) to future developments. Similarly, the expansion of the dataset (I5) was also deferred to a later phase, as the current data collection is sufficient for prototype testing. This decision allowed the project to prioritise system refinement and performance enhancement before introducing larger and more complex data sets.

Table 8 – Second iteration: improvement suggestions

ID		SUGGESTED BY	IMPLEMENTED
I1	Answers should be more accurate.	Interviewee 1	Yes
I2	Indicate clearly when no relevant information is available and avoid giving incorrect or incomplete answers.	Interviewee 1	Yes
I3	Provide more factual responses.	Interviewee 1	Yes
I4	Ability to compare data and provide list based answers	Interviewee 1	No
I5	Broader knowledgebase would be more useful.	Interviewee 1	No
I6	Personable and polite communication style.	Interviewee 1	Yes
I7	Proactive suggestions when no data are found, could improve the user experience.	Interviewee 1	No
I8	Data store should be persistent and not temporary	Author	Yes

4.2.1. Design and Development

Taking into account the interviewees' feedback and the identified technical improvement needs, the system was enhanced in four main areas: improving the quality and structure of the data processed by the workflow; optimising the embedding generation process within n8n to ensure better context representation; integrating a vector store to persist the data and improve the accuracy of vector-based searches; and testing and selecting an AI model that delivered the best overall performance for the chatbot.

The dataset was reviewed to identify and correct inconsistencies and inaccuracies, ensuring that all data was clearly structured and consistently formatted. The embedding was updated to include a recursive character text splitter and define that the data chunks should take 1000 characters with an overlap of 200 characters.

To address the limitations related to data persistence and improve the retrieval accuracy, Pinecone was integrated into the system as the vector database (Figure 9). This integration enables the embeddings generated within the n8n workflow to be stored and indexed efficiently. The first artefact was using temporary in-memory storage with the Pinecone integration the information remains accessible even after system restarts or workflow updates. Pinecone provides long-term persistence and scalability, allowing the system to handle larger datasets and maintain consistent performance over time.

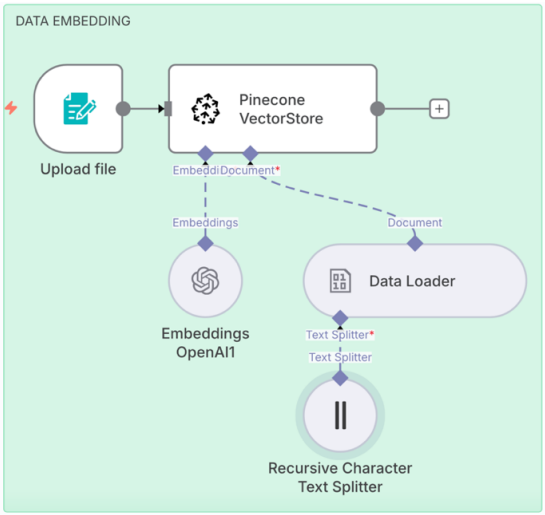


Figure 9 – Pinecone VectorStore integration for data embedding

The integration of Pinecone enhanced the semantic search capabilities of the chatbot. Its vector-based retrieval system enables the extraction of semantically relevant information rather than relying solely on keyword matching, leading to more accurate and meaningful responses (Figure 10).

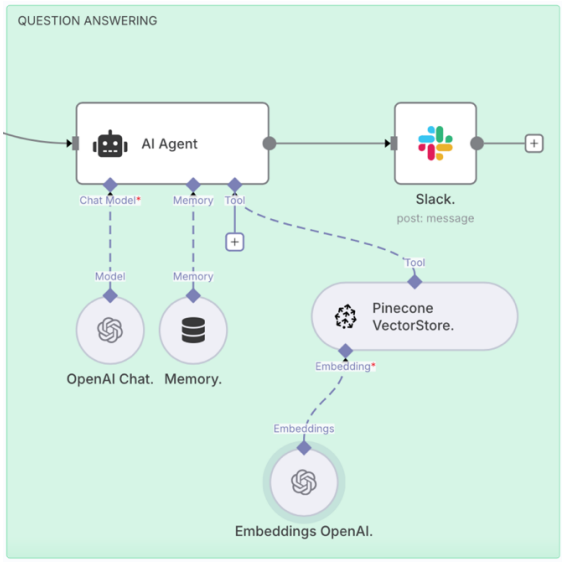


Figure 10 – Pinecone VectorStore integration for AI agent tool

To improve the accuracy and overall quality of the responses, as well as their clarity of expression, two alternative models were tested to compare with the GPT-4.1 Nano model used in the first iteration. For this evaluation, the n8n platform was once again employed, and a testing workflow was created using the existing modules, embeddings and system prompts. This setup enabled each question submitted through Slack to generate three responses simultaneously, one from each model, allowing for a direct comparative analysis. The assessment focused primarily on the accuracy of the answers. The models selected for this test were two of the most recent lightweight and cost-efficient options available from OpenAI and Google: GPT-5 Nano and Google Gemini.

The test was conducted using a set of thirty different questions. For each question, the responses generated by the three models were analysed and compared with the reference data previously entered the system (Table 9). This evaluation demonstrated that both GPT-5 Nano and Google Gemini 2.5 produced responses of higher quality and accuracy compared to GPT-4.1 Nano, which had been used in the first iteration. Further analysis show that incorrect answers on GPT-5 Nano were related to data comparison.

The results between GPT-5 Nano and Google Gemini were very similar; GPT-5 Nano was selected for subsequent iterations due to its easier integration with the existing prototype, as the previous version had already been based on an OpenAI model.

Table 9 – AI Model comparison

		CORRECT ANSWERS	
MODEL	**QUESTIONS**	**TOTAL**	**%**
GPT-4.1 Nano	30	19	63
GPT-5 Nano	30	27	90
Google Gemini 2.5 Flash	30	26	87

4.2.2. Demonstration and Evaluation

For the second iteration, testing was conducted with a senior interactive designer who has been with the company for over fifteen years. He has played a leading role in integrating the company's primary KM tool (Notion), has more than five years' experience with Slack, and uses multiple AI chatbot models daily.

The test results were highly positive, both in terms of usability and the ease with which the user was able to obtain responses and add new knowledge (Table 10). The participant expressed satisfaction with the quality and usefulness of the answers provided, as well as with the way the bot responded even when it did not have access to the exact information being requested. The answers given were consistently accurate, demonstrating an improvement in the system's precision and relevance.

Table 10 – Second iteration: interview results

ID	QUESTION	ANSWER
Q1	How easy was it to interact with TekBot when asking questions?	Very Easy
Q2	How easy was it to register new information?	Very Easy
Q3	Were the answers useful?	Very Satisfied
Q4	Were the answers accurate?	Very Satisfied
Q5	How satisfied are you with the response time?	Satisfied
Q6	Would you use TekBot as part of your daily work routine?	Yes
Q7	In which situations do you see yourself using TekBot?	Would use TekBot daily. It is a fast and easy way to access information.
Q8	Do you have any suggestions for improving TekBot?	Save answers into Notion, the current company KM. Ensure the bot always gives feedback to any user interaction.

During the testing process, there was one instance in which the user received no response from the bot. This occurred because the interaction was neither a question nor a request to save information. The lack of feedback in this situation created some uncertainty as to whether the bot was still processing the input or had stopped functioning. This was therefore identified and suggested as an area for improvement in the next iteration.

The participant also stated that TekBot would be a tool he would use daily, emphasising the ease and speed of accessing information as its main advantages. He highlighted that although most of the answers provided during the test were already available within the company's internal KM system, using TekBot as a first point of access would be considerably more practical and efficient.

Regarding improvements, it was suggested that the information stored by the bot should be saved directly into Notion, instead of using Google Sheets as in this testing phase. That integration would contribute to feeding both the bot itself and the KM system more efficiently.

4.3. Third DSR Iteration

The improvements identified during the second iteration focused on enhancing the process of information storage and the clarity of the bot's feedback (Table 11). Both suggestions were implemented during the design and development phase.

Table 11 – Third iteration : improvement suggestions

ID		SUGGESTED BY	IMPLEMENTED
I1	Ability to compare data and list-based answers	Interviewee 1	No
I2	Save answers into Notion, the current company KM platform.	Interviewee 2	Yes
I3	Ensure the bot always responds to user input.	Interviewee 2	Yes

4.3.1. Design and Development

The feature for storing information was initially based on Google Sheets. During the previous iteration, it was suggested that this functionality should be updated to integrate directly with Notion (I2). This change offered clear advantages for the company, as Notion is already the organisation's primary KM platform. It would also serve to enrich both the bot's internal knowledge and the information available on the Notion platform itself.

The implementation of this change was facilitated by n8n, which allowed for a quick and straightforward transition. The necessary access credentials were created, and a dedicated database was set up within Notion. On the n8n side, the Google Sheets module was replaced with a Notion module, enabling it to record the collected data, including the message content, author, and registration date.

The second improvement (I3) focused on ensuring that the bot always provides feedback whenever there is user interaction. In the previous version, the system only attempted to classify messages as questions or requests to store information. As a result, conversational exchanges such as greetings or expressions of gratitude (e.g., “thank you”) were not interpreted appropriately, and messages that could not be understood caused the system to remain silent, leaving the user uncertain about what was happening.

To make the bot consistently responsive and more cordial towards users, the message classification module was updated. This module was expanded to include a third classification category, identifying messages that are neither information requests nor data storage commands as “other” (Figure 11).

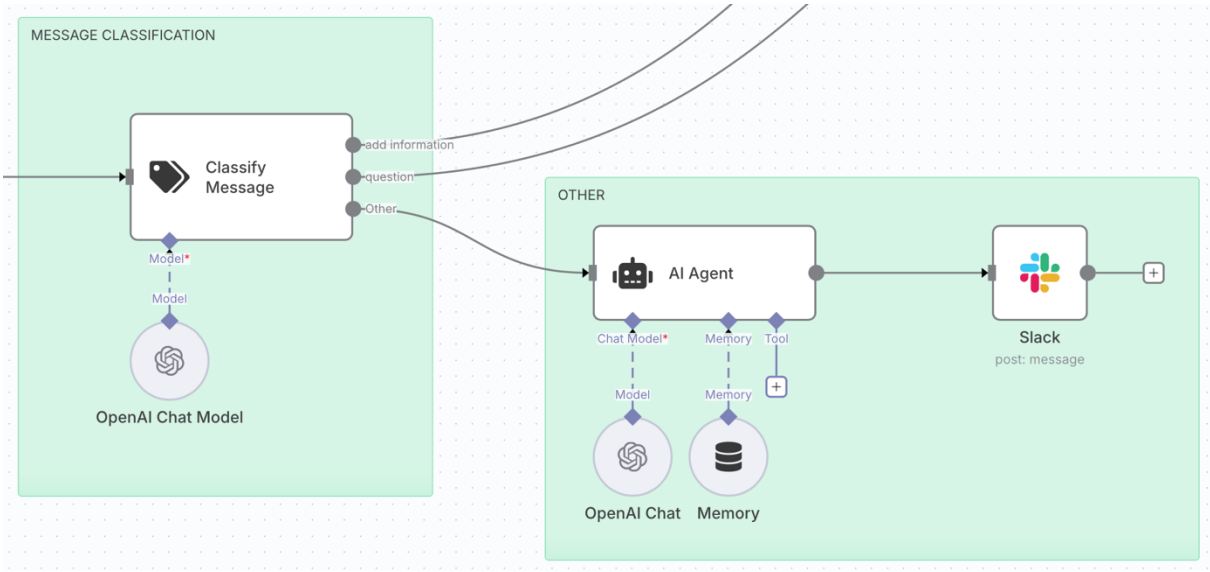


Figure 11 – “Other” AI agent subsystem

This new category led to the creation of an additional subsystem designed specifically to handle such interactions. To achieve this, an AI module based on the GPT-5 Nano model was implemented, consistent with the other AI assistants used within the system. To ensure that this assistant behaved as intended, the following system prompt was created: “You are an AI assistant for a technology company. Your goal is to provide concise and accurate answers in a friendly and helpful manner. You can answer questions about applications developed by the company or help users store information. To get answers, the user simply needs to ask a question. To store information, the user can describe in natural language what they want to save. Only suggest or provide information related to the applications produced by Tekuchi.”. This prompt ensured a coherent and user-friendly conversational flow throughout interactions (Figure 12).

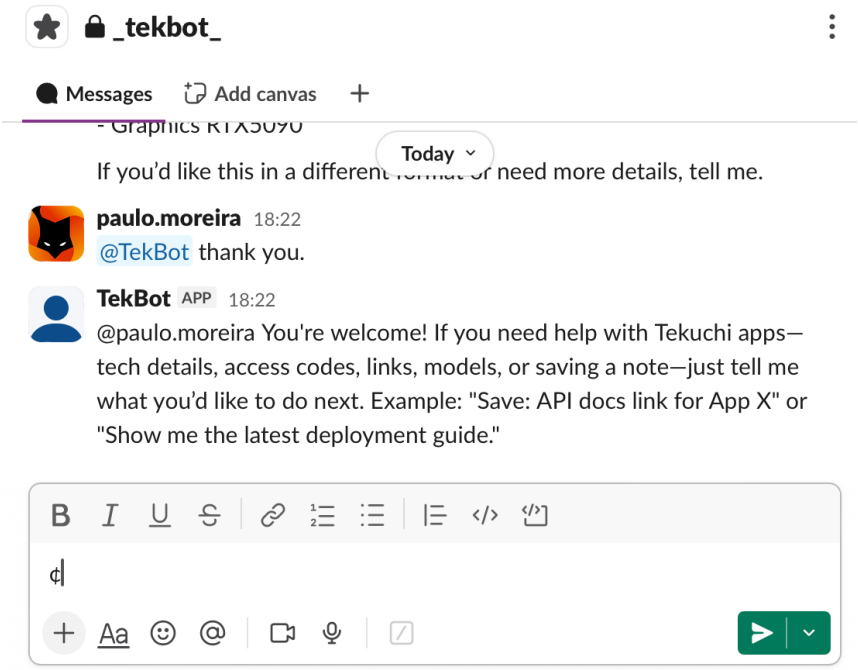


Figure 12 – Example of conversational flow between user and TekBot

4.3.2. Demonstration and Evaluation

The third iteration of testing was carried out with a senior software developer who has been part of the organisation for more than five years and has extensive experience in software engineering and frequent use of AI-based tools, including conversational agents and code-assistance models.

The participant submitted several questions and reported that both storing new data and interacting with the bot during QA process were very easy (Table 12). However, when asked to compare data, retrieve information that required combining multiple data points, or generate structured lists, the responses were often incomplete and, in some cases, incorrect. This issue was also indicated by another interviewee on the first iteration testing.

Although the interviewee indicated that he would see value in using the system regularly, particularly for storing information efficiently and retrieving data that he does not consult often, he also believed that the current limitations in response quality would discourage him from continuing to use the system.

Table 12 – Third iteration: interview results

ID	QUESTION	ANSWER
Q1	How easy was it to interact with TekBot when asking questions?	Very Easy
Q2	How easy was it to register new information?	Very Easy
Q3	Were the answers useful?	Very Satisfied
Q4	Were the answers accurate?	Dissatisfied
Q5	How satisfied are you with the response time?	Very Satisfied
Q6	Would you use TekBot as part of your daily work routine?	Yes
Q7	In which situations do you see yourself using TekBot?	To store information during my daily tasks. Retrieve information that I do not use regularly.
Q8	Do you have any suggestions for improving TekBot?	Improve answers accuracy. Allow the user to provide software code and receive responses related to it. Enable the creation of a personal bot that can store personal information for later retrieval.

4.4. Fourth DSR Iteration

In this iteration, the focus was on addressing an issue identified during the first and third iteration demonstrations (Table 13). While the chatbot performed well when responding to questions related to a single knowledge item, it showed limitations when queries required accessing and combining information from multiple items. Tasks involving data comparison or the generation of structured lists revealed weaknesses in the system's ability to process and organise information across different sources.

Table 13 – Fourth iteration: improvement suggestions

ID		SUGGESTED BY	IMPLEMENTED
I1	Improve quality of comparisons and lists.	Interviewees 1 &3	YES

4.4.1. Design and Development

To address this limitation, it was necessary to modify both the system architecture and the nature of the input data. Until this point, the knowledge base had been composed mainly of unstructured textual documents. It was therefore decided to transition towards the use of structured data to evaluate whether this change would improve response quality and resolve the issues identified above.

Given that a significant portion of the organisation's data is already stored within the Notion system, which includes internal databases that can be accessed via n8n, this platform was selected as the source of structured data. A system modification was implemented in which the AI agent module was directly connected to a Notion database as an external tool. Following this change, a new test was conducted using a set of comparative questions involving multiple items and requests for complete listings.

The system demonstrated a significant improvement in response quality, clearly outperforming the previous approach based on unstructured embedded documents.

An alternative approach was subsequently explored based on the use of structured data stored in relational tables and accessed through SQL queries. For this purpose, the Supabase platform was adopted, enabling the online hosting of the database and the creation of the required tables. In this configuration, the AI agent was connected to several external tools that allow it to read the database schema, access table contents, and retrieve data through SQL operations (Figure 13).

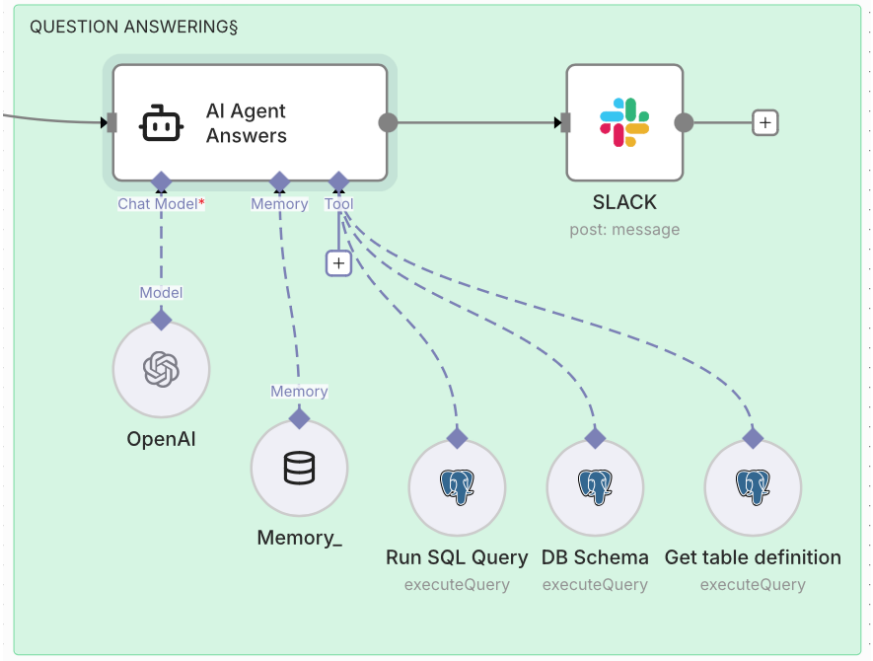


Figure 13 – AI Agent with SQL-based retrieval

The system instructions of the AI agent were then updated to explicitly guide the agent to translate user questions into SQL queries without deleting or altering any data. This allowed the agent to retrieve the necessary data directly from the database and use it to compose more accurate, structured, and context-aware responses for the user.

This approach presented clearly better results than the Notion-based retrieval approach (Table 14). The SQL-based architecture outperforms the Notion-based approach across all evaluated metrics.

Table 14 – Performance comparison between Notion-based and SQL-based retrieval

Metric	Notion-based retrieval	SQL-based retrieval
Correct and complete answers (out of 15)	10/15	14/15
Average response time (seconds)	36	24
Average tokens used per response	37.047	7.566

In terms of response correctness and completeness, the SQL-based solution achieved 14 correct answers out of 15, compared to 10 out of 15 with the Notion-based approach. The average response time was reduced from 36 to 24 seconds, demonstrating a substantial improvement in system responsiveness. An important factor is also the average number of tokens consumed per answer, decreased dramatically, from 37,047 to 7,566 tokens. This represents a reduction of approximately 80% in token usage, with direct implications in efficiency, user experience and operational cost.

These results empirically confirm that structured SQL-based retrieval is more scalable and cost-effective solution for AI-driven organisational knowledge access.

With this update, the system gained the ability to perform data comparisons (Figure 14) and generate structured lists from the stored information.

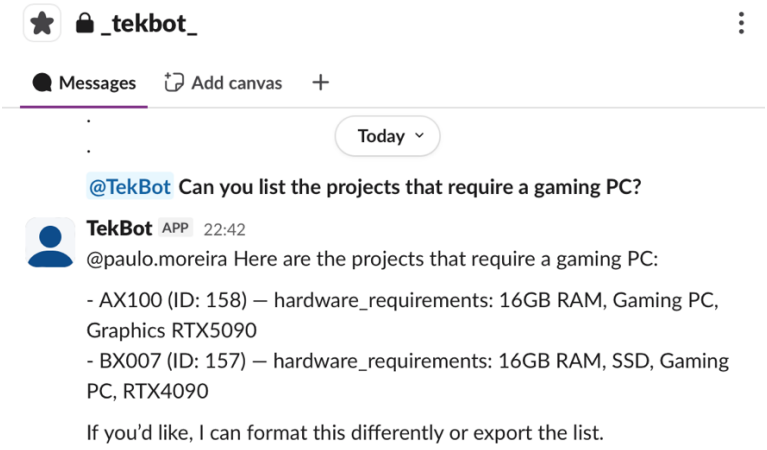


Figure 14 – Example of list generation and comparison

4.4.2. Demonstration and evaluation

The fourth iteration of testing was conducted with a Technical Project Manager who has been part of the organisation for more than twelve years. His role involves the technical design of sales centres and events, including hardware selection, system setup and installation, as well as the overall technical planning and execution of events. Although he has some general knowledge of chatbots, he does not use AI-based tools or conversational agents daily.

During the session, the interviewee submitted several questions and multiple requests to store new data. He expressed a positive impression of the chatbot's responses (Table 15), describing them as clear and accurate. However, he suggested that the bot should provide some form of feedback indicating that it is processing a request, so that users can better understand that any waiting time is normal and does not indicate a system failure.

Table 15 – Fourth iteration: interview results

ID	QUESTION	ANSWER
Q1	How easy was it to interact with TekBot when asking questions?	Very Easy
Q2	How easy was it to register new information?	Very Easy
Q3	Were the answers useful?	Very Satisfied
Q4	Were the answers accurate?	Very Satisfied
Q5	How satisfied are you with the response time?	Neutral
Q6	Would you use TekBot as part of your daily work routine?	Yes
Q7	In which situations do you see yourself using TekBot?	Every day, to find information needed.
Q8	Do you have any suggestions for improving TekBot?	Indicate when the bot is processing a request. Provide a direct link to where the stored data can be consulted and manually updated. Request user clarification when a data storage instruction is unclear or does not make sense.

The participant successfully carried out several data storage operations. It was observed that, in some cases, the instructions provided for saving information were incomplete or unclear, yet the bot still proceeded with the storage action. The interviewee indicated that he would like to know where the data are stored to allow future consultation or manual adjustment if necessary.

Additionally, he suggested that the bot should prompt the user for clarification whenever a request to store information does not appear to be meaningful or sufficiently well defined.

The user considered the system to be highly useful and stated that he would use it daily due to the ease and speed with which information can be accessed. He also expressed that he was particularly impressed by the system's capabilities and by the practical value that the tool demonstrated during the testing session.

4.5. Fifth DSR Iteration

Following the fourth iteration test and the analysis of the feedback presented in Table 16, the tasks performed in this iteration focused primarily on improving the quality of system feedback provided to the user. These improvements included informing the user when the system was actively processing a response, as well as introducing a quality control mechanism in the data storage process. In this context, users are now notified whenever the content submitted for storage is not properly understood, and explicit feedback is provided through the presentation of a direct link to the corresponding entry in Notion. This enhancement allows users to verify and, if necessary, edit the stored information.

Table 16 – Fifth iteration: improvement suggestions

ID		SUGGESTED BY	IMPLEMENTED
I1	Indicate when the bot is processing a request.	Interviewee 4	YES
I2	Provide a direct link to where the stored data can be consulted and manually updated.	Interviewee 4	YES
I3	Request user clarification when a data storage instruction is unclear or does not make sense.	Interviewee 4	YES

4.5.1. Design and Development

To improve the quality of the content being stored, the system was updated so that, upon receiving a message, the AI-based message classification module categorises storage requests into two distinct groups. The first category includes storage requests in which the content is considered relevant and clearly associated with valid technical information, such as application details, hardware configurations, how-to guides, operational instructions, and other technical data. The second category includes storage requests in which the submitted message does not appear to contain valid information or is visibly incomplete. To enable this improvement, the instructions of the AI classification node were updated in accordance with these validation rules.

When a message is classified as a relevant storage request, the information is stored automatically in Notion. When the message is classified as incomplete or invalid, the system sends a notification to the user requesting clarification before proceeding with storage.

Once the information is successfully stored, the user is informed that the request has been completed and is provided with a direct link to the corresponding Notion entry, allowing the user to review, edit, or enhance the recorded content if necessary (Figure 15).

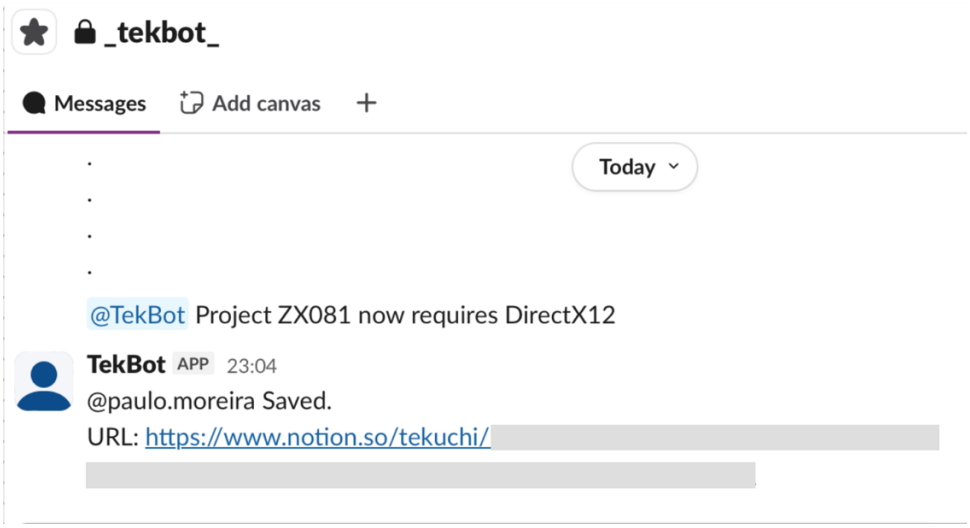


Figure 15 – Save request interaction

Table 17 presents the performance of the message classification module, based on a test set of 40 interactions distributed across the three main classifier categories and the alternative “other” category. The classifier demonstrated a high level of accuracy, correctly classifying 37 out of the 40 messages, corresponding to an overall accuracy of 92.5%.

Table 17 – Message Classifier analysis

MESSAGE TYPE	NUMBER OF MESSAGES	CORRECT	ACCURACY
Question	10	10	100%
Save request	10	10	100%
Save request with invalid data	10	7	70%
Other	10	10	100%
Total	**40**	**38**	**92.5%**

All messages classified as questions, valid save requests, and other interactions were correctly identified, achieving an accuracy of 100% in these categories. The only classification errors occurred in the save request with invalid data category, where three requests were incorrectly classified as valid. This result highlights a limitation of the classifier in detecting insufficient or incomplete input. In these cases, the information is still stored in Notion. By providing users with a direct link to the newly created Notion entry, the system enables immediate user control, allowing the submitted content to be reviewed, edited, or deleted if necessary.

Following the developments presented since Chapter 1, Table 18 provides an updated overview of the system components and the technologies used for each functionality.

Table 18 – Fifth iteration: main components and technology updated

COMPONENT	DESCRIPTION	TECHNOLOGY
Communication Interface	The channel through which users interact with TekBot. In this project, Slack is used as the primary interface for submitting questions and adding new information.	Slack
AI Agent	The language model responsible for classifing user messages and generating natural language responses.	GPT-5 Nano
Knowledge Repository	A SQL database that stores company data	Supabase
Memory Module	Component that retains previous interactions for conversational context	n8n : Memory Node
Information Storage	Used to store new knowledge contributed by users	Notion

4.5.2. Demonstration and evaluation

The final iteration of the evaluation process involved a Senior Interactive Developer with over twelve years of experience in the organisation. Her role covers both front-end and back-end development, requiring frequent documentation of technical work and regular access to internal knowledge to support daily tasks. She also makes regular use of AI-powered chatbots to assist with programming-related activities, ensuring a high level of familiarity with conversational AI tools and workflows.

The results of the fifth iteration evaluation, presented in Table 19, indicate a good level of user satisfaction across all different dimensions.

Table 19 – Fifth iteration: interview results

ID	QUESTION	ANSWER
Q1	How easy was it to interact with TekBot when asking questions?	Easy
Q2	How easy was it to register new information?	Very Easy
Q3	Were the answers useful?	Very Satisfied
Q4	Were the answers accurate?	Very Satisfied
Q5	How satisfied are you with the response time?	Very Satisfied
Q6	Would you use TekBot as part of your daily work routine?	Yes
Q7	In which situations do you see yourself using TekBot?	For asking quick questions. Storing new information.
Q8	Do you have any suggestions for improving TekBot?	Would prefer to interact with the bot via direct messaging instead of in a channel. Bigger knowledge base.

The participant reported that interacting with TekBot to ask questions was easy and that registering new information was very easy, confirming the effectiveness of the interaction design. Both the usefulness and accuracy of the answers were rated as very satisfactory, as was the system's response time, suggesting that the changes introduced in this iteration addressed previous limitations.

The participant confirmed that TekBot would be used as part of the daily work routine, particularly for asking quick questions and storing new information. Suggested improvements focused primarily on usability preferences, namely the option to interact with the chatbot via direct messaging rather than a public channel, as well as the expansion of the knowledge base to further increase the system's value. The suggestion to expand the data available to TekBot was raised again and represents an important direction for future work. Similarly, the request to enable direct interaction with the chatbot via direct private messages, rather than through a public channel, was identified as a potential usability improvement that could be included in future improvements.

CHAPTER 5

Conclusions

5.1. Contributions

A central contribution of this work lies in the way the Design Science Research (DSR) approach structured the evolution of the artefact. The iterative process in which each version of TekBot was designed, implemented, evaluated with real users, and subsequently refined proved essential for aligning the system with actual organisational needs.

The adoption of an orchestration platform such as n8n enabled the system to evolve in a practical and iterative manner, supporting successive tests and architectural changes, including the evaluation of different AI models and data retrieval approaches. Its low-code nature and ability to integrate multiple platforms facilitated rapid experimentation with minimal development overhead, making it especially suitable for iterative DSRM and for prototyping AI-driven solutions in organisational environments.

The results obtained throughout the iterative evaluation cycles demonstrate that TekBot successfully met the core objectives defined at the beginning of this research. The system improves access to internal technical knowledge by delivering context-aware answers directly within the organisation's primary communication platform, while also providing a simple, clear, and efficient mechanism for storing new information. User feedback consistently highlighted the ease of interaction, the usefulness of the responses, and the practical advantages of remaining within the tools already being used in the company, in this case Slack, eliminating the need to consult external documentation and resources.

The initial use of unstructured data and embedding-based retrieval produced positive results for general information access; however, it revealed clear limitations when addressing queries that required data comparison, aggregation, or the generation of comprehensive lists. These limitations directly reinforce the findings of the literature review, which emphasise the critical role of data quality, structure, and curation in the effectiveness of AI-powered conversational agents. In this context, the iterative transition towards a structured data architecture and the adoption of SQL-based reasoning proved to be a decisive design choice. This evolution significantly improved response accuracy enabled reliable comparative queries and list generation, and resulted in more consistent and factual answers, while simultaneously reducing token consumption and operational costs.

The ability of the solution to capture and store new information received consistently positive feedback during the evaluation phases, with participants explicitly expressing their intention to use this functionality as a simpler and more immediate way to document technical knowledge. The recurrence of requests to further enhance this capability indicates a clear willingness to adopt the system as part of daily work practices.

The demonstrated ease of use, combined with users' motivation to contribute knowledge, highlights the practical value of this functionality as a mechanism for supporting continuous knowledge creation within the organisation. To ensure long-term effectiveness, it will be necessary to define clear rules and procedures for validating, structuring, and curating the collected information, allowing the captured knowledge to feed both the chatbot and the organisation's KM system.

Overall, the results confirm that the proposed solution contributes positively to productivity, knowledge reuse, and the reduction of information loss within a real organisational context.

5.2. Limitations

Despite the positive results achieved, there are some limitations on this work. The evaluation was conducted within a single organisation and with a relatively small group of users. TekBot's performance is constrained by the quality and completeness of the information provided to it; gaps, ambiguities, or outdated content in the underlying knowledge base can directly affect the accuracy of responses. The system relies on continuous curation and timely knowledge registration.

This study did not explore long-term adoption, cultural resistance, or the impact of knowledge governance policies, which represent important opportunities for future refinement and research.

5.3. Future work

A first area for future work is the expansion of the knowledge base and data integration. The current dataset remains limited in scale and scope. Future iterations should aim to integrate additional data sources, data formats, and systems already in use within the organisation such as Notion, internal databases, and other existing documentation. This integration would enable better access to organisational knowledge, ensuring that information remains continuously updated, useful, and relevant to users' needs.

At present, the system offers a fast and efficient way of storing information directly into a Notion database within the existing KM system. Moving forward, it will be crucial to establish clear procedures for validating, refining, and structuring this information for integration into the KM system and the AI system.

Currently, TekBot works as a reactive system that responds to user prompts by providing answers or by storing new information. Future versions could introduce proactive functionalities, such as storing recommendations for new information topics or updates when no direct response is available. Another possibility could be TekBot asking questions to users related to the information topics flagged as missing. This feature would significantly enhance TekBot's functionality, moving from just a passive assistant to an active contributor within the organisation's KM system.

In conclusion, this research establishes a solid foundation for continued exploration and enhancement of AI-driven KM. Expanding TekBot's capabilities through broader data integration, proactive behaviours, and domain-specific applications will further strengthen its potential as a strategic tool for organisational learning, collaboration, and digital transformation.

References

[1] N. M. Lisandi, R. E. Tarigan, C. A. Haryani, Hery, and A. E. Widjaja, “Enhancing Organizational Performance Through Strategic Knowledge Sharing: Insights and Models,” in 2024 2nd International Conference on Technology Innovation and Its Applications (ICTIIA), Sept. 2024, pp. 1–6. doi: 10.1109/ICTIIA61827.2024.10761330.

[2] F. Abousweilem, “Driving remote team success through knowledge management practices in the Jordanian high-tech industry,” Probl. Perspect. Manag., vol. 22, no. 2, pp. 708–720, July 2024, doi: 10.21511/ppm.22(2).2024.55.

[3] E. D. Canedo, A. T. S. Calazans, G. R. S. Silva, E. T. S. Masson, and I. S. Brito, “On the Challenges to Documenting Requirements in Agile Software Development: A Practitioners’ Perspective,” in Congresso Ibero-Americano em Engenharia de Software (CIbSE), SBC, May 2024, pp. 286–300. doi: 10.5753/cibse.2024.28454.

[4] T. Theunissen, U. van Heesch, and P. Avgeriou, “A mapping study on documentation in Continuous Software Development,” Inf. Softw. Technol., vol. 142, p. 106733, Feb. 2022, doi: 10.1016/j.infsof.2021.106733.

[5] “TELETRABALHO PÓS-PANDEMIA: IMPACTOS E PERSPECTIVAS – ISSN 1678-0817 Qualis B2,” Rev. Ft, Accessed: June 20, 2025. [Online]. Available: https://revistaft.com.br/teletrabalho-pos-pandemia-impactos-e-perspectivas/

[6] W. Ben Arfi, J.-M. Sahut, L. Hikkerova, and E. Braune, “Virtual teams and knowledge sharing via digital platforms: Evidence from an inter-organizational network context,” Manag. Int. Int. Manag. Gestiòn Int., vol. 27, no. 6, pp. 18–31, 2023, doi: 10.59876/a-n1ve-zm1e.

[7] L. Yang et al., “The effects of remote work on collaboration among information workers,” Nat. Hum. Behav., vol. 6, no. 1, pp. 43–54, Sept. 2021, doi: 10.1038/s41562-021-01196-4.

[8] S. Morrison-Smith and J. Ruiz, “Challenges and barriers in virtual teams: a literature review,” SN Appl. Sci., vol. 2, no. 6, p. 1096, May 2020, doi: 10.1007/s42452-020-2801-5.

[9] H.-C. Lee, K. Hung, G. M.-T. Man, R. Ho, and M. Leung, “Development of an RAG-Based LLM Chatbot for Enhancing Technical Support Service,” presented at the IEEE Region 10 Annual International Conference, Proceedings/TENCON, 2024, pp. 1080–1083. doi: 10.1109/TENCON61640.2024.10902801.

[10] M. Alavi, D. Leidner, and R. Mousavi, “Knowledge Management Perspective of Generative Artificial Intelligence (GenAI),” Apr. 03, 2024, Social Science Research Network, Rochester, NY: 4782875. doi: 10.2139/ssrn.4782875.

[11] S. M. Inavolu, Exploring AI-Driven Customer Service: Evolution, Architectures, Opportunities, Challenges and Future Directions. 2024. doi: 10.13140/RG.2.2.19937.31841.

[12] K. Chhabria and M. Damle, “Evolving Journey of Chatbots: Insights into Business Decisions,” in 2022 International Interdisciplinary Humanitarian Conference for Sustainability (IIHC), Nov. 2022, pp. 102–107. doi: 10.1109/IIHC55949.2022.10060780.

[13] A. A. Tishchenko, P. A. Tishchenko, and T. A. Shestakov, “Developing and Optimizing a Model for Creating a Learning Chatbot Using the Example of Technical Support,” presented at the IOP Conference Series: Earth and Environmental Science, 2021. doi: 10.1088/1755-1315/720/1/012093.

[14] J. Benita, R. M. R. Kumar, E. E, G. M. Reddy, and N. R. K, “Multi-Language Conversational Agent for Tech Support: Design and Implementation,” in 2024 International Conference on Artificial

Intelligence and Quantum Computation-Based Sensor Application (ICAIQSA), Dec. 2024, pp. 1–7. doi: 10.1109/ICAIQSA64000.2024.10882255.

[15] K. Peffers, T. Tuunanen, M. Rothenberger, and S. Chatterjee, “A design science research methodology for information systems research,” J. Manag. Inf. Syst., vol. 24, pp. 45–77, Jan. 2007.

[16] “Scopus preview - Scopus - Welcome to Scopus.” Accessed: May 31, 2025. [Online]. Available: https://www.scopus.com/home.uri

[17] “IEEE Xplore.” Accessed: May 31, 2025. [Online]. Available: https://ieeexplore.ieee.org/Xplore/home.jsp

[18] “Document Search - Web of Science Core Collection.” Accessed: May 31, 2025. [Online]. Available: https://www.webofscience.com/wos/woscc/basic-search

[19] “Zotero | Your personal research assistant.” Accessed: May 31, 2025. [Online]. Available: https://www.zotero.org/

[20] “Software de Folha de Cálculo Online Gratuito: Excel | Microsoft 365.” Accessed: May 31, 2025. [Online]. Available: https://www.microsoft.com/pt-pt/microsoft-365/excel?market=pt

[21] “ChatGPT.” Accessed: May 31, 2025. [Online]. Available: https://chatgpt.com

[22] “The AI for Academic Research | SciSpace.” Accessed: May 31, 2025. [Online]. Available: https://scispace.com

[23] M. Poser and E. A. C. Bittner, “(Re)Designing IT Support: How Embedded and Conversational AI Can Augment Technical Support Work,” presented at the 42nd International Conference on Information Systems, ICIS 2021 TREOs: “Building Sustainability and Resilience with IS: A Call for Action,” 2021. [Online]. Available: https://www.scopus.com/inward/record.uri?eid=2-s2.0-85152125922&partnerID=40&md5=94bf8e0d71210e8c881d4a5a78ffb6f7

[24] J. O. Torres, V. X. Barraza, and P. S. Castaneda, “Chatbot Based on Conversational Artificial Intelligence for the Attention of an IT Service Desk in Florida,” presented at the 2024 6th International Conference on Natural Language Processing, ICNLP 2024, 2024, pp. 775–782. doi: 10.1109/ICNLP60986.2024.10692529.

[25] D. Tipe-Palomino and W. Auccahuasi, “Development of Helpdesk Chatbot for Incident Classification and Resolution using NPL and Deep Learning,” presented at the 3rd International Conference on Automation, Computing and Renewable Systems, ICACRS 2024 - Proceedings, 2024, pp. 785–791. doi: 10.1109/ICACRS62842.2024.10841729.

[26] P. Reinhard, M. M. Li, C. Peters, and J. M. Leimeister, “Generative AI in Customer Support Services: A Framework for Augmenting the Routines of Frontline Service Employees,” presented at the Proceedings of the Annual Hawaii International Conference on System Sciences, 2024, pp. 237–246. [Online]. Available: https://www.scopus.com/inward/record.uri?eid=2-s2.0-85195312271&partnerID=40&md5=55e93420aca04dd9ff29d28632b95d0e

[27] V. Scotti and M. J. Carman, “LLM Support for Real-Time Technical Assistance,” presented at the Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 2024, pp. 388–393. doi: 10.1007/978-3-031-70371-3_26.

[28] Z. Xu et al., “Retrieval-Augmented Generation with Knowledge Graphs for Customer Service Question Answering,” presented at the PROCEEDINGS OF THE 47TH INTERNATIONAL ACM SIGIR CONFERENCE ON RESEARCH AND DEVELOPMENT IN INFORMATION RETRIEVAL, SIGIR 2024, 2024, pp. 2905–2909. doi: 10.1145/3626772.3661370.

[29] R. Vasaniya, M. Visodiya, and A. K. Patel, “TechAssist: A RAG-Based Chatbot for Accessing Technical Information from StackOverflow,” presented at the 2025 IEEE International Students’ Conference on Electrical, Electronics and Computer Science, SCEECS 2025, 2025. doi: 10.1109/SCEECS64059.2025.10940184.
[30] “Dialogflow.” Accessed: July 12, 2025. [Online]. Available: https://dialogflow.cloud.google.com/#/login