

CHOIR: A Chatbot-mediated Organizational Memory Leveraging Communication in University Research Labs

Sangwook Lee
Computer Science
Virginia Tech
Blacksburg, Virginia, USA
sangwooklee@vt.edu

Adnan Abbas
Computer Science
Virginia Tech
Blacksburg, Virginia, USA
adnana99@vt.edu

Yan Chen
Computer Science
Virginia Tech
Blacksburg, Virginia, USA
ych@vt.edu

Young-Ho Kim
NAVER AI Lab
Seongnam, Republic of Korea
yghokim@younghokim.net

Sang Won Lee
Computer Science
Virginia Tech
Blacksburg, Virginia, USA
NAVER AI Lab
Seongnam, Republic of Korea
sangwonlee@vt.edu

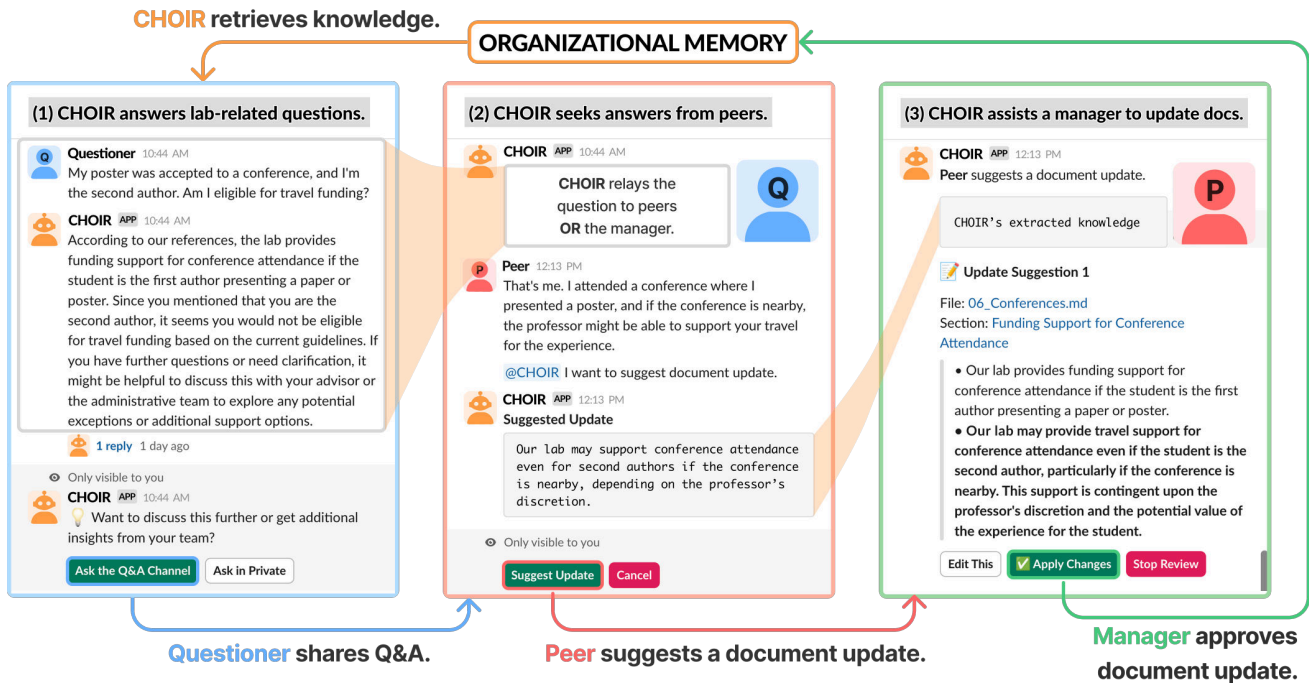


Figure 1: CHOIR's workflow across different communication channels

Abstract

University research labs often rely on chat-based platforms for communication and project management, where valuable knowledge surfaces but is easily lost in message streams. Documentation

can preserve knowledge, but it requires ongoing maintenance and is challenging to navigate. Drawing on formative interviews that revealed organizational memory challenges in labs, we designed CHOIR, an LLM-based chatbot that supports organizational memory through four key functions: document-grounded Q&A, Q&A sharing for follow-up discussion, knowledge extraction from conversations, and AI-assisted document updates. We deployed CHOIR in four research labs for one month (n=21), where the lab members asked 107 questions and lab directors updated documents 38 times in the organizational memory. Our findings reveal a privacy-awareness tension: questions were asked privately, limiting



directors' visibility into documentation gaps. Students often avoided contribution due to challenges in generalizing personal experiences into universal documentation. We contribute design implications for privacy-preserving awareness and supporting context-specific knowledge documentation.

CCS Concepts

• **Human-centered computing** → **Collaborative and social computing systems and tools.**

Keywords

chatbot, group chat, organizational memory, LLM, living document

ACM Reference Format:

Sangwook Lee, Adnan Abbas, Yan Chen, Young-Ho Kim, and Sang Won Lee. 2026. CHOIR: A Chatbot-mediated Organizational Memory Leveraging Communication in University Research Labs. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 23 pages. <https://doi.org/10.1145/3772318.3791314>

1 Introduction

Modern university research labs increasingly rely on workplace instant-messaging platforms such as Slack and Microsoft Teams to facilitate communication and collaboration among lab members, including those working remotely [79, 101]. Slack alone reported over 1.2 million users across more than 3,000 higher education institutions during the COVID-19 pandemic in 2020 [90]. These platforms serve as channels through which lab members share knowledge crucial and practical to their work—e.g., academic writing tips, internship opportunities, and travel grant information—on an ongoing basis [9, 51, 101]. However, this knowledge is often lost. Messages are stored chronologically and without semantic structure, making it difficult to retrieve relevant information once it is buried in long conversation threads [71, 78, 104, 105]. As a result, lab directors may receive the same questions repeatedly or give the same advice repeatedly, e.g., correcting punctuation in latex. Students are left to learn through trial and error [10].

Prior literature has suggested addressing this challenge through sociotechnical interventions that establish an *organizational memory*, which documents and preserves knowledge accumulated by an organization over time [4, 99]. Some research labs implement such organizational knowledge systems using collaborative documents such as Google Docs [94] or group wikis [47, 98], often in the form of lab handbooks or policy documents. However, maintaining these resources requires substantial manual effort [50, 58, 66, 77]. In practice, only professors find it difficult to consistently maintain such documents [92, 94]—a finding also confirmed by our formative study. As a result, much of the knowledge remains undocumented, residing instead in the memory of professors and lab members [34, 41]. Even when central documents exist, students often struggle to read them entirely or locate the right information when they need it [17, 102]. Consequently, they must turn to peers or professors, making knowledge transfer inefficient.

To address this challenge, we present CHOIR (Chat-based Helper for Organizational Intelligence Repository), a chatbot-based organizational memory management system embedded in Slack.

The primary goal of CHOIR is twofold: (1) to build and sustain a research lab's organizational memory, and (2) to enable lab members to easily access knowledge from it. To that end, our approach embeds organizational memory directly within their communication platform, such as Slack, allowing students to ask questions to a chatbot and allowing professors to modify documents within the platform. Specifically, CHOIR integrates four interconnected interactions: **Document-grounded Q&A**, which retrieves knowledge with relevant references from a collection of documents when a user asks a question; **Q&A Sharing for Follow-up Discussion**, which allows users to share Q&A with peers to address unanswered questions or extend discussions; **Knowledge Extraction from Conversation**, which captures new knowledge that naturally emerges from Slack messages; and **AI-assisted Document Update**, which facilitates document maintenance by helping managers locate relevant documents and suggest edits. In this way, CHOIR supports an organizational memory practice that leverages lab members' communication in Slack to retrieve and sustain knowledge, reducing lab directors' burden of continuously updating documents while providing members with immediate and convenient access to organizational memory. While this approach can be potentially applicable to many different types of organizations, such as an online community and a company, we chose research labs as our target environment because they constantly experience member turnover, rely on regular communication, and maintain various types of documentation that require collaborative updates [97].

To examine both current challenges of organizational memory for university research labs and the opportunities of a chatbot-driven organizational memory, we aim to answer the following two questions:

- **RQ1:** What are the current challenges of managing and utilizing organizational memory in research labs?
- **RQ2:** How do lab members engage with and experience a chatbot-driven system for managing and utilizing organizational memory in their online communication?

We followed a three-stage approach: (1) a formative study with faculty and graduate students (N=15) to understand their organizational memory practice, if any, challenges in transferring and retrieving knowledge among lab members; (2) the design and implementation of CHOIR guided by these findings; and (3) a one-month field deployment (N=21) across four university research labs, preceded by document collection and structured onboarding. Based on our interaction trace analysis and follow-up interviews after the field study, we found that CHOIR helped students obtain answers available in the documents. In addition, it supported professors with no prior experience in maintaining organizational memory to document new knowledge that students wanted to learn or needed to know when questions were shared publicly. However, our findings revealed a fundamental tension between students' privacy preferences and directors' need for awareness of knowledge gaps. Students faced psychological barriers to sharing their questions—either from concerns about bothering their peers and professors, or from reluctance due to how others might perceive them for not knowing the answers to seemingly basic questions. Similar barriers were noted for answering other people's questions and for contributing to documentation updates. Based

on the findings, we discuss the design implications for AI-mediated organizational memory systems, including the privacy-awareness trade-off, the role of AI as a facilitator for human communication, and approaches to expand the perceived boundary of documentable knowledge.

The research contributions of this paper include:

- (1) An empirical investigation of the challenges surrounding organizational memory practices in university research labs.
- (2) A socio-technical workflow that integrates conversational Q&A, knowledge sharing, and AI-assisted documentation to support organizational memory.
- (3) Findings from a month-long deployment that reveal real-world usage patterns, challenges and design implications for AI-mediated organizational memory systems.

2 Related Work

While prior work on personal knowledge management [20, 35, 44] focused on search-based retrieval across collected sources, an organization's knowledge management is challenging due to its scale and group dynamics. We explore prior research on organizational memory and knowledge management systems, examining how digital platforms have evolved to support knowledge preservation and sharing. We then investigate approaches to transforming conversational data into organizational knowledge, and finally review the emerging role of chatbots and large language models (LLMs) in facilitating knowledge work.

2.1 From Repository Models to Distributed Organizational Memory

Organizational memory (OM), defined as stored information from an organization's history that can be brought to bear on present decisions [99], has been a central concern in knowledge management research. Early approaches assumed that organizational knowledge could be systematically captured, stored, and retrieved through repository-based systems. Initial groupware systems attempted to implement this vision by extracting knowledge from individuals and storing it in centralized databases. However, deployments revealed fundamental social and organizational barriers that overshadowed technical challenges. Research identified that knowledge sharing requires addressing incentive structures [77, 92], organizational power dynamics [46], and the scattered nature of informal knowledge that often remains undocumented [28]. Competitive organizational cultures frequently inhibited knowledge sharing, as employees treated expertise as personal assets.

Meanwhile, a different approach emerged with Ackerman's Answer Garden series, which explored hybrid human-system methods. Answer Garden [5] allowed users to navigate branching questions and automatically routed unanswered queries to experts, who could then contribute responses to the knowledge base. Answer Garden 2 [6] introduced cascading routing from peers to specialists to experts. However, field studies revealed persistent challenges [1]: mismatched expectations between experts and users regarding answer detail, user reluctance to burden experts, and, in industrial deployments, conflicts between organizational silos and departmental politics [80]. These findings led to reconceptualizing

OM as a distributed cognitive system rather than a centralized repository. Ackerman [3] demonstrated through ethnographic research that OM encompasses individual memory, peer knowledge, documents, and customer interactions in interconnected forms. Knowledge crossing organizational boundaries transforms into boundary objects [91], requiring recontextualization for effective use [4].

With the advancement of web technologies, wiki systems demonstrated this potential through collaborative editing [98], but organizational deployments encountered familiar social barriers. Danis [31] deployed ResearchWiki in a research organization but found limited participation due to competitive culture and modification anxiety. Similarly, Munson [73] found in a MediaWiki deployment that while the system motivated documentation of past projects and enabled discovery of organizational work, members felt uncomfortable making substantial edits to others' content. Broader studies of enterprise wikis found that making workers participate in wiki is challenging and organization's culture that support collaboration is necessary [57]. Subsequent research identified systemic barriers: additional effort requirements, information sensitivity concerns, discomfort with unfinished work publication, and openness anxiety [50]. The collaborative editing model effective in public contexts created stress in organizational environments where authorship and accountability mattered [47]. Alternative platforms like Google Docs achieved better adoption through simplified user experiences that naturally facilitated role differentiation without wiki complexity [94], with supporting technologies focusing on transparent collaboration history to navigate social complexities [100].

2.2 Leveraging Conversational Platforms for Knowledge Management

Researchers have increasingly recognized conversations as valuable sources of organizational knowledge, particularly for capturing tacit knowledge that resists formalization [82]. The challenge is transforming conversational knowledge into explicit, reusable forms while preserving contextual richness [76]. Early digital communication platforms revealed this potential, with email evolving beyond simple messaging to serve multiple organizational functions, including task management, document archiving, and conversation tracking, though leading to information overload [103]. Instant messaging further demonstrated how informal channels could become knowledge repositories for relationship maintenance and work coordination [49, 51, 74].

As diverse platforms began supporting interpersonal conversations, researchers discovered that various platforms were serving knowledge management roles. Platform diversification revealed knowledge management opportunities in tools not originally designed for that purpose. Software issue trackers evolved into OM and coordination systems [16], while corporate social networking enabled help requests and work coordination through informal interactions [36]. Research showed these platforms improved employees' metaknowledge awareness through social mediation [65, 70, 75]. Team communication platforms like Slack provided unprecedented conversation visibility and multi-threaded collaboration [9], but also introduced information

overload and conversational inefficiencies from simultaneous topic progression [78]. In research environments, Slack served as both a knowledge-sharing hub and a culture formation space [79], with conversational data proving valuable for tasks like API recommendations from developer discussions [23].

Despite these discoveries, converting conversations into persistent knowledge required addressing the challenge of preserving natural communication flow while enabling knowledge extraction. Early approaches focused on visualization and threading for improved retrieval [32, 37, 54], but required manual navigation. More sophisticated systems automated knowledge extraction through annotation [7, 39] and collaborative summarization [105]. Chat-based knowledge management systems achieved greater sophistication with Tilda, which automatically summarizes tagged conversations [104], and tools combining topic learning with personalized summaries [96]. Social Q&A systems demonstrated how questions can surface knowledge gaps and how answers gradually accumulate into shared resources [84]. QWiki demonstrated that Q&A and documentation approaches could be mutually reinforcing in MOOC environments [88]. This trajectory confirms Ackerman's [2] prediction that OM should move toward communication-centric, people-oriented approaches, while raising questions about scalability and sustainability.

2.3 AI-Powered Knowledge Management: From Chatbots to LLMs

These scalability and sustainability challenges led to the emergence of AI-powered approaches to address the limitations of manual conversational knowledge extraction. As demonstrated by early systems like Collabot [96] and Tilda [104], many of these solutions took the form of task-based chatbots [8]. These systems provided basic topic extraction and unidirectional summarization capabilities. Chatbots evolved from passive observers to active conversation participants, taking specialized roles as archivists, social organizers, and moderators [55, 56, 87]. Research showed users felt more comfortable sharing knowledge with chatbots than humans, reducing social barriers [64].

With the advancement of chatbot technologies, LLMs enabled sophisticated conversational knowledge processing beyond basic pattern recognition. LLMs demonstrated superior performance in email and meeting summarization [106, 108] and enabled interactive refinement where users could update AI-generated summaries for improved performance [12]. Retrieval-Augmented Generation (RAG) addressed the challenge of grounding conversational AI in organizational knowledge bases [67], with RAG systems searching existing documents and incorporating relevant information into responses, improving accuracy and reducing hallucination [42]. Beyond retrieval and summarization, LLMs enable dynamic document updating based on conversational knowledge [61] and facilitate complex group interactions [24].

2.4 Summary and Connection to CHOIR

This review reveals consistent evolution in OM systems: from repository models to distributed cognitive systems, from formal documentation to conversational capture, and from manual curation to AI-assisted transformation. However, persistent

challenges have remained consistent across decades— incentive misalignment, expert burden, participation barriers, and information overload—indicating that technical advances alone are insufficient. Prior AI-based systems such as Tilda and Collabot primarily address subprocesses of the OM lifecycle, mainly capturing and summarizing conversation [96, 104], but do not support the full cycle: how extracted knowledge is validated, integrated into OM documents, and easily retrieved by the members. CHOIR addresses this gap by synthesizing three research streams: embedding OM in existing social systems, leveraging existing conversational interactions, and exploiting LLMs' capabilities. By operating within Slack where research teams already communicate, CHOIR avoids traditional adoption barriers while implementing document-grounded Q&A, conversational knowledge extraction, and AI-assisted documentation workflows. The system design incorporates lessons from prior research: mutually reinforcing Q&A and documentation cycles, integration with existing communication flows, and AI assistance that reduces rather than increases user burden. CHOIR thus represents both technological synthesis and an approach to addressing the persistent social challenges that have limited OM system adoption in collaborative research environments.

3 Formative Study

To address RQ1, i.e., understanding challenges of organizational memory in research labs, we conducted a formative study with university research lab members, both professors and students.

3.1 Method

We recruited participants through a screening survey distributed to faculty and graduate students at research-oriented universities. We advertised the study on the author's social media to recruit professors in other research-oriented universities. We recruited 15 participants: 11 faculty members (P1-P11) from various academic disciplines, primarily Computer Science (7) but also Industrial Engineering (2), Information Science (1), and Human Development and Family Science (1), and 4 graduate students (S1-S4) from different disciplines. We targeted professors who directed research labs with more than 4 Ph.D. students and had experience or interest in maintaining organizational memory. Faculty participants had been running their research labs for an average of 6.4 years (SD = 3.78, range: 4-13 years) and supervised an average of 5 Ph.D. students (SD = 1.67, range: 2-7 students). Some of the participants had a well-established organizational memory practice, whereas others relied on more informal organizational methods. This study was approved by the Virginia Tech Institutional Review Board (IRB 25-247).

We conducted semi-structured interviews lasting 45-60 minutes with each participant via Zoom. All interviews were audio-recorded with the participants' consent. The interview protocol explored the challenges that they face in addressing knowledge transfer and maintaining organizational memory, if any. We also presented early mock-up of CHOIR to elicit participants' reactions to possible features, rather than a fully implemented system. Participants were compensated with \$50 for faculty members and \$20 for graduate students upon completion.

Each interview was audio-recorded, and the first two authors transcribed the interview recordings, manually correcting the output of a transcription engine (Sonix AI¹). After familiarization with the data, the second author performed a thematic analysis [27], and initially generated 214 codes. The three authors conducted axial coding by revising, merging, and deleting codes and identifying emergent themes.

3.2 Result

Overall, both professors and students quickly recognized the benefits of documenting an organizational memory to help students access information. According to the professors interviewed, they often have to repeatedly address the same set of questions. These questions not only include administrative procedures like travel reimbursements or the Ph.D. qualification process but also involve insights or organizational memory gained over time, such as writing tips, information about equipment, research internships, and student expectations.

“We have all kinds of gadgets, devices, and equipment. We have many robots, driving simulators, laptops, and eye trackers. And so they ask the same questions again and again.” (P6)

“Even if the questions are slightly different, I think the answers could be the same. [...] And so I sort of find myself having to repeat a lot of these things” (P11)

When professors lack formal documentation, they often rely on senior students to address recurring questions and facilitate knowledge transfer. However, they reported that, as senior students typically graduate after completing their degree, ensuring continuity in knowledge transfer becomes a separate challenge if not documented.

Of the 11 professors we interviewed, 7 maintained or had a document or a collection of documents that allowed students to find answers independently. The level of detail in these documents also varied considerably, and most of them mentioned that they were not comprehensive or up-to-date. Some labs provided a two-page onboarding syllabus (P7), while others offered extensive and comprehensive documentation spanning over 80 pages (P8)

They also recognized that all the advice or tips they give in person is “ephemeral” and emphasized the benefits of documentation due to its persistent nature, allowing reuse and future reference. However, our participants reported that utilizing and maintaining such documents was both challenging and concerning, which we will outline below.

3.2.1 Challenges in Creating and Maintaining Lab Documentation. The primary challenge in documenting lab knowledge was the time-intensive nature of the task. Many participants admitted they were too “busy” or “lazy” to create such documents. In practice, documentation often takes a lower priority than more pressing responsibilities, such as research and teaching.

“I think it’s nice to have. I mean, I’m just lazy, and I don’t have the time and energy, and it’s just like, I’m not a type of person who could do that, but I think it’s a nice resource to have.” (P3)

“Documentation is usually one of those things that gets less attention, that gets deprioritized. [...] I have, I think, made 1 or 2 attempts over the years, but they never got fleshed out to the point where they would really be useful and worth sharing.” (P9)

While participants were concerned about the sheer amount of content they had to generate, keeping the document up to date with dynamic information was another significant challenge. As the document grows, identifying what needs to be updated and locating outdated information becomes increasingly difficult.

“I go back to it [the lab handbook] and I look at it, I’m like, ‘Oh, wow. That’s like no longer true.’ [...] The more content I add, the more I have to maintain.” (P8)

“There was one time there was like a whole step just completely missing, and I couldn’t get it to work. [...] The people who had written the entry into the Notion had already graduated.” (S4)

These challenges are consistent with prior work on knowledge management; studies have shown the high cost of creating and maintaining knowledge artifacts [2, 77]. Researchers have noted that documents tend to become outdated quickly, reducing their usefulness over time [46, 66].

Based on these findings, we aim to design a system that facilitates the documentation of organizational memory. Our approach leverages frequently used communication platforms, such as Slack or Microsoft Teams, where lab directors or senior students answer questions and share tips. We can effectively capture and preserve various types of knowledge for documentation. The following design goals are derived from these findings.

Design Goal 1 (DG1): Facilitate the process of documenting organizational memory, leveraging the existing conversation among lab members.

3.2.2 Need for Effective Knowledge Retrieval. Another significant challenge raised by participants was the retrieval of knowledge. Even professors who maintained centralized documentation often struggled to ensure that students knew what information was available, where they are, and how to access it.

“Sometimes they will email me and ask me where the link is for the survey to report their hours [...] which for me is a little frustrating because it’s linked in the syllabus for the lab.” (P7)

“There is [sic] a lot of safety protocols that we need to go through. So if they don’t pay too much attention during those [onboarding] meetings, then usually they struggle later on to find those resources or do things properly.” (S3)

For the undocumented knowledge, similar challenges persist, if not worsen, as they often need to locate files, links, emails, or Slack messages to share with students. While they actively used search features, these tools were limited as information was scattered across multiple locations.

“Sometimes I’m frustrated that I have to go dig [the information] out and didn’t have it nicely organized in one spot.” (P5)

¹<https://sonix.ai/>

Communication platforms' features, such as search function or pinned messages, provide a way to reuse existing messages to some extent. However, they reported difficulties using those features for knowledge retrieval and management.

"I probably don't use those [search functions] as effectively as I could. So, there's definitely been times when I'm trying to find an old conversation and I'm searching within a DM from maybe one student and it's not coming up, and I'm not really sure how I can find it." (P9)

"Pinning a message[...] it can get quite a large, like, a long list of items. And like the pinned message, you can only see, like, a couple." (P1)

While existing functions allow users to retrieve or curate past messages they authored, this approach does not scale as the communication history grows, and retrieving information through pinned messages is similarly unsustainable.

Building on these findings, we propose designing a conversational agent that allows students to ask questions directly. The chatbot will respond using information available in relevant documents. This chat-based assistance will delegate the *information search*, greatly reducing the effort required to retrieve information. Additionally, we design the system to guide students to the original document and the specific section from which the answer was sourced, encouraging them to read the source directly. Providing a reference to the original document will help them understand the context of the answer and become familiar with the document. We summarize this approach in the following design goal.

Design Goal 2 (DG2): Allow students to easily retrieve correct information and find the relevant document on demand via conversational interaction.

3.2.3 Importance of Interpersonal Communication. One notable finding from the formative study was that, although professors recognized the value of an organizational memory document in reducing the burden of answering questions, they still preferred to communicate directly with students and respond to questions for educational and mentoring purposes. In particular, when the questions concerned core skills they aimed to develop, professors expressed a strong desire to answer them directly, even if the questions were repetitive.

"I do worry sometimes that some things I do prefer to repeat in person, even if it's painful for me. You know, it takes time, etc., right? Because there's some weights carried with delivering it there and saying the same spiel over and over again. Sometimes it's a good wake-up call for people to realize it's important." (P11)

Furthermore, a few professors expressed concern that a chatbot might reduce mentoring opportunities if it answered all questions, thereby removing the need for students to communicate with their advisors. One professor, P3, who opposed the idea of organizational memory in general, explained why professors value direct communication with students when we presented multiple scenarios where a student asked a chatbot a question "How many hours is a student expected to work per week?"

"I actually don't think I need it, and I actually think this can be counterproductive, honestly. The reason being that, like for this kind of question, I really see it as an opportunity to have a conversation with a student personally because I want to, I don't think it's the number of hours that matter, but I really want to figure out why the student is asking that question." (P3)

Similarly, some professors also felt that delegating questions to senior students, even though "senior students will find it annoying (P1)," can provide valuable opportunities to build a sense of community and offer senior students mentoring experience.

These findings resonate with prior research, which has shown that organizational memory systems alone are insufficient without the social and contextual practices that make knowledge useful [4, 81]. Education research similarly highlights that repeated questions, while seemingly inefficient, play a key role in mentoring and learning by fostering trust, socio-emotional support, and community building [59, 62, 86].

Having organizational memory and fostering interpersonal communication in the lab are two seemingly conflicting goals, emphasizing efficiency and education, respectively. To address the professors' concerns, we design the system to actively promote communication between students and professors as well as among lab members. We summarize this design goal as follows:

Design Goal 3 (DG3): Facilitate interpersonal communication and integrate the retrieval and update of organizational memory within the context of communication.

3.2.4 Documentation as a Collective Effort with Barriers. Professors expressed that organizational memory should be a collective effort, rather than the sole responsibility of lab directors. This view was driven not only by the potential to reduce the effort required to answer questions, but also by the recognition that some students are better positioned to explain certain types of knowledge. Such knowledge includes hands-on experiences that professors may not be familiar with or processes that only students go through (e.g., submitting documents for Ph.D. requirements).

"If it's something like, how do I request reimbursement for conference travel? You know, it's not going to be more helpful for me to explain that than if my senior student explains it, so I might delegate something like that" (P9)

In addition, those who maintained some form of organizational memory wanted to encourage collaborative editing by giving students a certain level of write access, allowing them to contribute by adding and updating content in the documents.

"You know you have the login, go and if there's something wrong, or if your internship was different and you find a new rule, it's a Wiki that they can update to. [...] You're part of this lab. Go update the information if it's out of date. I find that to also be kind of a way to make them read it if they feel like they own it as well." (P11)

In both cases, they noted that students' *ownership* of the document could be beneficial, both in encouraging them to read it and in fostering a sense of contribution or achievement. However, students

expressed reluctance toward documentation work, viewing it as tedious.

“The documentation part is always tedious. Nobody wants to do that. [...] There are times when you feel like, okay, this could be automated or this could be a dedicated job for an AI agent or something like that.” (S2)

Still, professors wanted to track all changes to ensure that edits made by students were legitimate. Ultimately, professors wished to remain aware of the changes and retain the final decision on whether to approve them.

“I would probably give the students like suggest mode access where they can annotate the document and add a comment or add a suggested edit, and then I would be the one to actually approve it. It wouldn't become live until I made those approvals. And that way I could maintain the accuracy and quality of the documents, and I could be aware of any changes that are happening because I would have to approve them first.” (P9)

Overall, professors viewed it positively when students developed a sense of ownership and contributed to building organizational memory collectively. However, students expressed reluctance toward such documentation work, highlighting a tension between professors' expectations and students' actual willingness to participate.

Prior work has shown that granting editing privileges and fostering a sense of community ownership can encourage sustained contributions and effective use of shared documents [40, 72, 98]. At the same time, professors emphasized the importance of tracking all changes and retaining final approval rights to ensure accuracy, reflecting a recurring tension also observed in large-scale collaborative systems. In Wikipedia, for instance, the balance between openness and quality control has been identified as a central challenge [11, 43]. Building on these insights, our design aims to encourage lab members' participation in curating organizational memory while ensuring that professors can maintain oversight and quality control.

Design Goal 4 (DG4): Support members' participation in curating organizational memory and ensure accuracy.

4 CHOIR: Chat-based Helper for Organizational Intelligence Repository

Grounded in our formative study, we designed CHOIR to support students to obtain answers directly within Slack and to transform everyday conversations into a maintained organizational memory. CHOIR's primary objective of sustaining the growth of—and easy access to—organizational memory translates into several subgoals derived from our formative study. These subgoals include streamlining document updates (DG1), enabling easy and accurate information retrieval (DG2), facilitating communication among members (DG3), and supporting their participation in managing organizational memory (DG4). We accomplish these goals by embedding CHOIR as a chatbot within the lab's communication platform, i.e., Slack. CHOIR integrates four features for the subgoals: **Document-grounded Q&A**, **Q&A Sharing for Follow-up**

Discussion, **Knowledge Extraction from Conversation**, and **AI-assisted Document Update**. These features reflect the design goals identified in the formative study.

When a message is sent to CHOIR—either through a direct message or by mentioning its handle (i.e., @CHOIR) in any channel that it is part of, CHOIR processes it as either a question or an update request, corresponding to the two start nodes shown in Figure 2. The four features in CHOIR can support a feedback loop of how a question asked by a student can eventually update the lab handbook. For example, the member who posed a question may decide to share the question with their peers, and the peers can answer the question, suggesting updating the document. Finally, the manager will decide whether to apply updates. The flow may also begin outside of Q&A, when a member or the manager directly provides new knowledge to CHOIR. CHOIR's four features will be detailed in the remainder of this section, showing how each aligns with a design goal (DG1–DG4).

4.1 Document-grounded Q&A with CHOIR

CHOIR allows lab members to ask a question and retrieve information from the documentation in any Slack channel. When a question is submitted, CHOIR answers by composing a response grounded in those sources (DG2). To avoid hallucination, CHOIR includes an *answerability guard*: when the repository does not provide sufficient coverage, CHOIR abstains from inventing an answer, surfaces related passages with citations, and clarifies what information is available (e.g., related sections or partial guidance). This makes documentation gaps visible for further update, and keeps replies aligned with current documents.

Each Q&A exchange is posted with grounding and provenance so that the questioner can quickly verify sources. Figure 3 highlights the UI elements: the member's question (A), CHOIR's answer (B), an evidence panel listing the specific chunks used with anchored links (C), and a share banner (D) for further discussion. These references encourage the questioner to check out the original documents, thereby getting familiar with the organizational memory.

4.2 Q&A Sharing for Follow-up Discussion

After CHOIR provides an answer, the questioner can share the Q&A exchange with others to facilitate communication among members (DG3). The share banner (D in Figure 3) offers two options: posting to a public Q&A channel or sending via direct message to selected members. Questioners can configure additional options through one of two modal windows shown in Figure 4. The *Your Comment* field (A in Figure 4) allows the questioner to add context or a follow-up prompt.

For the “Ask in Private” option, the question asker can choose lab members that they will share the question with using a *People picker* (D in Figure 4), which is pre-populated with the manager by default but can be edited to include any Slack members (e.g., senior Ph.D. students). *Privacy Options* (B) enable anonymous sharing, in which case CHOIR will replace the questioner's name with “A team member” when the message is shared in a Q&A channel to stay anonymous. In addition, when anonymity is enabled in the private-sharing modal, the questioner is also excluded from the

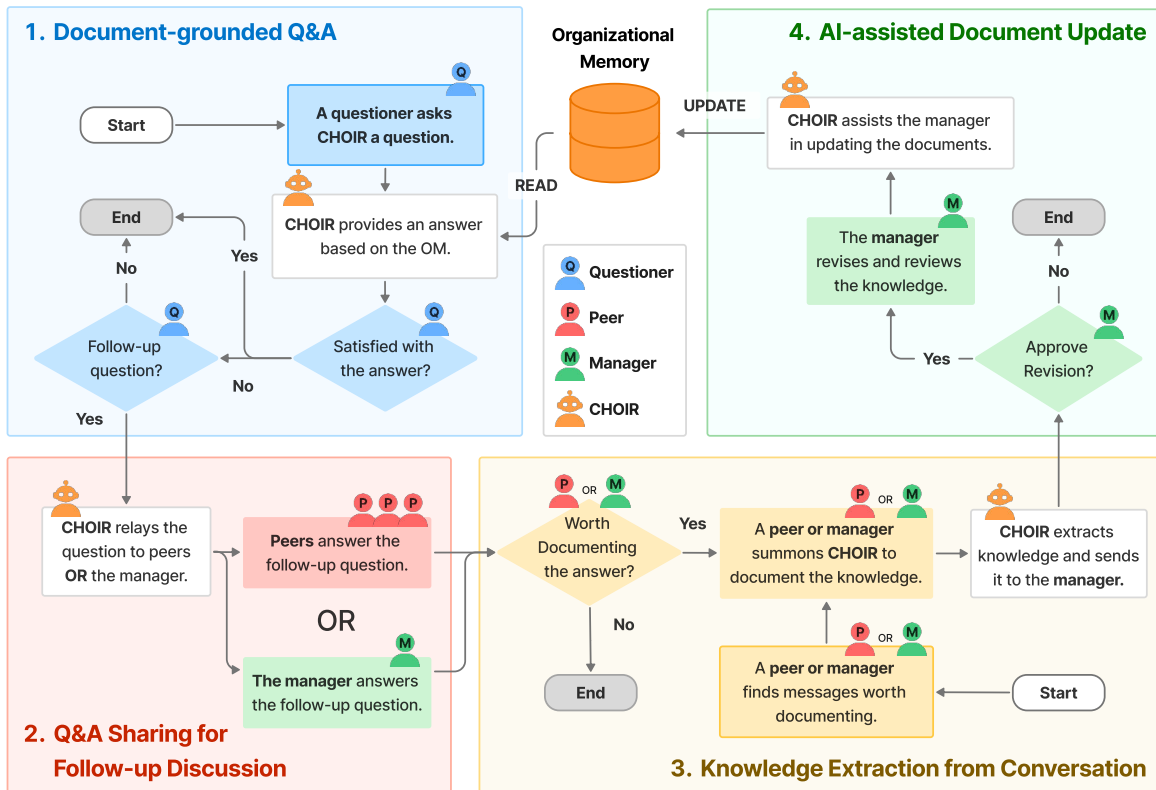


Figure 2: Overview of CHOIR’s workflow. A questioner may ask CHOIR and receive a document-grounded answer, optionally share the Q&A exchange for broader discussion, and from these interactions, knowledge can be extracted and reviewed by the manager. Approved updates are written back to the organizational memory repository, which also supplies sources for Q&A.

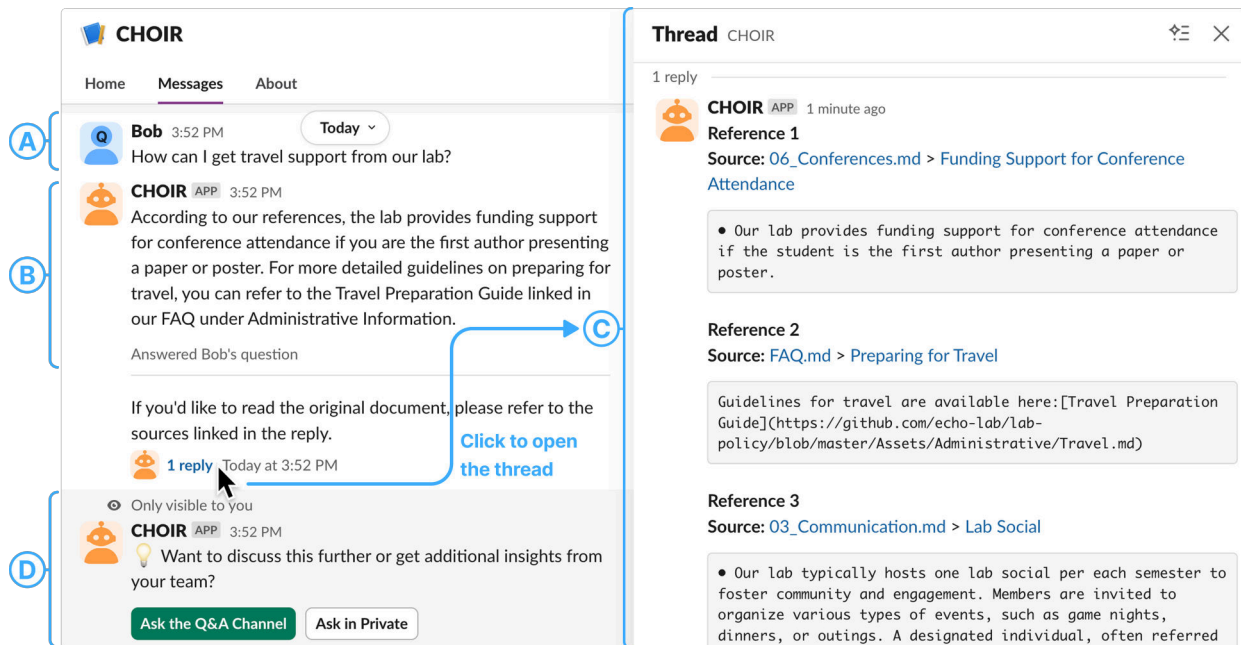


Figure 3: Document-grounded Q&A with CHOIR (DM example). Bob (questioner) asks CHOIR a question via direct message, CHOIR returns a document-grounded answer, and the thread is opened to view cited references.

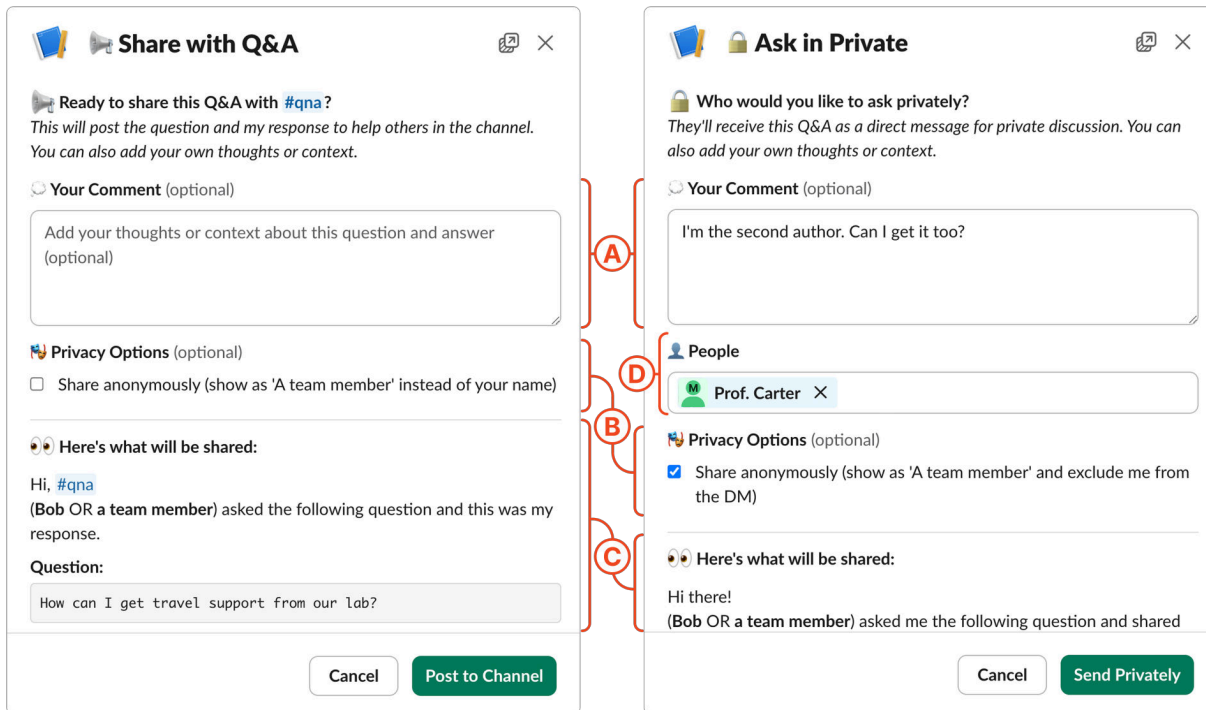


Figure 4: Two share modals in CHOIR. Left: "Share with Q&A" for posting a Q&A exchange to the designated channel. Right: "Ask in Private" for sending the same exchange via DM; in this example, the questioner adds a follow-up, the manager is pre-selected in the People picker, and anonymous sharing is enabled.

group chat conversation to preserve anonymity among recipients. The preview panel (C) shows exactly how the question will be shared, either as a channel post or as an outgoing DM.

The shared Q&A exchange includes the question, CHOIR's answer, the questioner's comment, and profile image as shown in Figure 5 (A), where other members can add follow-up comments (B). When shared anonymously in a private DM, CHOIR forwards replies to the original questioner without revealing their identity. When such discussions lead to new knowledge, they provide the entry point for *Knowledge Extraction from Conversation*.

4.3 Knowledge Extraction from Conversation

CHOIR supports turning any Slack messages into reusable knowledge (DG1), allowing members to participate in organizational memory maintenance (DG4). Follow-up discussions after Q&A sharing often contain answers worth documenting. When any lab member sees new knowledge worth keeping, they can summon CHOIR by mentioning @CHOIR (C) in Figure 5). CHOIR then examines the conversation, identifies *documentable* content, and drafts a suggested update (D). The requester can edit this draft and press *Suggest Update* to forward it to the manager (E). CHOIR can be summoned anywhere in Slack—for example, if a professor shares a writing tip while giving feedback to a student, the professor can summon CHOIR to document that knowledge, as represented by the new start node in the bottom right corner of Figure 2.

4.4 AI-assisted Document Update

This feature of AI-assisted update supports DG1 and DG4 in multiple ways. Once any member suggests an update from the Knowledge Extraction feature by pressing the "Suggest Update" button (E) in Figure 5), CHOIR delivers it to the manager as a Slack DM message (A) in Figure 6) from CHOIR. The message presents the proposed update together with the contributor's name and profile image, making the source and context explicit. At this stage, the manager can edit the suggestion, start the update process, or decline it, and these choices are visible not only in the DM but also as a message in the original channel so that other members are kept updated on how the update request is in progress. This final approval step, which the manager can only handle, is designed to help maintain the accuracy of documented information, which is relevant to DG4.

Once the manager starts the update process by pressing the "Start Update Process" button, CHOIR will send follow-up messages one step at a time, guiding the manager to access which file to update, and how. First, the manager can select a target file from the repository or create a new file if the knowledge does not fit into an existing one (B) in Figure 6). If not selected, CHOIR will use the most relevant file depending on the similarity of the suggested update and document content. CHOIR then shows which file is selected and the proposed change in a diff-style message with inline emphasis on additions and removals. For each suggestion, the manager may edit the content directly, apply the change to the document, skip it to show the next relevant document — as there

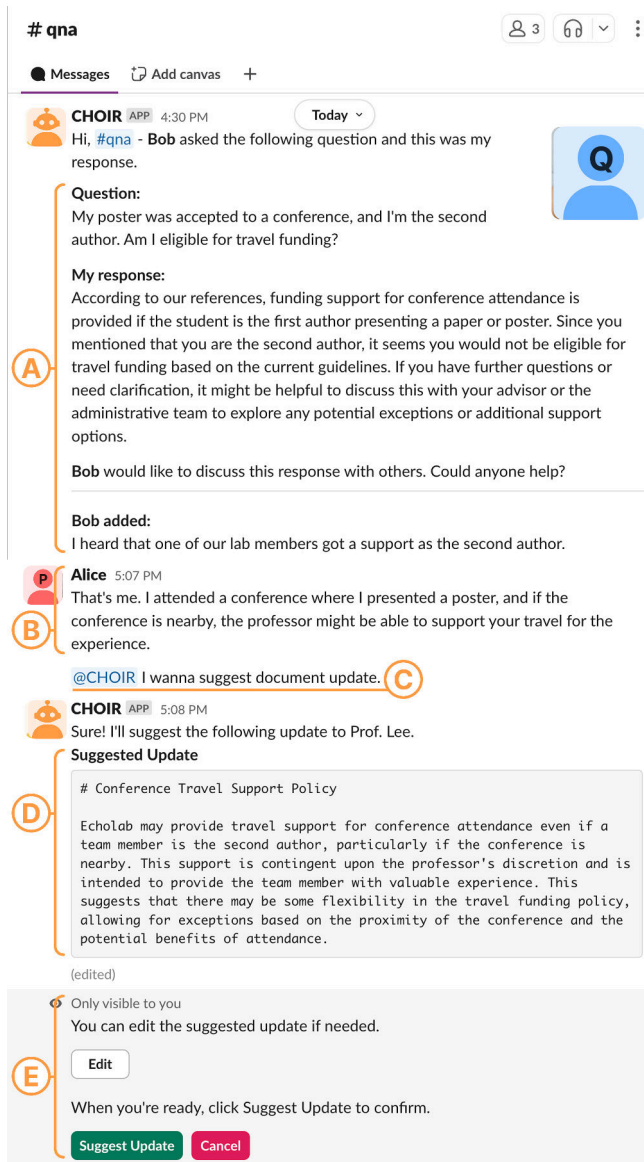


Figure 5: Example of knowledge extraction after a Q&A is shared. The Q&A exchange has been posted to the designated channel through the Share with Q&A modal. A peer (Alice) adds a follow-up comment and then mentions CHOIR to request an update, prompting CHOIR to extract documentable knowledge from the conversation.

can be multiple places to update in a document, or stop the review entirely (C). If the knowledge does not fit into an existing section, the interface provides an option to create a new section through a dedicated modal (D). Applied updates—whether to an existing file, a new file, or a new section—are written back to the repository, and CHOIR posts a message to the original channel acknowledging that the update has been finished.

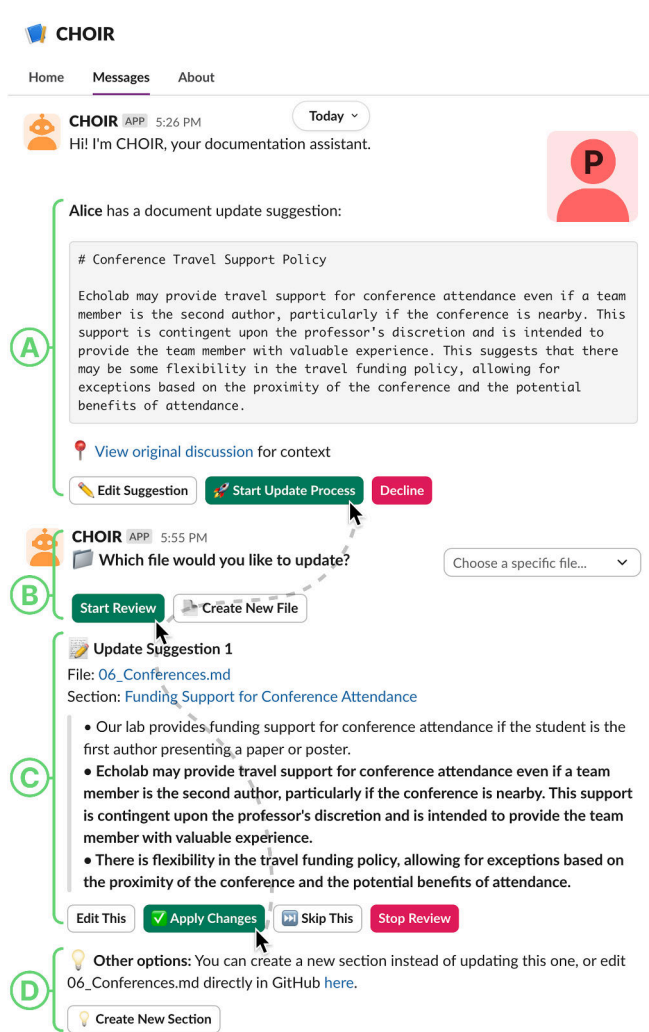


Figure 6: Manager-facing document update workflow in CHOIR. A member’s suggestion (from Alice) arrives as a Slack DM with attribution and a link to the original discussion. After starting the update process, the manager reviews the proposed change in context and decides how to incorporate it into the documentation.

This feature implements DG1 by automating key steps: recommending where updates are needed, proposing content, and committing changes. Managers can update documents directly in Slack without switching tools. CHOIR preserves managerial autonomy by allowing edits to proposed content or, for more complex tasks such as changing headings or reorganizing structure, enabling direct changes on GitHub via provided hyperlinks (“here” in D in Figure 6).

4.5 System Implementation

CHOIR uses Slack as the front-end client and is built with the Slack Bolt² framework. The backend connects to the lab’s GitHub repository through the GitHub REST API³, which serves as the persistent store for documentation in Markdown format. This stack allows Slack to serve as the entry point for everyday interactions, while GitHub provides version control for organizational memory.

For Q&A, CHOIR maintains an embedding index over the documentation repository and performs retrieval-augmented generation conditioned on the user’s question and the retrieved text. We use GPT-4o (temperature=0, max_tokens=1000) for text generation and text-embedding-3-small for document embeddings via the OpenAI API. Input handling is implemented through prompting: the LLM is instructed to decide whether an incoming message should be treated as a Q&A query or an update request. Responses to questions, knowledge extraction from conversations, and documentation update drafts are all generated via the OpenAI API. To protect privacy, user names are replaced with pseudonyms before any message is sent to the API. The complete system prompts are provided in Appendix A.

5 Field Study

To understand how lab members engage with and experience a chatbot-driven system for managing and utilizing organizational memory (RQ2), we conducted a month-long field deployment of CHOIR in the Slack workspaces of four university research labs, including the authors’ own.

5.1 Study Design

To explore these questions, we conducted an exploratory field deployment study that combined analysis of system usage logs with user interviews [30], enabling us to capture both behavioral patterns and the reasoning behind user actions. We deployed CHOIR in active research laboratories to observe usage and experiences within participants’ real work environments [85]. This field study approach reflects the often underestimated difficulty of evaluating CSCW applications: organizational systems must be assessed within authentic work contexts, as in-lab experiments that attempt to simulate ecologically valid environments cannot reproduce the social dynamics and group interactions that unfold over weeks and months [45].

Our evaluation focuses on understanding how CHOIR mediates organizational knowledge management practices, thereby demonstrating the feasibility of CHOIR and the benefits observed through user behaviors and interview responses. Field studies of organizational memory systems have long shown how such systems augment collective memory within real organizational contexts and workflows [1]. While field deployments cannot establish causal relationships due to the challenges of creating controlled baselines, they enable a holistic understanding via a *descriptive* approach [63] that characterizes how CHOIR mediates organizational communication and memory practices in situ. Many social computing studies that involve a group of users, an organization, or a community, similarly choose to evaluate

²<https://api.slack.com/bolt>

³<https://docs.github.com/en/rest>

Table 1: Participants in the field study: each lab’s director and members, including newcomer status. Lab age is shown in parentheses below each lab number.

Lab (years)	Department (Research Area)	Label (Gender)	Status
L1 (4 yrs)	Computer Science (Software Engineering)	D1 (M)	Director
		S1-1 (F)	New student
		S1-2 (F)	Existing student
		S1-3 (M)	Existing student
		S1-4 (M)	Existing student
L2 (4 yrs)	Computer Science (HCI)	D2 (F)	Director
		S2-1 (M)	New student
		S2-2 (M)	New student
		S2-3 (F)	Existing student
		S2-4 (M)	Existing student
		S2-5 (F)	Existing student
L3 (4 yrs)	Industrial Engineering (Physical Ergonomics)	D3 (F)	Director
		S3-1 (M)	New student
		S3-2 (F)	New student
		S3-3 (F)	Existing student
		S3-4 (F)	Existing student
		S3-5 (M)	Existing student
L4 (7 yrs)	Computer Science (HCI)	D4 (M)	Director (Author)
		S4-1 (M)	Existing student
		S4-2 (M)	Existing student
		S4-3 (M)	Existing student
		S4-4 (F)	Existing student

systems through deploying the system in the wild and takes the descriptive approach [14, 29, 53, 60, 68, 69, 73]. In studying sociotechnical interventions in organizational contexts, authentic workplace dynamics provide crucial insights into user engagement and experience with collaborative knowledge management systems.

5.2 Participants and Recruitment

We recruited university research lab directors through our social media and from formative study participants who expressed interest in the field study. Four research laboratories participated in the user study, comprising 21 participants: three external lab directors and 18 students, plus the corresponding author who served as a lab director. Upon confirmation from each director, we advertised the study in their Slack workspace and recruited students from the lab, endorsed by the lab director. We refer to each unit and participant based on the labels as specified in Table 1. The one-month deployment was conducted during the Summer of 2025, a transitional academic period when each lab had remote incoming graduate students joining. This timing allowed us to observe how CHOIR supports early-stage onboarding and knowledge exchange when new members have many initial questions.

All participants provided informed consent for interviews, system usage monitoring, and interaction log collection. The study received IRB approval (IRB 25-247). As compensation, lab directors received \$40 and students received \$20 upon completing the interviews. One participating lab (L4) was led by the corresponding

author, who served as the lab director. In addition to easing recruitment, the study was motivated by the director's experience with manually curating organizational memory and observing the limited use of these materials in practice. To mitigate any potential pressure that L4 participants might feel due to power dynamics or a desire to be favorable toward their peers, the director limited their role to answering questions and making changes to documents and did not otherwise intervene in the study procedures. Similarly, the first and second authors, who are members of L4, did not use CHOIR in this field deployment at all. All L4 participants voluntarily consented to the study independently through the standard procedure of advertisement, recruitment, and onboarding sessions, moderated by the first and second authors.

5.3 Deployment Process

The field deployment occurred in two phases: onboarding and active deployment.

5.3.1 Phase 1: Document Collection and Onboarding Sessions. Prior to the deployment period, we collected existing lab documents and conducted structured onboarding sessions with each participating lab.

Document Collection: Before the lab director onboarding sessions, we asked professors to share existing lab documents or materials they wished to integrate with CHOIR. We specifically requested only documents that could be made publicly available. The participating labs exhibited diverse documentation practices. D1 and D3 had no existing organizational memory systems and primarily relied on informal communication channels such as Slack messages, verbal exchanges, or external resources (e.g., department handbooks). To address this and avoid a cold-start problem, the research team worked with them to create initial seed documents based on available resources (e.g., lab websites and departmental handbooks), ensuring that CHOIR would contain useful content immediately upon deployment. In contrast, D2 and D4 already maintained web-based documentation, including a lab webpage and a GitHub repository containing the director's advice, mentoring plan, and lab policies. However, during the study period, they voluntarily used CHOIR rather than their existing systems. The research team, with the directors' approval and guidance, converted all provided materials into markdown files, uploaded them to a public GitHub repository for each lab, and granted lab directors write access and students read access.

Lab Director Onboarding Sessions (45–60 minutes) were conducted via Zoom with lab directors only. These sessions covered CHOIR's knowledge management capabilities, administrative features, the GitHub repository structure, and manager approval workflows for document updates.

Student Onboarding Sessions (45–60 minutes) were also conducted via Zoom to accommodate as many participants as possible. These sessions focused on CHOIR's question-answering functionality, conversation facilitation features, and integration with the existing Slack workspace. Both onboarding sessions included: (1) a system demonstration with examples drawn from typical lab environments, (2) hands-on practice with key features, and (3) a brief semi-structured interview.

5.3.2 Phase 2: Active Deployment. Following onboarding, CHOIR was deployed as a Slack App within each lab's existing workspace. Participants were registered as authorized users with two distinct roles: lab directors served as managers with document approval access, while students operated as CHOIR users. No specific usage requirements were imposed on participants, while awareness of CHOIR's availability was maintained by having the research team periodically ask questions to CHOIR in shared channels. Participant privacy was safeguarded through strict data collection protocols and automatic anonymization processes of personally identifiable information in compliance with FERPA requirements. When eligible messages were sent to external language models (OpenAI API), all usernames were replaced with pseudonyms, and responses were de-anonymized before delivery to users.

5.4 Data Collection

We collected two types of data during the one-month field deployment (Summer 2025).

5.4.1 System Interaction Logs. We collected interaction data within the CHOIR application, including message exchanges with CHOIR, feature usage patterns, button clicks, and document update activities recorded both in CHOIR and on GitHub. To protect participant privacy, messages were transferred to our servers or the OpenAI API only in three cases: (1) direct messages sent to CHOIR, (2) messages that mentioned CHOIR, and (3) the ten most recent messages when CHOIR was mentioned for knowledge extraction. We tracked usage metrics such as frequency and timing of CHOIR interactions per user, types of questions asked, conversation thread lengths, manager approval rates for suggested updates, and user feedback actions.

5.4.2 Log-Informed Semi-Structured Interviews. Following the deployment period, we conducted semi-structured interviews with each lab director and with students who had used CHOIR at least once. Rather than treating logs and interviews as separate data sources, we used individual usage patterns from the logs as prompts during interviews, enabling participants to explain their reasoning and experiences. This approach allowed us to understand not only what users did, but also why they made particular choices and how they experienced the system.

5.5 Data Analysis

We analyzed system interaction logs to first generate descriptive statistics of usage patterns of each lab and identify individual usage patterns, including feature adoption, engagement trends, and workflow progress rates. These patterns also served as the foundation for personalized interview discussions, allowing us to explore the reasoning behind specific user behaviors.

Transcripts from post-deployment log-informed interviews were analyzed using reflexive thematic analysis [18]. Two researchers independently coded a subset of transcripts to establish themes, then collaboratively developed a coding framework applied to all interviews. Key themes focused on how participants used CHOIR for organizational memory, the barriers they faced, and their suggestions for improvement.

6 Results

Overall, CHOIR supported knowledge retrieval and updates to organizational memory during the study period, while also revealing nuanced challenges and opportunities for chatbot-mediated knowledge management systems, which we present below.

6.1 Q&A Interactions with CHOIR

6.1.1 “I Consider CHOIR as a Senior in the Lab.”

During the one-month field study, participants submitted a total of 107 questions to CHOIR through both direct messages and public channels. Of these queries, 45 (42%) were answered using the connected lab documentation, while CHOIR could not answer the remaining questions and revealed gaps in the existing knowledge base. Participants valued CHOIR’s effectiveness and reliability, with S2-1 noting that “it definitely improved the speed of retrieving the information compared to me reading through the document or searching online.” S1-2 appreciated its accuracy with reference: “if the information is already in the handbook, it’s pretty straightforward to get those information, and I like that it provided a reference [...] That helped me to get the accurate part information.” D1 appreciated the quality of responses: “I think in general the answers were pretty good. When the response or when the like data was there, further details were there.”

Beyond its functional utility, deploying CHOIR in the social context of Slack led participants to perceive and interact with it as if it occupied familiar social roles. S1-3 conceptualized CHOIR as a knowledgeable lab member: “I can go to the CHOIR to ask him, confirm the question[...] I consider CHOIR as a senior in the lab.” This perspective was shared by S4-2, who viewed CHOIR “more like a librarian. A library librarian tells you what’s in the library and not if it’s not there.” These metaphors reflected how participants made sense of CHOIR’s role within their lab ecosystem, framing its consistent availability and non-judgmental responses in terms that helped them integrate AI assistance into their daily lab interactions. When prompted, all lab members expressed a desire to keep CHOIR, and the app remains active at the time of submission.

6.1.2 How Lab Culture Shapes Question Asking. One notable pattern that we found from L1, as shown in Figure 7, was that most of the questioners asked CHOIR in a public channel instead of via DM, contrasting sharply with the other three labs. L1 students asked 14 questions publicly, with none through private messages. Students’ motivation to ask public questions stemmed from altruistic and supportive attitudes to help others. S1-3 explained: “I think it can share to everybody about the answer[...] everybody can just check the Slack and see.” This visibility allowed students who had similar questions to benefit from others’ queries without needing to ask again. Some students also used CHOIR to support their lab. S1-4 stated, “I mostly used it to help [D1]. Add stuff that wasn’t already there.” The public nature of these interactions enabled collective awareness and knowledge sharing and also served as a catalyst for documentation improvements, with multiple students explicitly using CHOIR to identify gaps in the handbook.

This distinctive use of public Q&A perhaps stemmed from their altruistic and supportive lab culture. When D1 was asked to speculate why this pattern showed, he mentioned “I think all

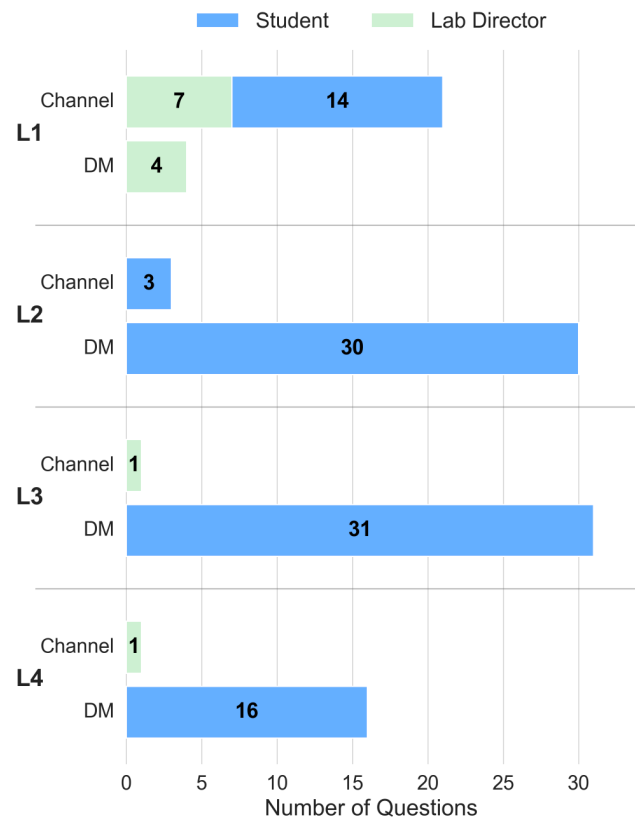


Figure 7: Distribution of questions asked to CHOIR across four research labs during the one-month field study. The graph shows the number of questions submitted through public channels versus direct messages (DM), broken down by students (blue) and lab directors (green). L1 shows a distinctive pattern with predominantly public channel usage, while L2, L3, and L4 primarily used private DMs for Q&A interactions.

the students get along pretty well[...] more comfortable with public questions and not being afraid to ask questions.” This comfort level was echoed by students, with S1-2 noting, “Our lab atmosphere is pretty good. We know each other well, so I don’t feel those kinds of questions cannot be asked publicly.” This mutual trust and familiarity created an environment where both directors and students felt comfortable sharing their uncertainties publicly, fostering a collaborative learning atmosphere. This collaborative approach resulted in D1 making more frequent updates than any other labs (Table 2), as questions surfaced with missing information that needed to be added.

In sharp contrast to L1, Figure 7 shows that L2, L3, and L4 predominantly used private messages. These participants sought a safe space to ask questions while avoiding the potential negative impression they may have from the question. S2-2 expressed, “I’m just afraid that my question could be like, could be dumb questions. So I just asked it first through DM because it’s more private.” Participants

also avoided public channels to prevent burdening others with potentially trivial questions, with S4-2 noting: “*when I don’t want to bother others,*” This sentiment was shared by lab directors.

6.1.3 Directors’ Personal Assistant to Convey Knowledge. While we did not expect lab directors to ask CHOIR questions beyond testing, lab directors repurposed CHOIR to communicate information to students. Directors Q&A to CHOIR showed different approaches to alternative usage patterns. D1 used public Q&A primarily as an acknowledgment about documentation updates: “*after I updated and asked the question, see if the response is correct. Also, maybe kind of to highlight to whoever asked the question that it is updated now.*” This approach transformed CHOIR into a broadcasting system, allowing D1 to signal to students that their previously unanswered questions had been addressed through documentation improvements. Similarly, D3 asked CHOIR a question once in a public channel to proactively share important information by asking questions such as “Can you tell how CITI training can be completed?” immediately after asking new students to take the CITI Training, to give them detailed instructions on which courses they needed to take. These cases demonstrate how lab directors leverage CHOIR as an assistant to acknowledge updates or disseminate information rather than simply seeking answers.

6.2 Knowledge Sharing Through Q&A

6.2.1 Q&A Sharing for Collecting Answers as well as Questions. CHOIR’s Share Q&A feature enabled users to share their questions and CHOIR’s responses with others, either in the designated Q&A channel or through private messages. Of the 95 questions asked by students, only 11 were shared (L1: 1, L2: 4, L3: 3, and L4: 3 with 1 shared privately with a director). Students appreciated the Sharing Q&A feature for being able to understand the context of someone’s question, as S2-3 noted “*I also like when it’s like sending the questions in the channel, like it’s also including what it had. So when we like, try to give our answers, we also have some reference to refer to.*” In the meantime, L1’s distinct pattern of public Q&A practice introduced in Section 6.1.2 effectively replaced the Share Q&A feature as most questions and answers were already visible to all students in L1 through their public interaction style.

Lab directors particularly appreciated the centralized Q&A channel for its organizational benefits. D2 valued the dedicated space for continuous interaction: “*I like the Q&A channel. So it’s like those that’s for people to ask questions and I can answer them. I think I would like to keep like Q&A channel [...] If people have questions, they can ask them there.*” D3 emphasized the channel’s role in managing communication flow: “*I think having this channel can kind of centralize all those messages in one place, so I like it.*” While Q&A sharing was primarily designed for getting answers for unanswered questions, it supported DG3 in a way that we did not anticipate. Overall, deploying CHOIR in these labs encouraged and fostered the culture of sharing knowledge through Q&A using various channels: plan Slack communication, asking CHOIR in a public channel, and sharing the Q&A feature, supporting our goal of facilitating interpersonal communication (DG3).

6.2.2 Having a Safe Space to Ask and Discuss. Beyond these patterns, CHOIR provided a safe channel for students to ask

questions anonymously, even when the contained anonymity came from the small scale of the participants. Notably, more than half of the shared questions (6/11) were posted anonymously, indicating users’ desire for privacy even when seeking broader input, which could have kept them from sharing it if it were not anonymous. S3-1 explained their use of anonymity: “*maybe thinking the question is kind of stupid question[...] if some senior student may think, okay, this question, maybe I’m supposed to know it, maybe it’s in the graduate manual, but I just didn’t read it. If I ask people going to think I’m stupid and lazy.*” Participants also cited finance- or budget-related topics as hypothetical examples that might warrant anonymous sharing, with S2-2 noting, “*sometimes I don’t want to reveal myself because I know there are some sensitive information like finance or budget.*” One such example involved a new Ph.D. student anonymously asking their peers when they would receive their first stipend—an important question that could have been awkward to pose to an advisor, who might not know the answer, and one the student might not feel comfortable asking of any peers.

Students had an additional motivation to share a question using CHOIR as a way to initiate discussions on topics where people might hold different opinions. Rather than seeking definite answers, participants shared questions to collaboratively explore a topic (e.g., “How many papers are needed to graduate?”), rather than expecting to get a definite answer. S2-2 stated, “*I just want to see like how others think about the topic,*” and S4-4 sought “*some suggestions from different people, their personal suggestions instead of an article.*” These examples demonstrate how CHOIR created a space for students to exchange perspectives and engage in open-ended conversations that they did not necessarily think should be documented in the organizational memory.

6.2.3 Question Sharing Complements, not Block, Interpersonal Communication. Despite the availability of sharing features, the vast majority of questions (86 out of 95) remained unshared, raising questions about both social and practical barriers to collaborative question and knowledge sharing. Instead, students chose to communicate directly with peers instead of CHOIR-mediated sharing. S3-3 explained, “*the most straight path for me was to just ask D3 in a direct message, rather than using CHOIR as, like a middleman.*” Similarly, S2-3 reported, “*I didn’t ask it anonymously in the Q&A channel because I know if there is someone who could answer my question, then it would be a specific senior PhD student, and I would probably just go to them and ask them directly.*” Overall, the limited activity of Q&A sharing was not necessarily due to barriers or fear; rather, students often chose traditional modes of communication to obtain answers more directly.

6.2.4 When Questions Stay Hidden: Barriers to Director Awareness. The limited activity in Sharing Q&A created a significant awareness gap for lab directors, particularly those in L2 and L3, who had limited visibility into what questions students were asking and which gaps existed in their documentation. During post-study interviews, in general, D2 and D3 wondered if anyone really used CHOIR, although the students were actively asking questions about it behind the scenes, which came to their surprise when we presented the statistics. D3 reflected on this limitation: “*I didn’t receive any request [to answer questions] directly [...] I also don’t know what kind of questions they’ve asked.*” While we anticipated

that the Sharing Q&A feature would raise awareness of gaps in the documentation, students' preference for direct communication limited the lab directors' view on what students want to know and CHOIR do not know. D2 recognized the value of knowing student questions: *"it might be helpful for me to know what questions are asked so that I know what they are confused about [...] if I don't know what they are confused about, then there's nothing I can do to clarify the confusion."* This lack of visibility prevented directors from identifying knowledge gaps, which is crucial for understanding student confusion, or updating documentation proactively. In contrast, L1's culture of public Q&A provided natural visibility into student needs and documentation gaps, enabling more responsive documentation maintenance.

6.3 Knowledge Extraction and Document Update Practices

6.3.1 Perceived Value of Document Update via CHOIR. Participants recognized significant value in collaborative documentation, particularly appreciating the potential for preserving and sharing knowledge within research labs. The organizational value of maintaining documentation was widely acknowledged, with D1, who did not have any organizational memory beforehand, expressing satisfaction: *"the documents. I'm very satisfied [...] I think it'll continue to live on and continue to be updated even after the study."* D2 emphasized the importance of capturing firsthand student experiences: *"because they have firsthand information [...] if they don't document this, then the next person would have to figure it out by themselves again."* S1-1 highlighted the cumulative value for future members: *"the more questions asked and updated, the more useful it becomes for future students."*

Directors particularly valued the AI-assisted update workflow, which reduced the effort required for documentation maintenance. D3 appreciated CHOIR's editing suggestions: *"CHOIR suggested me a sort of edit, which I liked [...] I don't have to spend a lot of time crafting."* D2 found the process straightforward, noting, *"I just click add and I think it asks me to select a document."* Directors who either received update requests from students or initiated updates themselves incorporated them into the lab documentation, suggesting the effectiveness of the workflow when participants chose to engage.

6.3.2 Directors' Update Practices and Member Engagement. During the study, lab directors exhibited varying update patterns, as summarized in Table 2. D3's only knowledge extraction attempt was made to update the document in a direct message with CHOIR. On the other hand, D2 proactively used CHOIR's knowledge extraction four times, completing four varying types of document updates. Two document updates originated from private Q&A sessions where S2-2 shared questions privately, prompting D2 to provide direct answers and mention CHOIR to trigger knowledge extraction. Another extraction occurred when S2-3 and S2-5 shared their experiences in response to an anonymous question shared in a channel, leading D2 to mention CHOIR to capture students' answers into the document. The fourth extraction involved a more complex collaborative process where D2 initially provided an answer while mentioning senior students to supplement the answer with additional information; when S2-3 subsequently provided the

Table 2: Documentation update patterns across research labs during the one-month field study, showing commits, files modified, and word changes for both direct GitHub updates and CHOIR-assisted updates.

Lab	Update			
	Method	Commits	Files	Words (+/-)
L1	Direct	19	5	+1153 / -63
	CHOIR	0	0	0
L2	Direct	1	1	+1 / -1
	CHOIR	4	1	+261 / -0
L3	Direct	0	0	0
	CHOIR	1	1	+109 / -1
L4	Direct	7	2	+122 / -29
	CHOIR	6	2	+348 / -2
Subtotal	Direct	27	8	+1276 / -93
	CHOIR	11	4	+718 / -3
Total	–	38	12	+1994 / -96

requested information via message, D2 mentioned CHOIR again to extract and document this enhanced response. This variety of cases exemplified how professors can leverage the collaborative and communicative nature of the workspace for documentation updates even when they lack hands-on knowledge in specific areas.

Beyond these expected uses, others also actively updated documents through the document editor available in GitHub. D4, one of the authors, showed a balanced approach, using both direct updates (7 commits) and CHOIR-assisted updates (6 commits) to maintain their documentation. D1 performed updates outside CHOIR through GitHub after identifying gaps via public Q&A, resulting in the highest volume of documentation changes with 19 commits across 5 files and 1,153 words added. When asked why he directly edited, D1 explained, *"Probably just out of convenience. It was easier. Or maybe I was just more familiar with pulling up the repo and making a change directly."* Given that he started with nearly empty documents, he would not have faced the burden of locating specific sections to update, as someone managing a longer document might have.

6.3.3 Barriers to Student Contributions in Documentation. Although some students contributed to the knowledge when their lab director specifically requested, as in D2's case above, students did not actively engage in initiating document updates; only three participants (S2-2, S2-3, S3-2) attempted to use the feature, all through private messages. These three members initiated a total of five updates during the study period. Although we expected students to request updates through the public Q&A channel, they primarily discovered documentation gaps during private Q&A with CHOIR, providing corrections based on their own knowledge before proposing updates to directors. For example, S3-2 noted, *"I was asking about which courses to take, and CHOIR talked about the*

two required courses, but the course code was incorrect [...] And so I just updated it.” This pattern suggests the feasibility of CHOIR in enabling students to contribute to documentation.

Despite recognizing its value, students’ participation in documenting updates remained limited due to several concerns. Knowledge-related concerns emerged as the primary obstacle. A lack of confidence in providing correct and accurate answers was a central issue, with S3-3 worrying about the permanence of contributions: “if I answered that question that would stay in the database forever [...] if it was wrong, then it would mislead all the students.” S2-3 expressed a similar hesitation: “sometimes I’m not sure about the answer[...] the wrong message would be there for some time and it might mislead.”

Students’ limited understanding also prevented them from generalizing personal experiences into knowledge useful for everyone, posing another challenge. Participants reported that not all knowledge was suitable for general documentation. S3-4 explained, “especially regarding the funding sources, some things are not generic, they can be very specific to someone.” At the stage of design, we believed that DG4 would be supported by having professors act as gatekeepers of how knowledge is documented in organizational memory. However, our findings indicate that students’ psychological barriers can be a larger issue that sheds lights on the need to foster a culture where documentation is viewed as a collective and welcoming practice rather than an individual risk.

6.4 Participant Suggestions for Improvement

Throughout the study, participants identified limitations in CHOIR’s capabilities and offered suggestions for improvement based on their experiences. These recommendations revealed both immediate needs and longer-term visions for AI-mediated knowledge management in research labs.

6.4.1 Manager Awareness and Communication. For lab directors, the predominance of private Q&A created a particularly problematic awareness gap as stated in Section 6.2.4. To address this critical gap, participants suggested several solutions for improving manager awareness while preserving student privacy. Participants suggested that students’ questions be aggregated for periodic high-level summaries. S2-4 proposed to “take all of them, strip the metadata, and then have a nice summary of the key points and send that as a weekly update to the professor.” The need for users’ option to opt out from being reported was also suggested, with S3-4 emphasizing that “there should be an option. If we want to send the question to the advisor or not.” These suggestions reflected a nuanced reluctance that they have for contributing to the organizational memory and preserving student autonomy and privacy.

6.4.2 Expanding Knowledge Base Access. Participants expressed frustration with CHOIR’s limitation to internal documentation—a constraint stemming less from the system itself than from the limited content available in the repository—and suggested multiple avenues for expanding its knowledge sources. For external resources, D3 noted that “it would be better that we can search relevant information from other sources,” while S3-3 proposed that “if it can access some websites or other miscellaneous web pages or

some forums, maybe it would be more helpful.” Integration with internal resources was considered equally important, with S2-5 expressing strong interest in “access to the Google Docs. I think in our group, the advisor mentioned everything in the Google docs.” Beyond documents, participants envisioned CHOIR as a bridge across labs and expertise sources. Cross-lab knowledge sharing emerged as particularly valuable, with D2 expressing interest in seeing “what questions students ask[...] if the answers are shareable, I would like to see their answers too.” Similarly, S2-5 suggested leveraging other labs’ experiences: “if there are datas of other lab policies[...] CHOIR could give suggestions on what is missing in our lab policy and elicit discussions.” These suggestions point toward a more interconnected knowledge ecosystem—one that bridges internal and external resources, integrates documents and people, and extends across different research groups.

6.4.3 Context-Aware and Personalized Support. Participants imagined a system that could adapt to different contexts and individual needs rather than providing one-size-fits-all responses. For situation-specific guidance, S3-3 suggested differentiated support: “textit if there can be a process where there are some situations created, for first-year or second-year PhD students, then maybe there can be some other ways to update that information.” S3-4 emphasized the need for flexibility: “there should be some flexibility in just being able to add our own situation right now, I think that is missing[...] whatever I will write, it will just be more generic. That will apply to everyone, which is not true.” These reflections highlighted participants’ recognition that research lab contexts vary significantly depending on factors such as career stage, research area, and individual circumstances. Providing personalized answers was expected to enhance CHOIR’s usefulness, pointing to the need for detailed, contextualized, and case-based content—beyond the lab’s formal policies in the repository.

7 Discussion

We reflect on our findings to understand implications for designing AI-mediated organizational memory systems. Our field study revealed several key tensions and opportunities that extend beyond CHOIR’s technical features to more fundamental questions about how knowledge management systems can support collaborative work in research organizations.

7.1 Privacy-Awareness Tension in AI-mediated Organizational Memory

Our findings revealed a fundamental tension between students’ privacy preferences and managers’ need for awareness in organizational memory systems. Because the majority of Q&A took place through direct messages to CHOIR without being shared, directors had limited visibility into knowledge gaps in the documentation. Directors expressed willingness to update information that their students could not access through interactions with CHOIR. This challenge reflects a broader issue in collaborative work: maintaining awareness of individual and group activities is essential for successful collaboration [33, 95], yet our study shows that AI-mediated knowledge retrieval can suffer from the same limitation.

Our analysis suggests that AI intermediaries such as CHOIR reconfigure organizational memory (OM) practices by redistributing visibility across members of an organization. By enabling workers to pose questions directly to an AI agent, these systems increase the visibility of existing organizational knowledge to questioners while reducing the workload of searching or manually locating information. However, this enhanced accessibility simultaneously diminishes members' organic awareness of knowledge gaps, which was previously supported through in-person interactions. This redistribution produces new asymmetries and hidden dynamics in a group, requiring members to negotiate trade-offs between individual privacy, exposure of information needs, and the collective benefits of shared organizational knowledge. Further, the interaction traces generated through AI-mediated OM—such as question-asking histories—can serve as a valuable resource for identifying missing knowledge or understanding evolving information needs within the organization. Yet leveraging such data also risks inadvertently creating forms of workplace surveillance, particularly when anonymity is contained in a smaller group. These dynamics reveal that AI-mediated OM systems do not merely support knowledge access; they actively reshape the visibility, accountability, and communication of knowledge sharing practice, offering a lens for understanding how algorithmic intermediaries influence sociotechnical dimensions of OM.

This redistribution of visibility intersects with a long-standing tension in communication research between psychological safety and organizational transparency. The tension stems from what communication researchers have identified as the fundamental trade-off between psychological safety and organizational transparency. Students' behaviors—seeking private channels to avoid judgment and protecting others from perceived disruption—represent manifestations of psychological safety-seeking in hierarchical organizations. Research has consistently demonstrated that anonymous communication reduces social pressure and enables more honest self-expression, particularly when interacting with authority figures [26, 93]. However, this creates what organizational behavior scholars have identified as a fundamental paradox: while privacy enables honest participation, complete transparency can also discourage participation by making members feel overly monitored [15].

These findings suggest that future organizational memory systems must balance individual privacy with the members' benefit coming from reinforced organizational memory. Our participants suggested privacy-preserving awareness mechanisms, one of which is to periodically report questions at high levels, which would inform managers about question patterns without compromising individual privacy. Given the students' preference for direction communication over CHOIR-mediated communication, this alternative could be more effective for lab directors to be aware of the gaps in organizational memory. This approach preserves student autonomy while providing managers with actionable insights about knowledge gaps and documentation needs.

7.2 Expanding the Perceived Boundary of Documentable Knowledge

Our study revealed a gap between the knowledge that lab members possessed and what they felt was appropriate to document. Many participants hesitated to share context-specific experiences that may be generalizable and beneficial to their peers. This reluctance appears to stem from an implicit expectation that documentation should be universal, authoritative, and permanent—creating a high bar for contribution that excludes valuable experiential knowledge and partial solutions. Participants also recognized value in case-based information and partial policies that could provide personalized guidance when comprehensive documentation was unavailable.

This challenge is well-studied: from early work on organizational memory systems [28] to research on knowledge sharing barriers [2], studies have consistently identified the difficulty of capturing informal knowledge due to cultural and technical barriers toward outcome-focused documentation. Moreover, research has shown that knowledge needs in document work are inherently role-dependent and context-specific [52], challenging assumptions of universal documentation. Prior work on enterprise wikis found that members only occasionally contributed despite recognizing documentation value, often requiring explicit requests from managers [57, 73]. CHOIR addressed this barrier by embedding documentation opportunities within natural conversation flows. Participants noted that the knowledge extraction feature made contributing informal knowledge feel less burdensome, as it transformed ongoing discussions into documentation without requiring separate writing effort. This suggests that conversational AI systems can provide lightweight pathways for capturing experiential knowledge that traditional documentation systems have struggled to elicit.

Building on these findings, we propose a more expansive model of organizational memory that embraces contextual knowledge alongside formal policies. Future systems should support lightweight contribution mechanisms that allow members to share experiences, partial solutions, and case-based knowledge without the burden of creating universally applicable documentation. Furthermore, systems can provide guided pathways that help users navigate between formal documentation and informal human expertise when contributing knowledge. This concept aligns with early systems like Answer Garden [5]. In addition, establishing common ground between lab directors and students, and fostering a welcoming environment even when the knowledge is canonical, may encourage students to contribute.

7.3 AI Agents as Mediator in Organizations

While recent research has demonstrated significant declines in human participation following AI introduction in knowledge communities [19], our findings demonstrate how AI-based chatbots can lead to markedly different outcomes for social dynamics within organizations. Rather than replacing human communication, CHOIR served as an intermediary that facilitated and preserved subsequent human interactions. Evidence for this facilitative effect was particularly strong in L1, where public Q&A usage corresponded with increased documentation activity, collaborative

knowledge identification, and peer learning. Students explicitly used CHOIR to help identify missing documentation for their director, demonstrating how AI-mediated knowledge work can complement human collaboration.

From a socio-technical perspective, CHOIR’s mediating role can be understood through the lens of boundary objects [91]. Organizational memory documents serve as boundary objects with interpretive flexibility, where directors view them as comprehensive repositories of lab knowledge that they want to convey to students, while students see them as practical guides for navigating their immediate questions and expectations. In our work, an AI agent mediates existing boundary objects to each group, and the AI agent itself is a new boundary object that each group can access and communicate with. Extending recent work on AI systems as boundary objects in creative and participatory contexts [13, 25, 48, 83], CHOIR shows how AI can serve this role in organizational memory, dynamically adapting the same knowledge base to serve different communities’ needs. This mediation extends organizational knowledge from context-dependent exchanges to persistent and reusable organizational memory that different members can retrieve and apply to their own situations.

These findings suggest an alternative to common framings of AI as either an assistant or a replacement for the human workforce. Our results point toward a model where AI serves as a facilitator that creates multiple channels for more effective human communication. This perspective aligns with prior work exploring conversational agents in collaborative settings [55, 89]. In our study, CHOIR reduced social barriers (through anonymity options), provided shared reference points for discussion (through document grounding), and surfaced implicit knowledge needs (through gap identification). The design implications extend beyond simple automation to consider how AI systems can actively foster human connection and collaboration.

7.4 Knowledge Access across Organizations

While our study focused on individual research labs, participants’ needs frequently extended beyond their immediate organizational boundaries. Participants expressed frustration with CHOIR’s limited access to external resources (e.g., a webpage on their department website) and desired connections to higher-level institutional documentation. This reflects a fundamental mismatch: users naturally seek knowledge at multiple levels and across the organization. Research on online communities demonstrates the value of such cross-organizational knowledge sharing, showing how similar governance structures across communities can provide valuable resources for new organizations lacking documentation [21, 22, 38]. This challenge would reflect the complexity found in other types of organizations, such as companies with many specialized teams performing different functions. Recent work on modular knowledge architectures offers promising technical approaches [107]. While such approaches primarily address retrieval challenges, they point toward architectures that could support the complex permission structures, review workflows, and cross-boundary access patterns that hierarchical organizations require.

7.5 Limitations

Our study has several limitations that may impact the generalizability. First, our field deployment involved a small sample of participants across four university research labs, all from STEM fields with high technical literacy. This limits generalizability to organizations with different technical backgrounds, hierarchical structures, or cultural norms around knowledge sharing. Second, the one-month deployment period was sufficient for observing initial adoption patterns but insufficient for evaluating the longer-term effectiveness of organizational memory systems. Our study focused on discovering potential effects rather than proving sustained efficacy.

Third, because the study was conducted during the summer, the findings primarily reflect the onboarding experiences of new remote students. Although this context provided insights into early-stage lab integration, usage patterns during regular academic semesters may differ. We plan to conduct a long-term follow-up with participating labs to examine if and how usage and perceptions evolve across different academic periods. Additionally, the study setup required preparing initial documentation for participating labs by creating new seed documents or converting existing materials into standardized formats. We acknowledge that these prepared documents may have affected ecological validity compared with naturally evolved documentation.

Fourth, one participating lab (L4) was led by the corresponding author, which may introduce potential bias despite our efforts to minimize researcher influence. L4 may also be particularly well aligned with the assumptions and practices that CHOIR supports, given the system’s origin in addressing needs observed in this context. While the corresponding author limited their role to managing the process, answering questions, and approving documents, and the first and second authors did not participate as study subjects, peer-pressure inherent in being in the same lab or the natural interest in a research project that is thematically similar to their own research could have influenced the L4 participants’ engagement patterns or feedback. However, the key findings and discussion points are valid as they are not disproportionately driven by L4 data alone.

Lastly, our study lacks a controlled baseline and quantitative performance measures beyond descriptive statistics, which limits our ability to make causal claims about CHOIR’s effectiveness compared to alternatives such as wiki pages or existing lab practices (e.g., Google Docs/Slack Pinned Messages). As a result, our findings are descriptive rather than comparative, and further research is needed to determine how specific features of CHOIR contribute to its impact. Future controlled experiments could compare an improved version of CHOIR against baseline conditions where shared documents and Slack are not connected through an AI mediator. Such studies could measure question response time, documentation update frequency and quality, knowledge accessibility, and participation in organizational memory maintenance.

8 Conclusion

In this paper, we presented CHOIR, a chatbot-mediated organizational memory system that integrates document-grounded Q&A, knowledge sharing, conversation-based knowledge extraction, and

AI-assisted document updates within workplace communication platforms. Through a formative study with 15 participants, we identified key challenges in managing organizational memory in research labs, which guided our design of CHOIR. We then deployed CHOIR across four university research labs for one month, observing that it supported knowledge retrieval and documentation practices, with 42% of questions answered from existing documentation and 1,994 words added to organizational memory. Our findings revealed a fundamental privacy-awareness tension in AI-mediated organizational memory systems, where students' privacy-seeking behaviors, driven by concerns about social evaluation, limited directors' awareness of knowledge gaps. We found that conversational agents can serve as facilitators rather than replacements for human communication, offering multiple channels for more effective human communication while identifying gaps in organizational documentation. Our work provides empirical observations and design implications for integrating AI-mediated organizational memory into workplace communication platforms, highlighting the importance of addressing psychological safety, power dynamics, and cultural norms in hierarchical organizations. These findings offer insights for designing knowledge management systems that balance individual privacy with organizational transparency and support capturing contextual knowledge that often remains undocumented.

Acknowledgments

We thank all professors and students who participated in the user studies, as well as the reviewers for their constructive feedback. Funding for conference presentation was provided by the Center for Human-Computer Interaction at Virginia Tech (CHCI@VT).

References

- Mark S. Ackerman. 1998. Augmenting Organizational Memory: A Field Study of Answer Garden. *ACM Transactions on Information Systems* 16, 3 (July 1998), 203–224. doi:10.1145/290159.290160
- Mark S. Ackerman, Juri Dachtera, Volkmar Pipek, and Volker Wulf. 2013. Sharing Knowledge and Expertise: The CSCW View of Knowledge Management. *Computer Supported Cooperative Work (CSCW)* 22, 4 (Aug. 2013), 531–573. doi:10.1007/s10606-013-9192-8
- Mark S. Ackerman and Christine Halverson. 1998. Considering an Organization's Memory. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work*. ACM, Seattle Washington USA, 39–48. doi:10.1145/289444.289461
- Mark S. Ackerman and Christine Halverson. 2004. Organizational Memory as Objects, Processes, and Trajectories: An Examination of Organizational Memory in Use. *Computer Supported Cooperative Work (CSCW)* 13, 2 (April 2004), 155–189. doi:10.1023/B:COSU.0000045805.77534.2a
- M. S. Ackerman and T. W. Malone. 1990. Answer Garden: A Tool for Growing Organizational Memory. *ACM SIGOIS Bulletin* 11, 2-3 (April 1990), 31–39. doi:10.1145/91478.91485
- Mark S. Ackerman and David W. McDonald. 1996. Answer Garden 2: Merging Organizational Memory with Collaborative Help. In *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work*. ACM, Boston Massachusetts USA, 97–105. doi:10.1145/240080.240203
- Mark S. Ackerman, Anne Swenson, Stephen Cotterill, and Kurtis DeMaagd. 2003. I-DIAG: From Community Discussion to Knowledge Distillation. In *Communities and Technologies*, Marleen Huysman, Etienne Wenger, and Volker Wulf (Eds.). Springer Netherlands, Dordrecht, 307–325. doi:10.1007/978-94-017-0115-0_16
- Eleni Adamopoulou and Lefteris Moussiades. 2020. An Overview of Chatbot Technology. In *Artificial Intelligence Applications and Innovations*, Ilias Maglogiannis, Lazaros Iliadis, and Elias Pimenidis (Eds.). Springer International Publishing, Cham, 373–383. doi:10.1007/978-3-030-49186-4_31
- Abram Anders. 2016. Team Communication Platforms and Emergent Social Collaboration Practices. *International Journal of Business Communication* 53, 2 (April 2016), 224–261. doi:10.1177/2329488415627273
- John R. Anderson, Lynne M. Reder, and Herbert A. Simon. 1996. Situated Learning and Education1. *Educational Researcher* 25, 4 (May 1996), 5–11. doi:10.3102/0013189X025004005
- Ofer Arazy, Oded Nov, Raymond Patterson, and Lisa Yeo. 2011. Information Quality in Wikipedia: The Effects of Group Composition and Task Conflict. *Journal of Management Information Systems* 27, 4 (April 2011), 71–98. doi:10.2753/MIS0742-1222270403
- Sumit Asthana, Sagi Hilleli, Pengcheng He, and Aaron Halfaker. 2025. Summaries, Highlights, and Action Items: Design, Implementation and Evaluation of an LLM-powered Meeting Recap System. *Proc. ACM Hum.-Comput. Interact.* 9, 2 (May 2025), CSCW176:1–CSCW176:29. doi:10.1145/3711074
- A. Ayobi, K. Stawarz, D. Katz, P. Marshall, T. Yamagata, R. Santos-Rodríguez, P. Flach, and A. A. O'Kane. 2021. Machine Learning Explanations as Boundary Objects: How AI Researchers Explain and Non-Experts Perceive Machine Learning. In *2021 Joint ACM Conference on Intelligent User Interfaces Workshops, ACM/IUI-WS 2021*. <http://ceur-ws.org/Vol-2903/IUI21WS-TESS-3.pdf>
- Shreya Bali, Pranav Khadpe, Geoff Kaufman, and Chinmay Kulkarni. 2023. Nooks: Social Spaces to Lower Hesitations in Interacting with New People at Work. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–18. doi:10.1145/3544548.3580796
- Ethan S. Bernstein. 2012. The Transparency Paradox: A Role for Privacy in Organizational Learning and Operational Control. *Administrative Science Quarterly* 57, 2 (June 2012), 181–216. doi:10.1177/0001839212453028
- Dane Bertram, Amy Volda, Saul Greenberg, and Robert Walker. 2010. Communication, Collaboration, and Bugs: The Social Nature of Issue Tracking in Small, Collocated Teams. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (CSCW '10)*. Association for Computing Machinery, New York, NY, USA, 291–300. doi:10.1145/1718918.1718972
- Richard Boardman and M. Angela Sasse. 2004. "Stuff Goes into the Computer and Doesn't come out": A Cross-Tool Study of Personal Information Management. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. Association for Computing Machinery, New York, NY, USA, 583–590. doi:10.1145/985692.985766
- Virginia Braun and Victoria Clarke. 2019. Reflecting on Reflexive Thematic Analysis. *Qualitative Research in Sport, Exercise and Health* 11, 4 (Aug. 2019), 589–597. doi:10.1080/2159676X.2019.1628806
- Gordon Burtch, Dokyun Lee, and Zhichen Chen. 2024. The Consequences of Generative AI for Online Knowledge Communities. *Scientific Reports* 14, 1 (May 2024), 10413. doi:10.1038/s41598-024-61221-0
- R.G. Capra and M.A. Perez-Quinones. 2005. Using Web Search Engines to Find and Refind Information. *Computer* 38, 10 (Oct. 2005), 36–42. doi:10.1109/MC.2005.355
- Eshwar Chandrasekharan, Chaitrali Gandhi, Matthew Wortley Mustelier, and Eric Gilbert. 2019. Crossmod: A Cross-Community Learning-based System to Assist Reddit Moderators. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW (Nov. 2019), 174:1–174:30. doi:10.1145/3359276
- Eshwar Chandrasekharan, Mattia Samory, Shagun Jhaver, Hunter Charvat, Amy Bruckman, Cliff Lampe, Jacob Eisenstein, and Eric Gilbert. 2018. The Internet's Hidden Rules: An Empirical Study of Reddit Norm Violations at Micro, Meso, and Macro Scales. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW (Nov. 2018), 32:1–32:25. doi:10.1145/3274301
- Preetha Chatterjee, Kostadin Damevski, Lori Pollock, Vinay Augustine, and Nicholas A. Kraft. 2019. Exploratory Study of Slack Q&A Chats as a Mining Source for Software Engineering Tools. In *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. 490–501. doi:10.1109/MSR.2019.00075
- Chun-Wei Chiang, Zhuoran Lu, Zhuoyan Li, and Ming Yin. 2024. Enhancing AI-Assisted Group Decision Making through LLM-Powered Devil's Advocate. In *Proceedings of the 29th International Conference on Intelligent User Interfaces (IUI '24)*. Association for Computing Machinery, New York, NY, USA, 103–119. doi:10.1145/3640543.3645199
- John Joon Young Chung and Eytan Adar. 2023. Artinter: AI-powered Boundary Objects for Commissioning Visual Arts. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference (Pittsburgh, PA, USA) (DIS '23)*. Association for Computing Machinery, New York, NY, USA, 1997–2018. doi:10.1145/3563657.3595961
- Cathlin V. Clark-Gordon, Nicholas D. Bowman, Alan K. Goodboy, and Alyssa Wright. 2019. Anonymity and Online Self-Disclosure: A Meta-Analysis. *Communication Reports* 32, 2 (May 2019), 98–111. doi:10.1080/08934215.2019.1607516
- Victoria Clarke and Virginia Braun. 2014. Thematic Analysis. In *Encyclopedia of Critical Psychology*. Springer, New York, NY, 1947–1952. doi:10.1007/978-1-4614-5583-7_311
- Jeff Conklin. 1997. Designing Organizational Memory: Preserving Intellectual Assets in a Knowledge Economy. (1997).
- Justin Cranshaw, Emad Elwany, Todd Newman, Rafal Kocielnik, Bowen Yu, Sandeep Soni, Jaime Teevan, and Andrés Monroy-Hernández. 2017. Calendar.help: Designing a Workflow-Based Scheduling Agent with Humans

- in the Loop. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 2382–2393. doi:10.1145/3025453.3025780
- [30] John W. Creswell and Vicki L. Plano Clark. 2017. *Designing and Conducting Mixed Methods Research*. SAGE Publications.
- [31] Catalina Danis and David Singer. 2008. A Wiki Instance in the Enterprise: Opportunities, Concerns and Reality. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work (CSCW '08)*. Association for Computing Machinery, New York, NY, USA, 495–504. doi:10.1145/1460563.1460642
- [32] Judith Donath, Karrie Karahalios, and Fernanda Viégas. 1999. Visualizing Conversation. *Journal of Computer-Mediated Communication* 4, 4 (June 1999), JCMC442. doi:10.1111/j.1083-6101.1999.tb00107.x
- [33] Paul Dourish and Victoria Bellotti. 1992. Awareness and Coordination in Shared Workspaces. In *Proceedings of the 1992 ACM Conference on Computer-Supported Cooperative Work - CSCW '92*. ACM Press, Toronto, Ontario, Canada, 107–114. doi:10.1145/143457.143468
- [34] Scott B. Droege and Jenny M. Hoobler. 2003. Employee Turnover And Tacit Knowledge Diffusion: A Network Perspective. *Journal of Managerial Issues* 15, 1 (2003), 50–64. jstor:40604414
- [35] Susan Dumais, Edward Cutrell, JJ Cadiz, Gavin Jancke, Raman Sarin, and Daniel C. Robbins. 2003. Stuff I've Seen: A System for Personal Information Retrieval and Re-Use. In *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Informaion Retrieval (SIGIR '03)*. Association for Computing Machinery, New York, NY, USA, 72–79. doi:10.1145/860435.860451
- [36] Kate Ehrlich and N. Shami. 2010. Microblogging Inside and Outside the Workplace. *Proceedings of the International AAAI Conference on Web and Social Media* 4, 1 (May 2010), 42–49. doi:10.1609/icwsm.v4i1.14020
- [37] Siamak Faridani, Ephrat Bitton, Kimiko Ryokai, and Ken Goldberg. 2010. Opinion Space: A Scalable Tool for Browsing Online Comments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. Association for Computing Machinery, New York, NY, USA, 1175–1184. doi:10.1145/1753326.1753502
- [38] Casey Fiesler, Jialun Jiang, Joshua McCann, Kyle Frye, and Jed Brubaker. 2018. Reddit Rules! Characterizing an Ecosystem of Governance. *Proceedings of the International AAAI Conference on Web and Social Media* 12, 1 (June 2018). doi:10.1609/icwsm.v12i1.15033
- [39] David Fono and Ron Baecker. 2006. Structuring and Supporting Persistent Chat Conversations. In *Proceedings of the 2006 20th Anniversary Conference on Computer Supported Cooperative Work (CSCW '06)*. Association for Computing Machinery, New York, NY, USA, 455–458. doi:10.1145/1180875.1180944
- [40] Andrea Forte and Amy Bruckman. 2005. Why Do People Write for Wikipedia? Incentives to Contribute to Open-Content Publishing. (2005).
- [41] Nataliya Galan. 2023. Knowledge Loss Induced by Organizational Member Turnover: A Review of Empirical Literature, Synthesis and Future Research Directions (Part II). *The Learning Organization: An International Journal* 30, 2 (Feb. 2023), 137–161. doi:10.1108/TLO-09-2022-0108
- [42] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. Retrieval-Augmented Generation for Large Language Models: A Survey. doi:10.48550/arXiv.2312.10997 arXiv:2312.10997 [cs]
- [43] R. Stuart Geiger and David Ribes. 2010. The Work of Sustaining Order in Wikipedia: The Banning of a Vandal. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (CSCW '10)*. Association for Computing Machinery, New York, NY, USA, 117–126. doi:10.1145/1718918.1718941
- [44] Jim Gemmell, Gordon Bell, and Roger Lueder. 2006. MyLifeBits: A Personal Database for Everything. *Commun. ACM* 49, 1 (Jan. 2006), 88–95. doi:10.1145/1107458.1107460
- [45] Jonathan Grudin. 1988. Why CSCW Applications Fail: Problems in the Design and Evaluation of Organizational Interfaces. In *Proceedings of the 1988 ACM Conference on Computer-Supported Cooperative Work - CSCW '88*. ACM Press, Portland, Oregon, United States, 85–93. doi:10.1145/62266.62273
- [46] Jonathan Grudin. 1994. Groupware and Social Dynamics: Eight Challenges for Developers. *Commun. ACM* 37, 1 (Jan. 1994), 92–105. doi:10.1145/175222.175230
- [47] Jonathan Grudin and Erika Shehan Poole. 2010. Wikis at Work: Success Factors and Challenges for Sustainability of Enterprise Wikis. In *Proceedings of the 6th International Symposium on Wikis and Open Collaboration (WikiSym '10)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/1832772.1832780
- [48] Jose A. Guridi, Angel Hsing-Chi Hwang, Duarte Santo, Maria Goula, Cristobal Cheyre, Lee Humphreys, and Marco Rangel. 2025. From Fake Perfects to Conversational Imperfects: Exploring Image-Generative AI as a Boundary Object for Participatory Design of Public Spaces. *Proc. ACM Hum.-Comput. Interact.* 9, 2, Article CSCW014 (may 2025), 33 pages. doi:10.1145/3710912
- [49] Mark Handel and James D. Herbsleb. 2002. What Is Chat Doing in the Workplace?. In *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work (CSCW '02)*. Association for Computing Machinery, New York, NY, USA, 1–10. doi:10.1145/587078.587080
- [50] Lester J. Holtzblatt, Laurie E. Damianos, and Daniel Weiss. 2010. Factors Impeding Wiki Use in the Enterprise: A Case Study. In *CHI '10 Extended Abstracts on Human Factors in Computing Systems (CHI EA '10)*. Association for Computing Machinery, New York, NY, USA, 4661–4676. doi:10.1145/1753846.1754208
- [51] Ellen Isaacs, Alan Walendowski, Steve Whittaker, Diane J. Schiano, and Candace Kamm. 2002. The Character, Functions, and Styles of Instant Messaging in the Workplace. In *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work (CSCW '02)*. Association for Computing Machinery, New York, NY, USA, 11–20. doi:10.1145/587078.587081
- [52] Farnaz Jahanbakhsh, Elnaz Nouri, Robert Sim, Ryan W. White, and Adam Fournay. 2022. Understanding Questions That Arise When Working with Business Documents. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2 (Nov. 2022), 341:1–341:24. doi:10.1145/3555761
- [53] Mahmood Jasim, Pooya Khaloo, Somn Wadhwa, Amy X. Zhang, Ali Sarvghad, and Narges Mahyar. 2021. CommunityClick: Capturing and Reporting Community Feedback from Town Halls to Improve Inclusivity. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW3, Article 213 (jan 2021), 32 pages. doi:10.1145/3432912
- [54] B. Kerr. 2003. Thread Arcs: An Email Thread Visualization. In *IEEE Symposium on Information Visualization 2003*. 211–218. doi:10.1109/INFVIS.2003.1249028
- [55] Soomin Kim, Jinsu Eun, Changhoon Oh, Bongwon Suh, and Joonhwan Lee. 2020. Bot in the Bunch: Facilitating Group Chat Discussion by Improving Efficiency and Participation with a Chatbot. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3313831.3376785
- [56] Soomin Kim, Jinsu Eun, Joseph Seering, and Joonhwan Lee. 2021. Moderator Chatbot for Deliberative Discussion: Effects of Discussion Structure and Discussant Facilitation. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1 (April 2021), 87:1–87:26. doi:10.1145/3449161
- [57] Sarah Kiniti and Craig Standing. 2013. Wikis as Knowledge Management Systems: Issues and Challenges. *Journal of Systems and Information Technology* 15, 2 (April 2013), 189–201. doi:10.1108/13287261311328895
- [58] Kalle Koivisto and Toni Taipalus. 2025. Pitfalls in Effective Knowledge Management: Insights from an International Information Technology Organization. *International Journal of Knowledge Management Studies* 16, 1 (2025), 1–26. doi:10.1504/IJKMS.2025.146083 arXiv:2304.07737 [cs]
- [59] Kathy E. Kram and Lynn A. Isabella. 1985. Mentoring Alternatives: The Role of Peer Relationships in Career Development. *The Academy of Management Journal* 28, 1 (1985), 110–132. doi:10.2307/256064 jstor:256064
- [60] Travis Kriplean, Jonathan Morgan, Deen Freelon, Alan Borning, and Lance Bennett. 2012. Supporting reflective public thought with considerit. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (Seattle, Washington, USA) (CSCW '12)*. Association for Computing Machinery, New York, NY, USA, 265–274. doi:10.1145/2145204.2145249
- [61] Philippe Laban, Jesse Vig, Marti Hearst, Caiming Xiong, and Chien-Sheng Wu. 2024. Beyond the Chat: Executable and Verifiable Text-Editing with LLMs. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*. Association for Computing Machinery, New York, NY, USA, 1–23. doi:10.1145/3654777.3676419
- [62] Jean Lave and Etienne Wenger. 1991. *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press.
- [63] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. 2017. *Research Methods in Human-Computer Interaction*. Morgan Kaufmann.
- [64] Yi-Chieh Lee, Naomi Yamashita, Yun Huang, and Wai Fu. 2020. "I Hear You, I Feel You": Encouraging Deep Self-disclosure through a Chatbot. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3313831.3376175
- [65] Paul M. Leonardi. 2015. Ambient Awareness and Knowledge Acquisition: Using Social Media to Learn ?Who Knows What? And ?Who Knows Whom? *MIS Quarterly* 39, 4 (2015), 747–762. jstor:26628649
- [66] T.C. Lethbridge, J. Singer, and A. Forward. 2003. How Software Engineers Use Documentation: The State of the Practice. *IEEE Software* 20, 6 (Nov. 2003), 35–39. doi:10.1109/MS.2003.1241364
- [67] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2021. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. doi:10.48550/arXiv.2005.11401 arXiv:2005.11401 [cs]
- [68] Kaitlin Mahar, Amy X. Zhang, and David Karger. 2018. Squadbox: A Tool to Combat Email Harassment Using Friendsourced Moderation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–13. doi:10.1145/3173574.3174160
- [69] Narges Mahyar, Michael R. James, Michelle M. Ng, Reginald A. Wu, and Steven P. Dow. 2018. CommunityCrit: Inviting the Public to Improve and Evaluate Urban Design Ideas through Micro-Activities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal QC, Canada) (CHI '18)*.

- Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3173574.3173769
- [70] Lena Mamykina, Bella Manoim, Manas Mittal, George Hripesak, and Björn Hartmann. 2011. Design Lessons from the Fastest Q&a Site in the West. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. Association for Computing Machinery, New York, NY, USA, 2857–2866. doi:10.1145/1978942.1979366
- [71] Britta Meixner, Matthew Lee, and Scott Carter. 2017. Chat2Doc: From Chats to How-to Instructions, FAQ, and Reports. In *Proceedings of the 2017 ACM Workshop on Multimedia-Based Educational and Knowledge Technologies for Personalized and Social Online Training (Mountain View, California, USA) (MultiEdTech '17)*. Association for Computing Machinery, New York, NY, USA, 27–30. doi:10.1145/3132390.3132392
- [72] Jonathan T. Morgan, Siko Bouterse, Heather Walls, and Sarah Stierch. 2013. Tea and Sympathy: Crafting Positive New User Experiences on Wikipedia. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work (CSCW '13)*. Association for Computing Machinery, New York, NY, USA, 839–848. doi:10.1145/2441776.2441871
- [73] Sean A. Munson. 2008. Motivating and enabling organizational memory with a workshop wiki. In *Proceedings of the 4th International Symposium on Wikis (Porto, Portugal) (WikiSym '08)*. Association for Computing Machinery, New York, NY, USA, Article 18, 5 pages. doi:10.1145/1822258.1822283
- [74] Bonnie A. Nardi, Steve Whittaker, and Erin Bradner. 2000. Interaction and Outaction: Instant Messaging in Action. In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work (CSCW '00)*. Association for Computing Machinery, New York, NY, USA, 79–88. doi:10.1145/358916.358975
- [75] Tahir M. Nisar, Guru Prabhakar, and Lubica Strakova. 2019. Social Media Information Benefits, Knowledge Management and Smart Organizations. *Journal of Business Research* 94 (Jan. 2019), 264–272. doi:10.1016/j.jbusres.2018.05.005
- [76] Ikujiro Nonaka. 1998. The Knowledge-Creating Company. In *The Economic Impact of Knowledge*. Routledge.
- [77] Wanda J. Orlikowski. 1992. Learning from Notes: Organizational Issues in Groupware Implementation. In *Proceedings of the 1992 ACM Conference on Computer-supported Cooperative Work - CSCW '92*. ACM Press, Toronto, Ontario, Canada, 362–369. doi:10.1145/143457.143549
- [78] Viral Pasad, Boyuan Wang, and Sang Won Lee. 2020. Understanding the Challenges of Online Group Chat for Productive Discourse at Scale. In *Companion Publication of the 2020 Conference on Computer Supported Cooperative Work and Social Computing (CSCW '20 Companion)*. Association for Computing Machinery, New York, NY, USA, 365–369. doi:10.1145/3406865.3418335
- [79] Jeffrey M. Perkel. 2017. How Scientists Use Slack. *Nature* 541, 7635 (Jan. 2017), 123–124. doi:10.1038/541123a
- [80] Volkmar Pipek and Volker Wulf. 2003. Pruning the Answer Garden: Knowledge Sharing in Maintenance Engineering. In *ECSCW 2003*, Kari Kuutti, Eija Helena Karsten, Geraldine Fitzpatrick, Paul Dourish, and Kjeld Schmidt (Eds.). Springer Netherlands, Dordrecht, 1–20. doi:10.1007/978-94-010-0068-0_1
- [81] Volkmar Pipek and Volker Wulf. 2009. Infrastructuring: Toward an Integrated Perspective on the Design and Use of Information Technology. *Journal of the Association for Information Systems* 10, 5 (May 2009). doi:10.17705/1jais.00195
- [82] Michael Polanyi. 1997. The Tacit Dimension. In *Knowledge in Organisations*. Routledge.
- [83] Julien Porquet, Sitong Wang, and Lydia B. Chilton. 2025. Copying style, Extracting value: Illustrators' Perception of AI Style Transfer and its Impact on Creative Labor. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 202, 16 pages. doi:10.1145/3706598.3713854
- [84] Matthew Richardson and Ryan W. White. 2011. Supporting Synchronous Social Q&a throughout the Question Lifecycle. In *Proceedings of the 20th International Conference on World Wide Web (WWW '11)*. Association for Computing Machinery, New York, NY, USA, 755–764. doi:10.1145/1963405.1963511
- [85] Yvonne Rogers. 2011. Interaction Design Gone Wild: Striving for Wild Theory. *Interactions* 18, 4 (July 2011), 58–62. doi:10.1145/1978822.1978834
- [86] Rod D. Roscoe and Michelene T. H. Chi. 2007. Understanding Tutor Learning: Knowledge-Building and Knowledge-Telling in Peer Tutors' Explanations and Questions. *Review of Educational Research* 77, 4 (Dec. 2007), 534–574. doi:10.3102/0034654307309920
- [87] Joseph Seering, Michal Luria, Geoff Kaufman, and Jessica Hammer. 2019. Beyond Dyadic Interactions: Considering Chatbots as Community Members. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, Glasgow Scotland Uk, 1–13. doi:10.1145/3290605.3300680
- [88] Simran Setia, S. R.S. Iyengar, and Amit Arjun Verma. 2020. QWiki: Need for QnA & Wiki to Co-exist. In *Proceedings of the 16th International Symposium on Open Collaboration*. ACM, Virtual conference Spain, 1–12. doi:10.1145/3412569.3412576
- [89] Donghoon Shin, Soomin Kim, Ruoxi Shang, Joonhwan Lee, and Gary Hsieh. 2023. IntroBot: Exploring the Use of Chatbot-assisted Familiarization in Online Collaborative Groups. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3544548.3580930
- [90] Slack. 2020. Distance Learning Thrives in Slack. <https://slack.com/blog/collaboration/distance-learning-in-slack>.
- [91] Susan Leigh Star and James R. Griesemer. 1989. Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907–39. *Social Studies of Science* 19, 3 (Aug. 1989), 387–420. doi:10.1177/03063128919003001
- [92] Susan Leigh Star and Anselm Strauss. 1999. Layers of Silence, Arenas of Voice: The Ecology of Visible and Invisible Work. *Computer Supported Cooperative Work (CSCW)* 8, 1 (March 1999), 9–30. doi:10.1023/A:1008651105359
- [93] John Suler. 2004. The Online Disinhibition Effect. *CyberPsychology & Behavior* 7, 3 (June 2004), 321–326. doi:10.1089/1094931041291295
- [94] Yunting Sun, Diane Lambert, Makoto Uchida, and Nicolas Remy. 2014. Collaboration in the Cloud at Google. In *Proceedings of the 2014 ACM Conference on Web Science (WebSci '14)*. Association for Computing Machinery, New York, NY, USA, 239–240. doi:10.1145/2615569.2615637
- [95] James Tam and Saul Greenberg. 2006. A Framework for Asynchronous Change Awareness in Collaborative Documents and Workspaces. *International Journal of Human-Computer Studies* 64, 7 (July 2006), 583–598. doi:10.1016/j.ijhcs.2006.02.004
- [96] Naama Tepper, Anat Hashavit, Maya Barnea, Inbal Ronen, and Lior Leiba. 2018. Collabot: Personalized Group Chat Summarization. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining (WSDM '18)*. Association for Computing Machinery, New York, NY, USA, 771–774. doi:10.1145/3159652.3160588
- [97] Trevor T. Tuma, John D. Adams, Benjamin C. Hultquist, and Erin L. Dolan. 2021. The Dark Side of Development: A Systems Characterization of the Negative Mentoring Experiences of Doctoral Students. *CBE—Life Sciences Education* 20, 2 (June 2021), ar16. doi:10.1187/cbe.20-10-0231
- [98] Christian Wagner. 2004. Wiki: A Technology for Conversational Knowledge Management and Group Collaboration. *Communications of the Association for Information Systems* 13 (2004). doi:10.17705/1CAIS.01319
- [99] James P. Walsh and Gerardo Rivera Ungson. 1991. Organizational Memory. *The Academy of Management Review* 16, 1 (1991), 57–91. doi:10.2307/258607
- [100] Dakuo Wang, Judith S. Olson, Jingwen Zhang, Trung Nguyen, and Gary M. Olson. 2015. DocuViz: Visualizing Collaborative Writing. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. Association for Computing Machinery, New York, NY, USA, 1865–1874. doi:10.1145/2702123.2702517
- [101] Dakuo Wang, Haoyu Wang, Mo Yu, Zahra Ashktorab, and Ming Tan. 2022. Group Chat Ecology in Enterprise Instant Messaging: How Employees Collaborate Through Multi-User Chat Channels on Slack. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW1 (April 2022), 94:1–94:14. doi:10.1145/3512941
- [102] Steve Whittaker. 2011. Personal Information Management: From Information Consumption to Curation. *Annual Review of Information Science and Technology* 45, 1 (Jan. 2011), 1–62. doi:10.1002/aris.2011.1440450108
- [103] Steve Whittaker and Candace Sidner. 1996. Email Overload: Exploring Personal Information Management of Email. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems Common Ground - CHI '96*. ACM Press, Vancouver, British Columbia, Canada, 276–283. doi:10.1145/238386.238530
- [104] Amy X. Zhang and Justin Cranshaw. 2018. Making Sense of Group Chat through Collaborative Tagging and Summarization. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (Nov. 2018), 1–27. doi:10.1145/3274465
- [105] Amy X. Zhang, Lea Verou, and David Karger. 2017. Wikum: Bridging Discussion Forums and Wikis Using Recursive Summarization. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, Portland Oregon USA, 2082–2096. doi:10.1145/2998181.2998235
- [106] Shiyue Zhang, Asli Celikyilmaz, Jianfeng Gao, and Mohit Bansal. 2021. EmailSum: Abstractive Email Thread Summarization. doi:10.48550/arXiv.2107.14691 arXiv:2107.14691 [cs]
- [107] Dora Zhao, Diyi Yang, and Michael S. Bernstein. 2025. Knoll: Creating a Knowledge Ecosystem for Large Language Models. doi:10.48550/arXiv.2505.19335 arXiv:2505.19335 [cs]
- [108] Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan Awadallah, Asli Celikyilmaz, Yang Liu, Xipeng Qiu, and Dragomir Radev. 2021. QMSum: A New Benchmark for Query-based Multi-domain Meeting Summarization. doi:10.48550/arXiv.2104.05938 arXiv:2104.05938 [cs]

A System Prompts

This appendix provides the complete system prompts used in CHOIR for reproducibility. Model configuration details are provided in Section 4.5.2.

A.1 Question Answering Prompt

The following prompt is used for document-grounded Q&A (Feature 1). User names are replaced with pseudonyms before being sent to the API for privacy protection.

You are CHOIR, a helpful AI assistant for [Organization Name].

Think of yourself as a knowledgeable senior student or friendly professor who's always ready to help with questions.

I have access to the organization's documentation and knowledge base, so I can help you find information and provide guidance based on what we have documented.

Organization Information:

- Organization: [Organization Name]
- About: [Organization Description]
- Today's date: [Current Date]
- Workspace: [Workspace Name]

IMPORTANT: First, determine if you can answer the user's question based ONLY on the provided references. Do NOT use your general knowledge or make assumptions beyond what's explicitly stated in the references.

Then respond with a JSON object containing:

1. ``canAnswer``: true/false - whether the references contain sufficient information to answer the question
2. ``response``: your answer based on the references, or a friendly explanation that you couldn't find the information

Guidelines for answering:

- Be friendly, approachable, and helpful - like a senior colleague who genuinely wants to help
- Answer ONLY based on the provided references below - do NOT add general knowledge
- If multiple references contain conflicting information, prioritize the first reference in the list
- If you cannot answer based on the references, encourage the user to ask others or start a discussion to help improve our documentation
- When you cannot fully answer a question, mention which related references you found and what information they contain
- Use a warm, academic tone - professional but not overly formal

==== REFERENCES ====

[Retrieved document chunks]

==== CONVERSATION HISTORY ====

[Previous messages in the thread]

==== CURRENT QUESTION ====

[User's question]

Analyze whether you can answer based on the documentation and provide your response as JSON.

A.2 Knowledge Extraction Prompt

The following prompt is used to extract documentable knowledge from conversations (Feature 3). User names are replaced with pseudonyms before being sent to the API.

You are CHOIR, a documentation specialist.

Extract organizational knowledge from the conversation that would be valuable for future reference.

Only extract information that establishes policies, procedures, or reusable knowledge for the organization. Do NOT extract personal preferences, individual decisions, or casual conversation.

Start with a descriptive markdown section title (using heading format with [Topic Name]),

then write the information in natural paragraph format.

Only include facts that are directly stated in the conversation - do not add explanations, interpretations, or implications.

Do not include personal conversation details like who said what. Always preserve any URLs mentioned.

Conversation:

[Conversation messages]

What information is shared in the conversation that should be documented?

A.3 Document Update Prompt

The following prompt is used to integrate extracted knowledge into existing documentation (Feature 4).

You are a document editor. Improve existing content by integrating the provided knowledge.

Existing content type: [paragraph or list]

Rules:

- PRIORITIZE updating/replacing existing content when knowledge provides better, more accurate, or more comprehensive information
- If knowledge contradicts existing content, prefer the knowledge (assume it's more current/accurate)
- If knowledge complements existing content without contradiction, add it in matching format
- If knowledge provides more specific details about existing points, merge them into improved versions
- Preserve all URLs from the knowledge
- Use single-level lists only (no nested bullets)
- No headings or section titles
- Return original only if knowledge adds no meaningful value

Approach: Update first, then add if needed. Create the most accurate and comprehensive version.

Wrap your response in <markdown> tags.

File: [File Name] - Section: [Section Name]

Existing content:

<markdown>

[Current section content]

</markdown>

Knowledge to integrate:

[Extracted knowledge]