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

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

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

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

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

Sang Won Lee
sangwonlee@vt.edu
Virginia Tech
Blacksburg, Virginia, USA

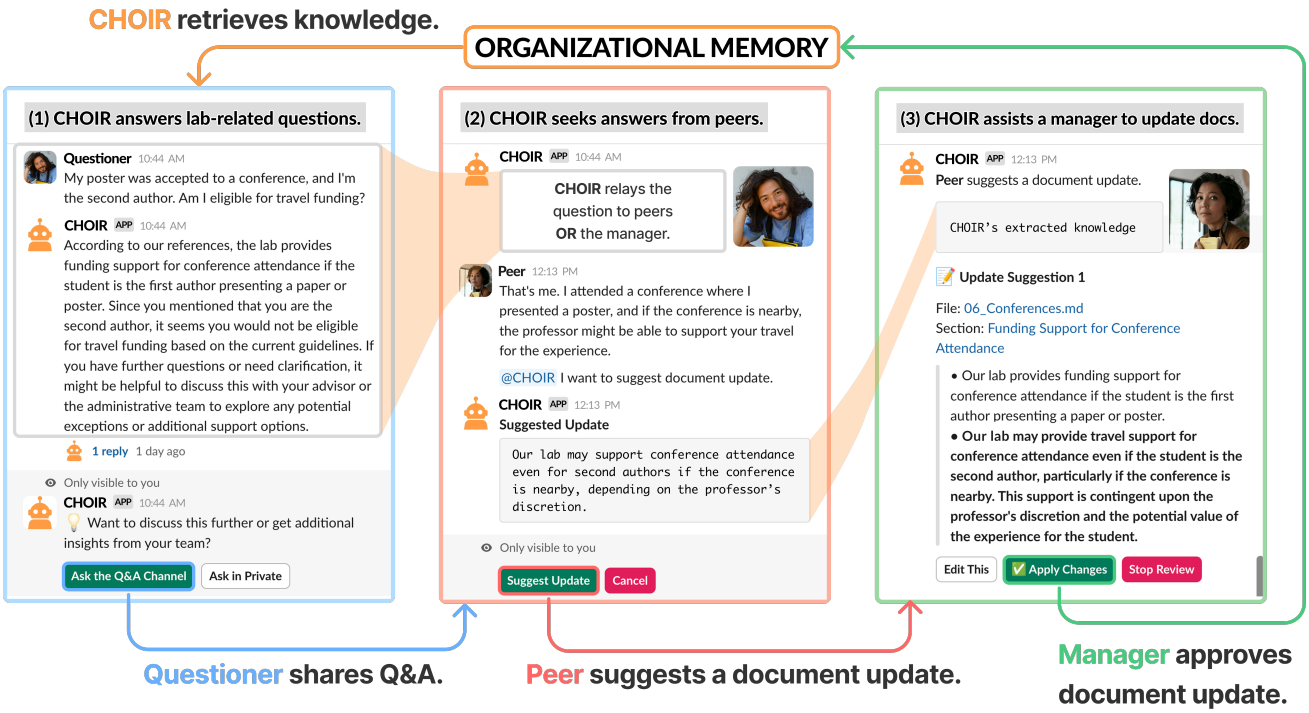

Figure 1: CHOIR's workflow across different communication channels

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

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 [61, 81]. Slack alone reported over 1.2 million users across more than 3,000 higher education institutions during the COVID-19 pandemic in 2020 [70]. 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, 42, 81]. 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 [60, 84, 85]. 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, 79]. Some research labs implement such organizational knowledge systems using collaborative documents such as Google Docs [74] or group wikis [39, 78], often in the form of lab handbooks or policy documents. However, maintaining these resources requires substantial manual effort [41, 46, 52, 59]. In practice, only professors find it difficult to consistently maintain such documents [72, 74]—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 [29, 35]. Even when central documents exist, students often struggle to read them entirely or locate the right information when they need it [15, 82]. 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 goal of CHOIR is to give students access to organizational memory within their communication platform, i.e., Slack, and to allow professors to maintain documents within the platform, leveraging communication that they have in the platform to retrieve, extract, and preserve knowledge. To accomplish this, 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 establishes 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 [77].

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 organizational memory 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 transfering and retreiving knowledge amont 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 confirm that CHOIR effectively helped students obtain answers available in the documents. In addition, it encouraged 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, 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 was noted for answering other people's questions. Based on the findings, we discuss the social and technical aspects of how organizational memory can be effectively integrated into chat communication platforms for collective contributions, while taking into account students' hesitation and fear of updating organizational memory documents.

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) Empirical evidence demonstrating the effectiveness of a chat-based organizational memory system, along with insights from a month-long deployment that reveal real-world challenges and design implications.

2 Related Work

In this section, 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 [79], 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 [59, 72], organizational power dynamics [38], and the scattered nature of informal knowledge that often remains undocumented [24]. 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 [62]. 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 [71], requiring recontextualization for effective use [4].

With the advancement of web technologies, wiki systems demonstrated this potential through collaborative editing [78], but organizational deployments encountered familiar social barriers. Danis [26] deployed ResearchWiki in a research organization but found limited participation due to competitive culture and modification anxiety. Subsequent research identified systemic barriers: additional effort requirements, information sensitivity concerns, discomfort with unfinished work publication, and openness anxiety [41]. The collaborative editing model effective in public contexts created stress in organizational environments where authorship and accountability mattered [39]. Alternative platforms like Google Docs achieved better adoption through simplified user experiences that naturally facilitated role differentiation without wiki complexity [74], with supporting technologies focusing on transparent collaboration history to navigate social complexities [80].

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 [64]. The challenge is transforming conversational knowledge into explicit, reusable forms while preserving contextual richness [58]. 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 [83]. Instant messaging further demonstrated how informal channels could become knowledge repositories for relationship maintenance and work coordination [40, 42, 56].

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 [14], while corporate social networking enabled help requests and work coordination through informal interactions [30]. Research showed these platforms improved employees' metaknowledge awareness through social mediation [51, 54, 57]. 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 [60]. In research environments, Slack served as both a knowledge-sharing hub and a culture formation space [61], with conversational data proving valuable for tasks like API recommendations from developer discussions [20].

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 [27, 31, 43], but required manual navigation. More sophisticated systems automated knowledge extraction through annotation [7, 33] and collaborative summarization [85]. Chat-based knowledge management systems achieved greater sophistication with Tilda, which automatically summarizes tagged conversations [84], and tools combining topic learning with personalized summaries [76]. QWiki demonstrated that Q&A and documentation approaches could be mutually reinforcing in MOOC environments [68]. 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 [76] and Tilda [84], 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 [44, 45, 67]. Research showed users felt more comfortable sharing knowledge with chatbots than humans, reducing social barriers [50].

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 [86, 88] 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 [53], with RAG systems searching existing documents and incorporating relevant information into responses, improving accuracy and reducing hallucination [36]. Beyond retrieval and summarization, LLMs enable dynamic document updating based on conversational knowledge [48] and facilitate complex group interactions [21].

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. 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 focused on recruiting professors and university students to understand the needs and challenges of organizational memory. 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 professors who directed research labs with more than 4 Ph.D. students and had experience or interest in maintaining organizational memory for their research lab. We adapted the eligibility criteria for recruiting students accordingly. From the survey responses, we contacted and recruited 14 participants: 11 faculty members (P1-P11) and 4 graduate students (S1-S4) representing various academic disciplines, primarily from Computer Science 7) but also from other disciplines: Industrial Engineering (2), Information Science (1), and Human Development and Family Science (1). 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). We also recruited 4 graduate students from 4 different disciplines. 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 internal review board of the author's university([IRB-XXXXX]).

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 AI[1]). After familiarization with the data, the second author performed a thematic analysis [23], 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

[1] https://sonix.ai/

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)

These challenges are consistent with prior work on knowledge management; studies have shown the high cost of creating and maintaining knowledge artifacts [2, 59]. Researchers have noted that documents tend to become outdated quickly, reducing their usefulness over time [38, 52].

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 Retreival.*** 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)

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, 63]. 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 [47, 49, 66].

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. 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. 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 [34, 55, 78]. 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, 37]. 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.

These design goals guided the design and implementation of CHOIR, which we describe in the following section.

4 CHOIR: Chat-based Helper for Organizational Intelligence Repository

Grounded in our formative study, we designed CHOIR to enable students to obtain answers directly within Slack and to transform everyday conversations into a maintained organizational memory. CHOIR integrates four features essential for intelligent knowledge retrieval and documentation: **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, and the underlying implementation utilizes LLM-powered multi-agents that update and access documents stored in GitHub via GitHub and Slack API.

When a message is sent to CHOIR—either through a direct message or by mentioning its handle (i.e., @CHOIR) in a channel any channel that it is part of, the message will be classified as either a question or an update request, corresponding to the two start nodes shown in Figure 2. The four features in CHOIR can establish

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 to it and retrieve information from the document in any Slack channel. When a question is submitted, CHOIR answers by composing a response grounded in those sources (**DG2**). Members can submit questions to CHOIR either by sending a direct message or by mentioning it in a Slack channel.

Internally, CHOIR retrieves topically relevant chunks and performs generation conditioned on the member's question and the retrieved text. 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 and decide whether to continue the discussion. Figure 3 highlights four UI elements: the member's question (A in Figure 3, submitted via direct message in this example), CHOIR's answer in the main message (B in Figure 3), and an evidence panel in the thread that lists the specific chunks used with anchored links (C in Figure 3). These references encourage the questioner to check out the original documents available in a GitHub repository, thereby getting familiar with the organizational memory.

If a question is asked in a public channel, all members can naturally see the CHOIR's answers, establishing a common ground that supports further discussion or follow-up questions and thereby facilitates communication among members (**DG3**).

After CHOIR provides a response, users are given the option to continue the discussion with their peers if they have follow-up questions, regardless of whether the original question was fully answered. A private message that is only visible to the questioner (D in Figure 3) provide two buttons with which the questioner can choose either to post the Q&A exchange to the designated Q&A channel for lab-wide follow-up or to send it privately to selected peers or the manager via DM, creating a natural handoff to the next feature.

4.2 Q&A Sharing for Follow-up Discussion

CHOIR encourages questioners to relay a question—whether answered or unanswered— to a Slack channel or to a specific person via a group chat, thereby encouraging communication among members (**DG3**). After CHOIR posts an answer, the questioner sees a share banner offering two options to share Q&A either in a public channel (Green button in Figure 3-D) or send a direct message ("Ask in Private"D, Figure 3)) to a private group chat with selected members, CHOIR, and the questioner.

Questioners can specify more options in 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 in Figure 4) 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 group chat conversation to preserve anonymity among recipients. The preview panel (C in Figure 4) shows exactly how the question will be shared, either as a channel post or as an outgoing DM.

The full *Q&A exchange* will be shared as shown in the left half of the Figure 5, and it appears as a message that includes the question and CHOIR's answer, along with the questioner's comment and profile image (A in Figure 5). Other members or the manager can then add follow-up comments (B in Figure 5). The same posting format also applies when the Q&A is shared in a private DM with selected people via the "Ask in Private" modal. While answers are posted in a public channel or if a questioner is part of the group chat, the questioner can see how others responded to the relayed question. If a message is shared anonymously in a private DM, CHOIR forwards any follow-up replies to the original questioner without revealing their identity. CHOIR is designed to encourage Q&A sharing so that it supplements rather than replaces interpersonal communication in the lab, enabling members to collectively answer questions. More importantly, when such discussions lead to new or missing knowledge, Q&A sharing can provide the entry point for the next feature, *Knowledge Extraction from Conversation*.

4.3 Knowledge Extraction from Conversation

CHOIR turns any Slack messages into a piece of reusable knowledge from conversation (DG1), allowing any members to participate in organizational memory maintenance(**DG4**). The follow-up discussion after a Q&A exchange is shared is likely to have answers to the question that the original questioner could not gain from CHOIR, which may be worth documenting.

When any lab member sees new knowledge worth keeping, they can summon CHOIR to request a documentation update by mentioning its handle, i.e., @CHOIR in a message (C in Figure 5). At that point, CHOIR examines the previously shared Q&A together with the conversation up to the mention, identifies *documentable* content (filtering out personal or social chatter), and drafts a suggested update (D in Figure 5). The requester can then edit this draft and press *Suggest Update* to forward it to the manager (E in Figure 5). CHOIR can be summoned anywhere in Slack. For example, if a professor shared a writing tip as part of giving feedback to a student, a professor can summon CHOIR to document the knowledge, which is represented as a new start node in the bottom right corner of Figure 2. This feature simplifies the document update process,

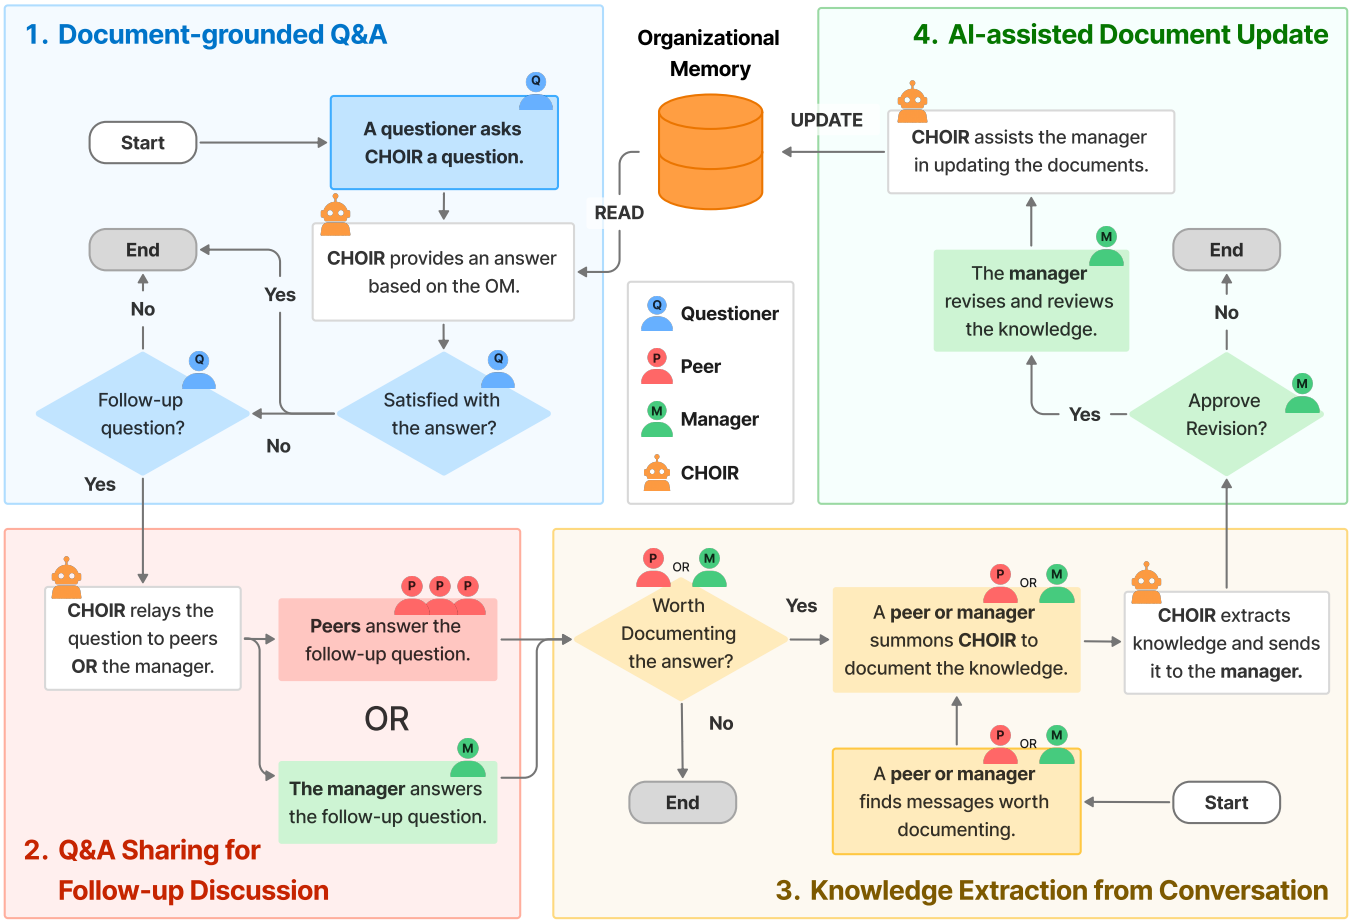

Figure 2: A questioner may ask CHOIR and receive a document-grounded answer, optionally share the Q&A exchange for broader discussion, and from these interactions, knowledge can be extracted and reviewed by the manager. Approved updates are written back to the organizational memory repository, which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

supporting DG1, by allowing any member to initiate updates and refine the knowledge rather than placing the entire responsibility on the manager. It lowers the barrier to documenting knowledge, thereby facilitating all lab members' participation in organizational memory practices (DG4).

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 in Figure 5-E, 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, ensures that only accurate information is documented, 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 can be multiple places to update in a document, or stop the review entirely (C in Figure 6). 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 in Figure 6). Applied updates—whether to an existing file, a new file, or a new section—are written back to

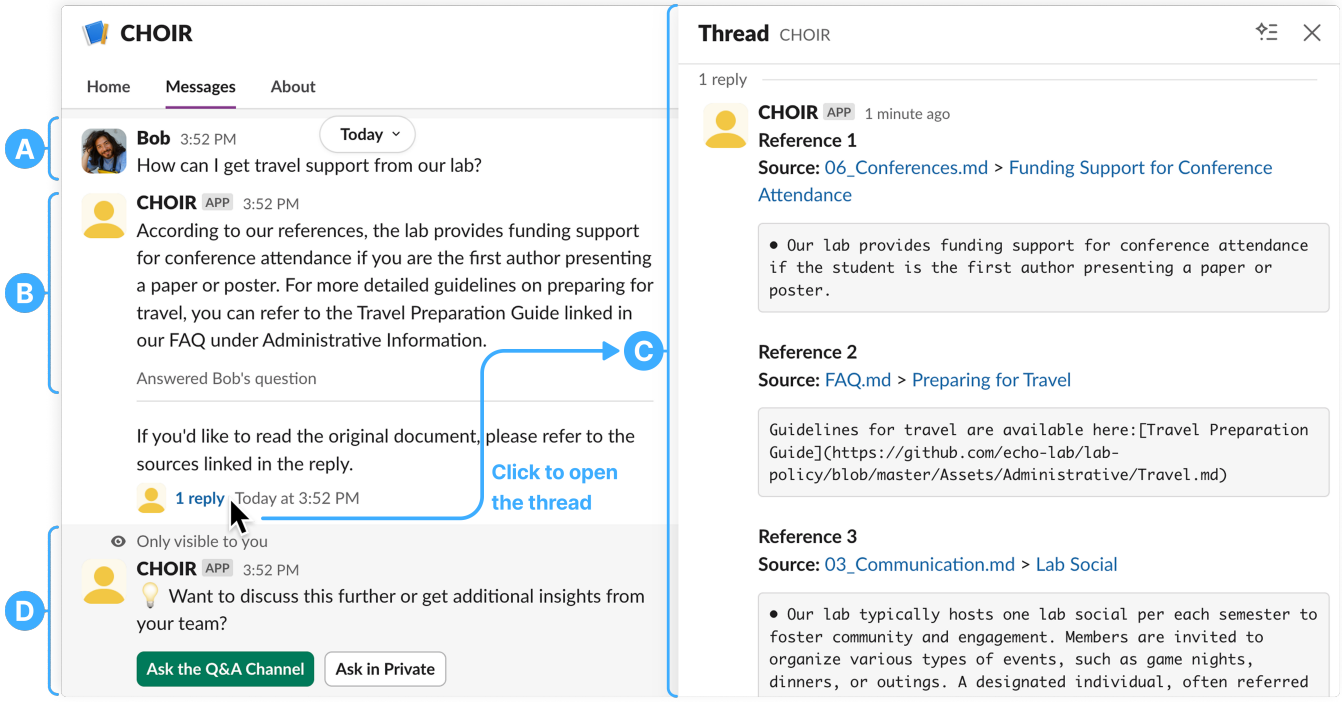

Figure 3: Document-grounded Q&A with CHOIR (DM example). Bob (questioner) asks CHOIR a question via direct message, CHOIR returns a document-grounded answer, and the thread is opened to view cited references.

the repository, and CHOIR posts a message to the original channel acknowledging that the update has been finished.

This feature implements **DG1** by simplifying the document update process and automating many of the steps involved: (1) recommending where an update is needed, (2) proposing updated content and checking for conflicts with existing material, and (3) committing the changes at the point where new knowledge becomes available. In this way, managers can update documents directly within the flow of communication in Slack, without switching tools. While CHOIR automates much of the update process, it also preserves flexibility and managerial autonomy by allowing managers to edit the proposed content or make changes directly on GitHub. For example, if an update requires more than adding or revising text—such as changing headings or reorganizing the structure—CHOIR provides hyperlinks that open the file editor directly ("here" in D in Figure 6). Overall, this AI-assisted update feature enables managers to keep documents up to date with minimal effort while retaining full control over what becomes part of the organizational memory.

4.5 System Implementation

4.5.1 Architecture. CHOIR uses Slack as the front-end client where members interact with the system. The backend server is built with the JavaScript-based Slack Bolt[2] framework on top of Express. In this setup, the server acts as an HTTP service that receives events from the Slack client as requests and returns messages as HTTP responses. The server is also connected to the lab's GitHub repository through the GitHub REST API[3], which serves as the canonical backend store for documentation. All documents in the repository are maintained in Markdown format, providing a lightweight and widely supported format. This stack allows Slack to serve as the entry point for everyday interactions, while GitHub provides persistent storage and version control for organizational memory.

4.5.2 Knowledge Retrieval and Generation. 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. 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 exact prompts used for classification and generation are included in the Appendix for transparency and reproducibility.

5 Field Study

To address RQ2, understanding how lab members engage with and experience a chatbot-driven organizational memory system, we conducted a month-long field deployment of CHOIR in the Slack

[2]https://api.slack.com/bolt

[3]https://docs.github.com/en/rest

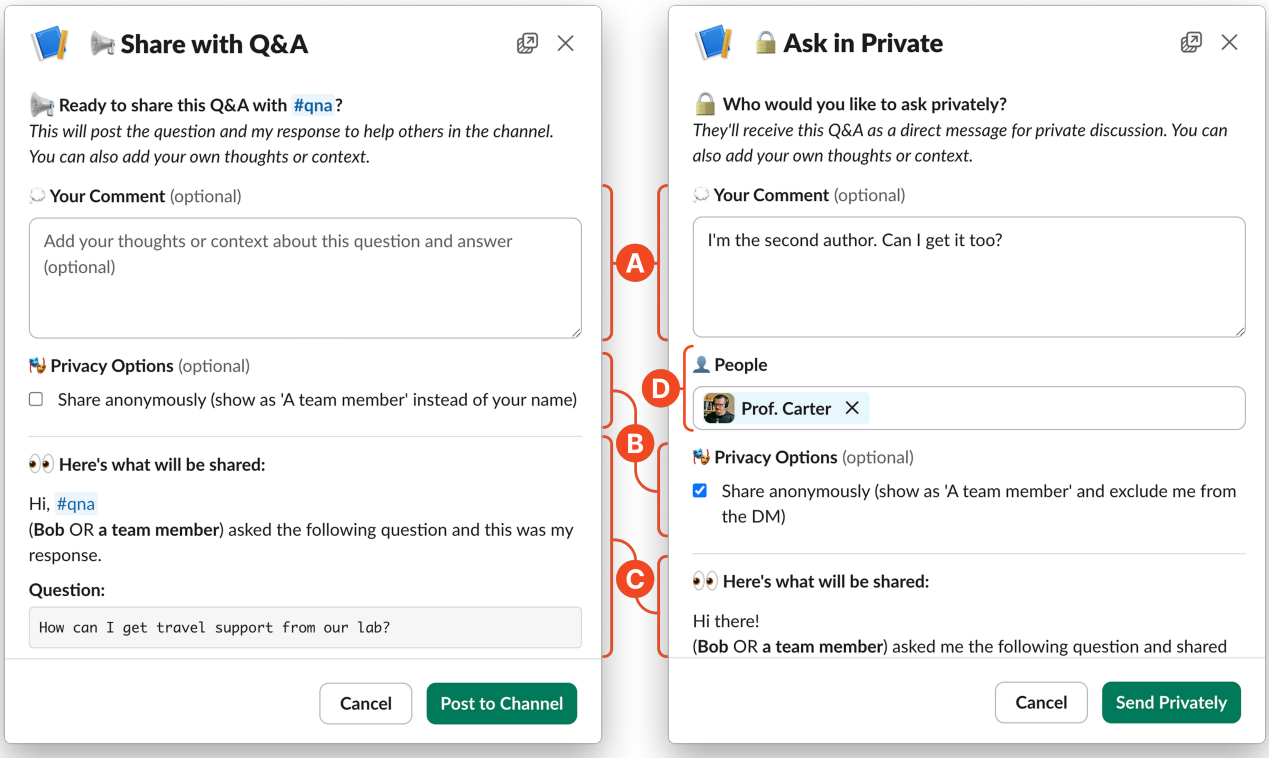

Figure 4: Two share modals in CHOIR. Left: "Share with Q&A" for posting a Q&A exchange to the designated channel. Right: "Ask in Private" for sending the same exchange via DM; in this example, the questioner adds a follow-up, the manager is pre-selected in the People picker, and anonymous sharing is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

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 system usage logs with user interviews [25], 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 [65]. This approach is particularly appropriate given our research focus on understanding user engagement and experience with a sociotechnical intervention in organizational contexts, where authentic workplace dynamics are crucial for meaningful insights.

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 lab directors initially agreed to participate; however, one later declined full deployment due to concerns about potential human-subject data leakage caused by human error. Upon confirmation from the director, we advertised the study in their slack worksapce and recruited the students from the lab, endorsed by the lab director. In total, four research laboratories participated in the user study, comprising 21 participants: three lab directors and 18 students. We refer to each unit and participant based on the labels as specified in Table 1. Since our study was conducted in the Summer of 2025, each lab had remote incoming graduate students joining during this period, allowing us to observe newcomer experiences.

All participants provided informed consent for interviews, system usage monitoring, and interaction log collection. The study was approved by our institution's IRB. As compensation, lab directors received $40 and students received $20 upon completing the interviews.

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

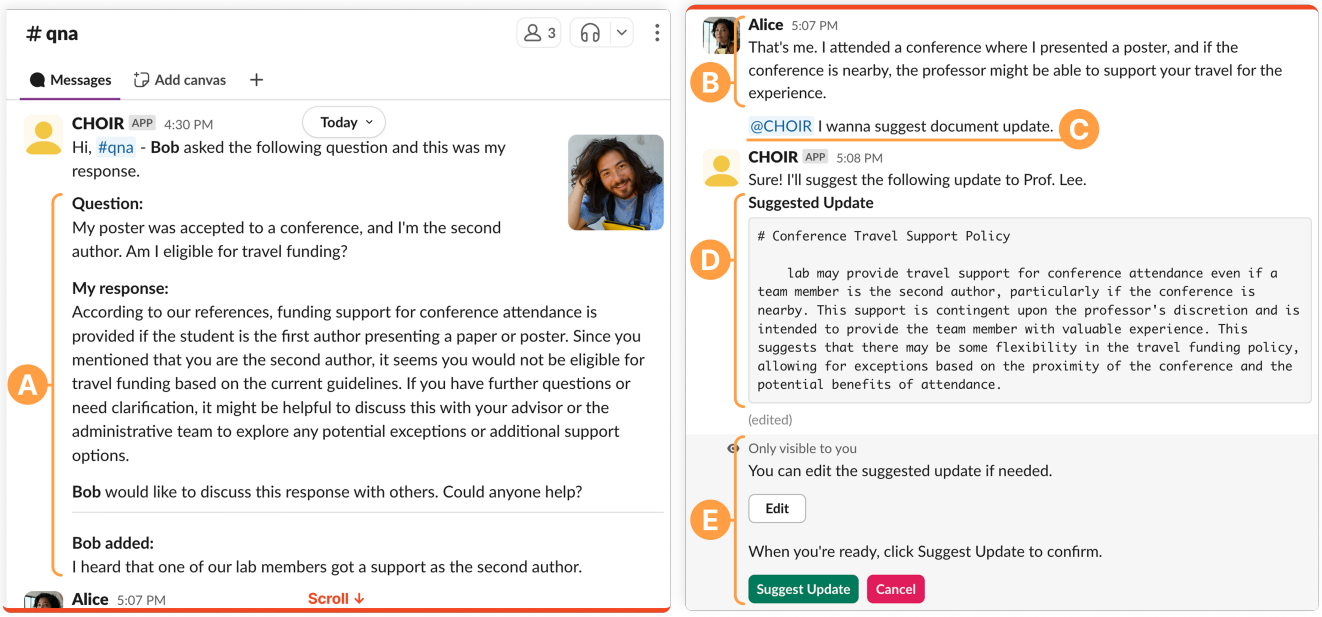

Figure 5: Example of knowledge extraction after a Q&A is shared. The Q&A exchange has been posted to the designated channel (#qna) through the Share with Q&A modal. A colleague (Alice) adds a follow-up comment and then mentions CHOIR to request an update, prompting CHOIR to extract documentable knowledge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

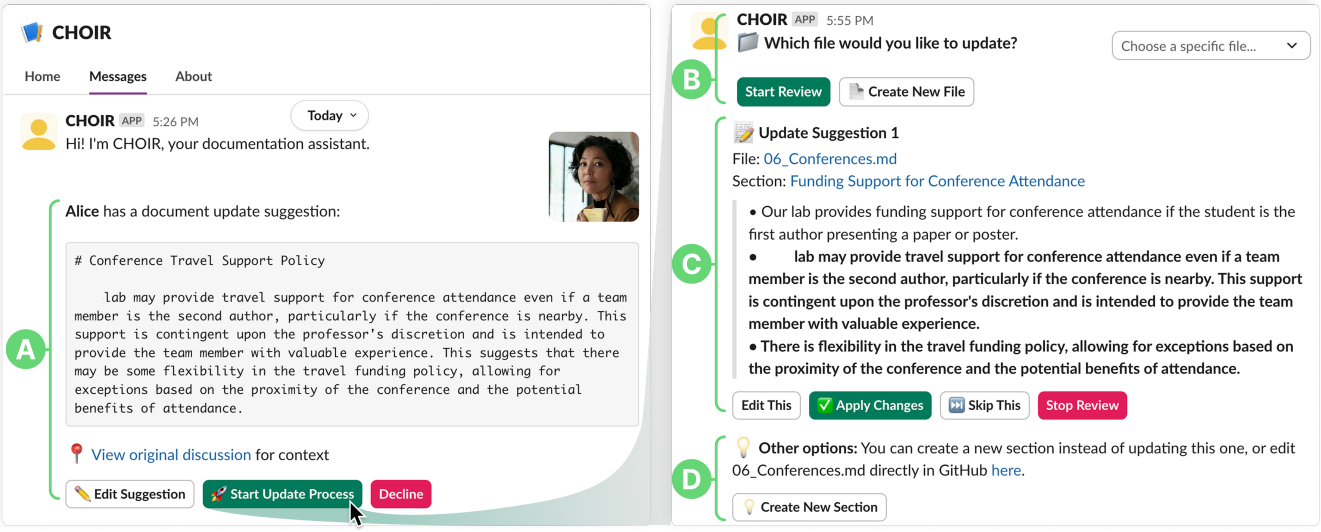

Figure 6: Manager-facing document update workflow in CHOIR. Left: a member's suggestion (from Alice) arrives as a Slack DM with attribution and a link to the original discussion. Right: after starting the update process, the manager reviews the proposed change in context and decides how to incorporate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

only documents that could be made publicly available. D1 and D3 did not have any documents to start with, while D2 had a 1,650-word-long document that explained the lab director's mentoring plan. To avoid a cold-start problem, we generated some documents based on information from each lab's website and the graduate handbook that each university has, which included requirements, milestones, and timelines for graduate students. The research team converted these documents into markdown format, uploaded them to a public GitHub repository for each lab, and granted lab directors write access and students read access.

Table 1: Participants in the field study: each lab's director and members, including newcomer status.

Lab #.	Department (Research Area)	Lab Age	Label (Gender)	Status
L1	Computer Science (Software Engineering)	4 yrs	D1 (M)	Director
			S1-1 (F)	New student
			S1-2 (F)	Existing student
			S1-3 (M)	Existing student
			S1-4 (M)	Existing student
L2	Computer Science (Human-Computer Interaction)	4 yrs	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	Industrial Eng. (Physical Ergonomics)	4 yrs	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	Computer Science (Human-Computer Interaction)	7 yrs	D4 (M)	Director (Author)
			S4-1 (M)	Existing student
			S4-2 (M)	Existing student
			S4-3 (M)	Existing student
			S4-4 (F)	Existing student

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. 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 [16]. 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. We also compared interview themes with patterns in the system logs to understand not only what participants did, but also why they made particular choices and how they experienced the system.

6 Field Study Results

Overall, CHOIR successfully 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 successfully 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.*" S4-2 emphasized the system's accuracy: "*if I asked a question and if it was there, you know, it would always give me a right answer or like summary of the documents.*" 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.

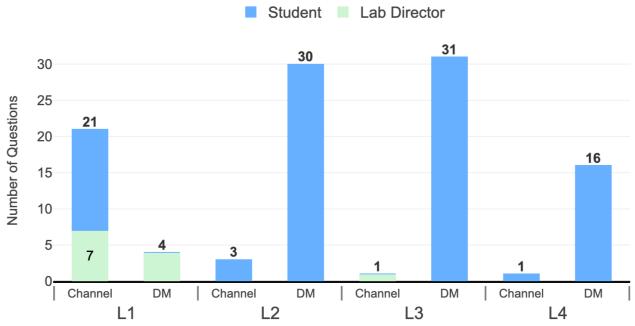

Figure 7: Distribution of questions asked to CHOIR across four research labs during the one-month field study. The graph shows the number of questions submitted through public channels versus direct messages (DM), broken down by students (blue) and lab directors (green). L1 shows a distinctive pattern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

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 21 questions publicly compared to only 5 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 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, successfully 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.*" Students also appreciated the value of documentation, with S4-1 noting that "*documentation is good. And another kind of documentation is all of our Slack messages. And it really sucks that those go away after three months or whatever. Because I feel like there's actually a lot of organizational knowledge that is embedded within.*"

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.*" All directors who either received update requests from students or initiated updates themselves successfully incorporated them into the lab documentation, demonstrating 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 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 validated 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

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 Added	Words Deleted
L1	Direct Update	19	5	1153	63
	CHOIR	0	0	0	0
L2	Direct Update	1	1	1	1
	CHOIR	4	1	261	0
L3	Direct Update	0	0	0	0
	CHOIR	1	1	109	1
L4	Direct Update	7	2	122	29
	CHOIR	6	2	348	2
Subtotal	Direct Update	**27**	**8**	**1276**	**93**
	CHOIR	**11**	**4**	**718**	**3**
Total	–	**38**	**12**	**1994**	**96**

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 reluctnace that they have for contributing to the organizaionl 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: `textitif 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

In this section, we reflect on our findings to understand the broader implications for designing AI-mediated organizational memory systems. Our field study revealed several key tensions and opportunities that extend beyond the specific technical features of CHOIR to more fundamental questions about how knowledge management

systems can effectively support collaborative work in research organizations.

7.1 Privacy-Awareness Tension in Organizational Memory Systems

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, lab directors had limited visibility into knowledge gaps in the documentation. Directors expressed both their willingness and desire 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 [28, 75], yet our study shows that AI-mediated knowledge retrieval can suffer from the same limitation.

This privacy-awareness 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 [22, 73]. 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 [13].

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 can 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 significant 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. As an alternative, participants also recognized significant value in case-based information and partial policies that could provide personalized guidance even when comprehensive documentation was unavailable. This challenge is well-studied: from early work on organizational memory systems [24] 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. As a result, much valuable knowledge remains trapped in individual experience rather than being shared with future lab members.

To address this limitation, 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—an aspect that was not captured in our field study.

7.3 AI Agents as Mediator in Organizations

While recent research has demonstrated significant declines in human participation following AI introduction in knowledge communities [17], 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.

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 [44, 69]. 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 [18, 19, 32]. 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 [87]. 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 of the findings. First, our field deployment involved a relatively small sample of participants (21 total) 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, conducting the study during the summer when lab activities were less active than usual may not reflect typical year-round organizational memory practices.

References

[1] 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

[2] 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

[3] 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

[4] 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

[5] 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

[6] 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

[7] 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

[8] 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

[9] 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

[10] 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

[11] 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

[12] 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

[13] 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

[14] Dane Bertram, Amy Voida, 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

[15] 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

[16] 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

[17] 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

[18] 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

[19] 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

[20] 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

[21] 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

[22] 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

[23] 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

[24] Jeff Conklin. 1997. Designing Organizational Memory: Preserving Intellectual Assets in a Knowledge Economy. (1997).

[25] John W Creswell and Vicki L Plano Clark. 2017. *Designing and conducting mixed methods research*. Sage publications.

[26] 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

[27] 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

[28] 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

[29] 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

[30] 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

[31] 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

[32] 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

[33] 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

[34] Andrea Forte and Amy Bruckman. 2005. Why Do People Write for Wikipedia? Incentives to Contribute to Open-Content Publishing. (2005).

[35] 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

[36] 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]

[37] 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

[38] Jonathan Grudin. 1994. Groupware and Social Dynamics: Eight Challenges for Developers. *Commun. ACM* 37, 1 (Jan. 1994), 92–105. doi:10.1145/175222.175230

[39] 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

[40] 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

[41] 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

[42] 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

[43] B. Kerr. 2003. Thread Arcs: An Email Thread Visualization. In *IEEE Symposium on Information Visualization 2003.* 211–218. doi:10.1109/INFVIS.2003.1249028

[44] 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

[45] 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

[46] 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]

[47] 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

[48] 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

[49] Jean Lave and Etienne Wenger. 1991. *Situated Learning: Legitimate Peripheral Participation.* Cambridge University Press.

[50] 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

[51] 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

[52] 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

[53] 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]

[54] Lena Mamykina, Bella Manoim, Manas Mittal, George Hripcsak, 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

[55] 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

[56] Bonnie A. Nardi, Steve Whittaker, and Erin Bradner. 2000. Interaction and Outeraction: 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

[57] 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

[58] Ikujiro Nonaka. 1998. The Knowledge-Creating Company. In *The Economic Impact of Knowledge*. Routledge.

[59] 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

[60] 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

[61] Jeffrey M. Perkel. 2017. How Scientists Use Slack. *Nature* 541, 7635 (Jan. 2017), 123–124. doi:10.1038/541123a

[62] 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

[63] 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 (2009), 1.

[64] Michael Polanyi. 1997. The Tacit Dimension. In *Knowledge in Organisations*. Routledge.

[65] Yvonne Rogers. 2011. Interaction Design Gone Wild: Striving for Wild Theory. *Interactions* 18, 4 (July 2011), 58–62. doi:10.1145/1978822.1978834

[66] 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

[67] 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

[68] 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

[69] 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

[70] Slack. 2020. Distance Learning Thrives in Slack. https://slack.com/blog/collaboration/distance-learning-in-slack.

[71] 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:10.1177/030631289019003001

[72] 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

[73] John Suler. 2004. The Online Disinhibition Effect. *CyberPsychology & Behavior* 7, 3 (June 2004), 321–326. doi:10.1089/1094931041291295

[74] 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

[75] 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

[76] 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

[77] 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

[78] 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

[79] James P. Walsh and Gerardo Rivera Ungson. 1991. Organizational Memory. *The Academy of Management Review* 16, 1 (1991), 57–91. doi:10.2307/258607 jstor:258607

[80] 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

[81] 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

[82] 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

[83] 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

[84] 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

[85] 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

[86] 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]

[87] 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]

[88] 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]