

Redistributing Interdependence in Organizational Knowledge Work: Lessons from Deploying an LLM Mediator in Research Labs

SANGWOOK LEE, Virginia Tech, USA

SANG WON LEE, Virginia Tech, USA

Large Language Models (LLMs) are increasingly embedded in collaborative workflows, yet their structural effects on the relationships between human stakeholders in data-intensive knowledge work remain underexplored. In this position paper, we reinterpret findings from a month-long field deployment of an LLM-based chatbot that mediates organizational memory across four university research labs (N=21) through the lens of Interdependence Theory. Our analysis reveals two key dynamics. First, the LLM redistributes the *dependence structure* between students and lab directors: while students gain autonomous access to institutional knowledge, directors lose visibility into knowledge gaps, shifting bilateral dependence toward unilateral dependence. This redistribution is moderated by organizational culture, amplifying *mutual responsiveness* in psychologically safe environments while dampening it where students default to private interaction. Second, the system fails to support the *transformation of motivation* needed for collaborative knowledge stewardship: students consistently avoid contributing to documentation due to role perceptions and temporal asymmetry between individual costs and collective benefits. We derive design implications including privacy-preserving awareness mechanisms, graduated contribution pathways, and designing for the LLM’s dual role as a boundary object across stakeholder groups.

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

Additional Key Words and Phrases: human-LLM collaboration, interdependence theory, organizational memory, chatbot, knowledge management

1 Introduction

Modern data and knowledge workflows are rarely solo efforts. They bring together people with different expertise, vocabularies, and degrees of trust in automation [? ?]. University research labs exemplify this dynamic: lab directors maintain institutional knowledge, while students must locate, interpret, and contribute to that knowledge as part of their academic work [?]. These workflows are inherently collaborative and data-intensive, involving the ongoing collection, curation, and retrieval of organizational knowledge distributed across documents, conversations, and human memory [? ?]. At the same time, Large Language Models (LLMs) are increasingly embedded in tools for data processing, knowledge retrieval, and documentation [? ?]. While prior work has shown that AI intermediaries can reshape group dynamics [? ? ?], how these reconfigurations alter the structural dependencies among stakeholders in collaborative knowledge work remains less understood.

In this position paper, we draw on findings from a month-long field deployment of an LLM-based chatbot designed to mediate organizational memory in university research labs [?]. Rather than treating the deployment solely as a systems evaluation, we revisit its findings through the lens of *Interdependence Theory* [? ?], a framework from social psychology that analyzes how actors’ outcomes are mutually contingent on one another’s actions. We argue that Interdependence Theory offers a productive lens for understanding human-LLM collaboration in knowledge work because it foregrounds questions that purely technical evaluations often overlook: *How does the introduction of an LLM change who depends on whom? What happens to mutual responsiveness when an AI mediates communication? How do shared outcomes shift when individual actors can bypass collective channels?* We present an interdependence-theoretic analysis of these findings and derive design implications for LLM-mediated collaborative systems.

Authors’ Contact Information: Sangwook Lee, Virginia Tech, Blacksburg, Virginia, USA, sangwooklee@vt.edu; Sang Won Lee, Virginia Tech, Blacksburg, Virginia, USA, sangwonlee@vt.edu.

Manuscript submitted to ACM

1

2 Background

2.1 The Deployed System and Key Findings

The system analyzed in this paper is an LLM-based chatbot embedded in Slack that supports organizational memory in university research labs [?]. It provides document-grounded Q&A via retrieval-augmented generation [?], peer knowledge sharing with optional anonymity, conversational knowledge extraction, and director-approved document updates. All documentation changes require explicit director approval, ensuring that the director retains final authority over the organizational memory.

The system was deployed for one month across four university research labs (N=21). Participants submitted 107 questions, 42% of which were answered using existing documentation. Lab directors made 38 documentation commits, adding 1,994 words to the organizational memory. The key findings that motivate our analysis include: (a) a *privacy-awareness tension*, where the majority of student questions were asked through private direct messages, limiting directors’ visibility into knowledge gaps; (b) *culture-dependent usage*, where one lab’s culture of psychological safety led to predominantly public questioning and significantly more documentation updates; and (c) *persistent barriers to student contribution*, where students hesitated to suggest documentation updates despite the system’s collaborative features.

2.2 Interdependence Theory

Interdependence Theory, originating from the work of Kelley and Thibaut [?], provides a framework for analyzing situations in which two or more actors’ outcomes are mutually contingent on one another’s actions [?]. The theory characterizes the *dependence structure* between actors, meaning the degree and direction of influence each party’s actions have on the other’s outcomes. It further examines *mutual responsiveness*, the extent to which actors attend and adapt to each other’s needs, and *transformation of motivation*, the process by which actors shift from self-interested behavior (“given preferences”) toward behavior that accounts for the partner’s or group’s welfare (“effective preferences”). The theory also considers *correspondence of outcomes*, whether actors’ interests are aligned or in conflict. While originally formulated for human–human dyads, the theory’s emphasis on structural features of interaction situations, rather than on cognitive capacities, makes it applicable to analyzing human–LLM configurations.

3 An Interdependence-Theoretic Analysis

3.1 Redistribution of Dependence and Disrupted Mutual Responsiveness

Prior to the introduction of the LLM chatbot, the *dependence structure* in research labs was relatively straightforward. Students depended on lab directors for access to institutional knowledge (policies, procedures, implicit norms), and lab directors depended on students’ questions to understand their knowledge needs and to maintain awareness of documentation gaps. This created a bilateral dependence structure: each party’s outcomes (effective knowledge access for students; well-maintained documentation for directors) were contingent on the other’s actions.

The introduction of the LLM mediator fundamentally restructured this dependence. Students’ dependence on directors for routine knowledge retrieval was partially displaced onto the system, as 42% of their questions could be answered directly from existing documentation. This displacement was welcomed by students, who described the chatbot as “a senior in the lab” or “a librarian,” a dependable knowledge source that reduced the social cost of asking questions. However, this redistribution simultaneously attenuated directors’ dependence on student questions as a source of awareness. When students obtained answers privately, the informational signal that directors relied upon (observing what students asked and where they struggled) was no longer available. This asymmetric redistribution created a

new interdependence structure in which students' outcomes (knowledge access) improved while directors' outcomes (documentation awareness) were inadvertently diminished. In Interdependence Theory terms, this represents a shift from bilateral to more unilateral dependence: the LLM, in absorbing one pathway of dependence, inadvertently severed a feedback loop that sustained the collaborative maintenance of organizational knowledge.

Critically, this redistribution manifested differently depending on organizational culture. In one lab (Lab A), where a culture of psychological safety led members to ask questions publicly, the LLM amplified *mutual responsiveness*: directors could observe student needs, discussions fed into documentation updates, and Lab A's director made the most commits (19). In the remaining three labs, where students predominantly used private direct messages, the LLM dampened mutual responsiveness; one director expressed surprise upon learning how actively students had been using the system privately. The same LLM system can either amplify or dampen mutual responsiveness depending on the organizational culture in which it is embedded. From an interdependence-theoretic perspective, organizational culture functions as a moderator of the situation structure [?], producing different effective situations depending on the interpersonal norms, trust levels, and power dynamics of the group.

3.2 Failures of Transformation in Knowledge Contribution

The findings reveal systematic failures of what Interdependence Theory calls *transformation of motivation*, the process by which actors shift from self-interested behavior ("given preferences") to prosocial behavior that considers the group's welfare ("effective preferences"). Students recognized that contributing their knowledge to organizational documentation would benefit future lab members, yet they consistently prioritized their given preferences: avoiding the risk of documenting inaccurate information, protecting themselves from social evaluation, and minimizing the effort of generalizing personal experiences. As one student noted, the concern was that contributed knowledge "would stay in the database forever" and could "mislead all the students."

Notably, this reluctance persisted even though the system's design included a safeguard: all documentation updates required explicit approval from the lab director before being applied. Students were aware of this review process, yet the knowledge that a director would see and evaluate their contribution did not alleviate their hesitation; if anything, it added a layer of social evaluation. More fundamentally, students perceived documentation maintenance as the director's prerogative and responsibility rather than a shared endeavor. This framing positioned students as passive consumers of organizational knowledge and directors as sole curators, reinforcing an asymmetric dependence structure in which students depended on directors for both knowledge access and knowledge maintenance. Even when the system provided technical pathways for student contribution, the perceived role boundary limited the transformation of motivation necessary for collaborative knowledge stewardship.

This failure of transformation also reflected a structural property of the interdependence situation. Contributing to documentation creates a temporal asymmetry: the costs (effort, risk of error, social exposure) are borne immediately by the contributor, while the benefits accrue to future, often unknown, lab members. In Interdependence Theory terms, this is a *low-correspondence* situation where the contributor's immediate interests diverge from the collective's long-term interests. The chatbot's Knowledge Extraction feature partially addressed this by lowering the cost of contribution, transforming ongoing conversations into documentation without requiring separate writing effort. However, the psychological barriers persisted, suggesting that LLM-mediated systems must go beyond reducing technical friction and actively address the role perceptions and social dynamics that inhibit prosocial transformation.

4 Design Implications

Our interdependence-theoretic analysis yields several design implications for LLM-mediated collaborative systems, particularly those operating in knowledge-intensive, data-rich organizational contexts.

Privacy-preserving interdependence mechanisms. The privacy-awareness tension represents a fundamental design challenge [?]: how to maintain the informational interdependence that directors need for effective documentation maintenance without compromising students' autonomy and psychological safety. We recommend *aggregated awareness mechanisms*, such as periodic, anonymized summaries of question patterns, knowledge gaps, and documentation usage, that preserve the interdependence signal without exposing individual behavior. This approach maintains the structural interdependence between stakeholders while respecting individual preferences for privacy, aligning with what participants in the study explicitly suggested.

Graduated transformation support. To address the observed failures of transformation [?], LLM-mediated systems should provide graduated pathways for knowledge contribution [?]. Rather than requiring contributors to immediately produce universally applicable documentation, systems should support a spectrum, from low-commitment annotations ("this worked for me") to collaborative refinement (peer discussion and director review) to formal documentation. Automatically extracting experiential knowledge from ongoing conversations and drafting documentation from them can lower the barrier from "authoring documentation" to "sharing what happened." Crucially, the design must also address role perceptions: when students view documentation as exclusively the director's domain, even well-designed contribution mechanisms may go unused. Systems could reframe contribution as a normal part of membership in the organization rather than an encroachment on managerial authority.

Designing for the LLM's dual role as boundary object. The deployment revealed that users simultaneously experienced the LLM as a social agent—students described it as "a senior in the lab"—and as a communication medium through which directors broadcast documentation updates. This duality suggests that LLM mediators in organizational contexts function as boundary objects [?]: artifacts that maintain shared identity across different social worlds while serving each group's situated needs [? ? ?]. Designers should account for this dual role, ensuring that the same system can support private, low-stakes knowledge retrieval for individual users while also enabling collective knowledge maintenance and organizational awareness for managers.

Temporary page!

LaTeX was unable to guess the total number of pages correctly. As there was some unprocessed data that should have been added to the final page this extra page has been added to receive it.

If you rerun the document (without altering it) this surplus page will go away, because LaTeX now knows how many pages to expect for this document.