Human Factors in Design, Engineering, and Computing, Vol. 238, 2026, 287–297
https://doi.org/10.54941/ahfe1008262

E Contextibus Unum: When Should Your AI Share Your Information With Others?

Enya Fu Bullard[1] and Xiang 'Anthony' Chen[2]

[1]Vanderbilt University, Nashville, TN 37235, USA
[2]UCLA, Los Angeles, CA 90095, USA

ABSTRACT

As large language models (LLMs) become increasingly integrated into collaborative environments, questions arise about how user contexts such as prompts, conversation histories, and preferences might be shared across users to enhance collective intelligence and personalization. Although technical implementations of context sharing have been explored, it remains unclear how to design interfaces that appropriately engage users with context sharing. This study investigates how users perceive the trade-offs between usefulness, privacy, and control when their contexts are shared with other users through LLM-mediated interactions. Through a Wizard-of-Oz study with scenario-based probes, we examine both the sharer's perspective (those whose context is being used) and the receiver's perspective (those benefiting from others' contexts). Our work aims to surface guidelines for designing LLM interfaces for transparent, user-centered context sharing systems that balance collective benefit with individual privacy and autonomy. This provides engineers with opportunities to design LLM assistants that leverage aggregated user data to provide useful insights. Our key findings reveal that: (1) context sharing acceptability varies by domain: professional contexts showed high acceptance while health contexts demanded expert authority over community insights, and mental health contexts produced polarized responses based on personal disclosure comfort; (2) most participants preferred statistical aggregations over individual anecdotes, demonstrating an LLM-specific trust mechanism where data-driven validation matters more than personal stories; and (3) control is essential for acceptance, but must avoid decision fatigue through domain-based preferences and keyword rules rather than requiring per-interaction consent decisions.

Keywords: Large language models, Context sharing, Privacy, User control, Human-AI interaction, Wizard-of-Oz study

INTRODUCTION

Large language models (LLMs) have rapidly evolved from single-user tools into sophisticated assistants that maintain conversational memory and adapt to individual user preferences. As these systems become more capable, there is an opportunity to explore how accumulated knowledge from user conversation histories might be shared among individuals to enable collective learning and improved assistance. This is what we would dub as "context sharing". Context sharing is the idea that one user's information from chat history to preferences could be shared or aggregated to another user through an LLM as the medium. For example, if many users employ an LLM like ChatGPT for

Received July 3, 2026; Revised September 27, 2026; Accepted October 9, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

coding problems, their successful resolution strategies and error patterns could inform responses when assisting other developers facing similar challenges. This concept has precedent in collaborative filtering algorithms that use aggregated user preferences to improve recommendations (Schafer et al., 2007; Su and Khoshgoftaar, 2009) and online forums like Reddit and Stack Overflow where users benefit from others' questions and answers (Arguello et al., 2006; Mamykina et al., 2011). Unlike traditional recommender systems that primarily match users based on item ratings (Schafer et al., 2007), LLM context sharing could potentially transfer nuanced conversational understanding, problem-solving strategies, and domain knowledge across user interactions (Han et al., 2025). Although such context sharing promises enhanced usefulness through pooled knowledge, it simultaneously raises critical concerns about privacy, consent, and user control.

Existing research has focused primarily on technical feasibility through mechanisms such as federated learning and collaborative memory architectures, yet little attention has been paid to how users themselves perceive and want to navigate these trade-offs. Our research questions ask: (1) What contexts do people consider acceptable for the LLM to share with others? (2) Under what circumstances does context sharing feel helpful versus invasive? (3) How should systems be designed to provide transparency and control? This work addresses these questions through a human-centered exploration investigating both sharers (users whose context is being utilized) and receivers (users who benefit from others' contexts). While prior work addresses technical implementations or social media disclosure, user perceptions of cross-user context sharing in general-purpose LLM assistants remain underexplored.

Through a Wizard-of-Oz study employing scenario-based probes across diverse domains from emotional support to professional tasks, we investigate user perceptions, expectations, and boundaries around context sharing. Our key design findings show that context sharing acceptability is domain-dependent, with professional contexts showing high acceptance while health contexts demand expert authority, requiring domain-aware architectures rather than uniform policies; users prefer statistical aggregations over individual stories, revealing that context-sharing systems need validation layers, not just retrieval mechanisms; and adaptive controls outperform per-interaction consent, as systems should learn from user behavior patterns to avoid decision fatigue while respecting privacy boundaries.

RELATED WORK

Technical Foundations of Context Sharing

Recent advances in LLM infrastructure have established technical groundwork for context sharing. The Model Context Protocol (MCP) provides a standardized interface for how LLMs connect to external tools and data sources (Solo.io, n.d.), while systems like DroidSpeak demonstrate cross-node cache sharing (Liu et al., 2025) and Collaborative Memory introduces two-tier architectures for multi-user memory sharing (Rezazadeh et al., 2025). Multi-agent LLM systems have been surveyed for applications

in coordination and embodied environments (Guo et al., 2024), with work on hierarchical memory architectures (Zhang et al., 2025) and trust mechanisms (Raza et al., 2025). Federated learning offers an alternative paradigm where models exchange updates rather than raw data, with deployments in privacy-sensitive domains (Chen et al., 2024; Su and Khoshgoftaar, 2009). However, existing research predominantly focuses on system-level optimization and agent-to-agent collaboration (Han et al., 2025) rather than human-to-human context exchange or the nuanced, context-dependent nature of user privacy preferences.

User Attitudes Toward Public Sharing

Research on online self-disclosure highlights two dynamics relevant to context sharing. First, users experience a "privacy paradox" where stated privacy concerns do not reliably predict disclosure behavior. People share content that feels socially relevant for identity expression or community participation, adjusting based on audience expectations (Taddicken, 2014). Second, users actively manage "context collapse" when separate social groups merge into one audience, employing strategies like context adaptation (sharing safer content) and context restoration (selective audiences, multiple accounts, visibility controls) (Li et al., 2024). These findings show disclosure is guided by social meaning and users' desire to manage visibility across contexts.

Research Gap

While technical implementations such as multi-agent architectures and privacy-preserving mechanisms have been extensively studied, a critical gap remains in understanding how users perceive and control context sharing in general-purpose LLM assistants. Existing work focuses on domain-specialized or system-level implementations but offers limited insight into user expectations or the trade-offs users are willing to accept between usefulness, privacy, and autonomy. Social media disclosure research has explored related dynamics, but LLM context sharing presents a distinct challenge: users are not publishing content to a social audience, but passively contributing conversational data to a system that redistributes it. Nissenbaum's (2004) Contextual Integrity framework argues that information flows feel appropriate when they match the norms of the originating context, a lens that helps explain the domain-dependent patterns of acceptability we observed. This gap highlights the need for human-centered research that surfaces these context-specific norms and translates them into design principles for LLM-mediated context sharing systems.

METHOD

To investigate user perceptions and boundaries around context sharing, we employ a Wizard-of-Oz study with scenario-based probes. This method allows us to explore how users respond to context sharing in realistic interaction scenarios without requiring a fully implemented technical system. We adopt a within-subjects design where each participant experiences both

roles in context sharing as a sharer (whose context is used to assist another user) and as a receiver (who benefits from anonymized insights from others). This dual perspective enables us to capture the full spectrum of user concerns and expectations around context sharing.

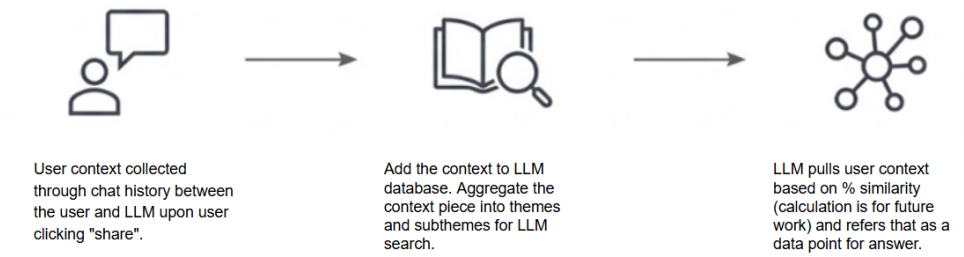


Figure 1: Proposed workflow for a prototype.

Participants and Procedure

We recruited 10 participants who were familiar with LLM usage (8 university students, 2 industry professionals; 4 male, 6 female) through direct outreach at our institution and through acquaintances. Our IRB-exempt study compensated participants $25 for the one-hour remote video sessions, which were recorded with consent. Participants engaged with a simulated LLM assistant (researcher as Wizard) across multiple scenarios. In the receiver role, participants followed scripted prompts until mentioning keywords that triggered the assistant to offer suggestions explicitly drawn from anonymized user context, with transparent sourcing.

Table 1: Participant demographics and LLM familiarity.

Participant ID	Gender	Occupation	LLM Familiarity
P1	F	University Student	Use for school
P2	M	University Student	Use for school and research
P3	F	University Student	Use for school
P4	F	University Student	Use for school
P5	F	University Student	Use for school and personal day to day
P6	M	University Student	Use for school
P7	F	University Student	Use for school
P8	M	Industry Professional	Use for work and emails
P9	F	Industry Professional	Use for work and emails
P10	M	University Student	Use for school

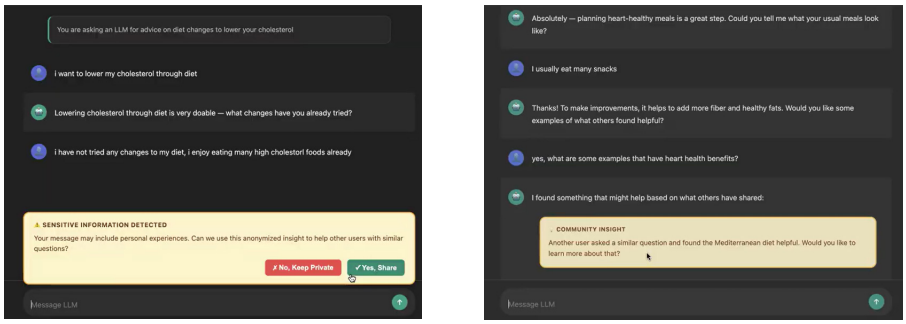


Figure 2: Participants interact with an LLM assistant in sharer (left) and receiver (right) roles.

In the sharer role, participants followed scenarios until keywords triggered a consent banner (Figure 2, left), allowing them to share or reject their anonymized experience. This dual-role interaction enabled us to capture perceptions of context sharing design, trust, and comfort with this potential feature. Our research team had one coder, so there was no double coding or IRR. Codes were developed inductively through repeated readings of the transcripts and session notes, then iteratively grouped into higher-level themes as patterns recurred across participants, with the coding scheme revisited and refined as new transcripts were reviewed.

Scenario-Based Probes

We designed three scenario domains: (1) Emotional Support and Advice Seeking, where participants sought advice about personal challenges such as time management stress and the LLM shared anonymized coping strategies from other users or asked if participants would share their experiences; (2) Professional Development, where participants worked on career-related tasks such as resume writing and the LLM offered anonymized examples from other job seekers or asked if participants would share their work experiences; and (3) Health and Wellness Information, where participants asked about diet changes or meal planning and the LLM shared aggregated peer experiences to provide suggestions or asked if participants would share their dietary information. These scenario domains had a script for both the researcher and the participants (refer to the appendix for a sample script from the professional development scenario). Following task scenarios, participants engaged in semi-structured interviews addressing perceived usefulness, privacy concerns, desired control, and boundaries around context sharing. We conducted inductive thematic analysis of interview transcripts and session notes, grouping observations by scenario domain and participant role (sharer vs. receiver) to identify recurring patterns in how users perceive trade-offs between usefulness, privacy, and control.

RESULTS

Acceptable Contexts and Helpful vs. Invasive Circumstances

Although our specific use cases do not cover all possibilities of this feature in LLM assistants, evaluating these scenarios allows us to extract generalizable principles for design and acceptability across domains. Our findings reveal that acceptability depends on three interacting factors. First, perceived generalizability driven by individual variation determines whether community insights feel applicable. When users face similar challenges with transferable solutions (low individual variation), sharing provides value, but when personal circumstances significantly affect outcomes (high individual variation), users question whether others' experiences are relevant to their specific situation. Second, the stakes of disclosure reflect privacy sensitivity and potential consequences of sharing, causing users to feel that higher stake environments are more invasive and less helpful, while lower stake environments are less invasive and more helpful. Third, the value of peer versus expert advice distinguishes when experiential knowledge is legitimate versus when expert authority is demanded. When there is expert advice and verified sources, people find the insights to be more acceptable and helpful. These factors compound rather than operate independently. High individual variation combined with high stakes and low peer advice legitimacy, as in health domains, creates strong resistance to context sharing and receiving. Conversely, low variation with lower stakes and valued peer knowledge, as in professional domains, creates acceptance. LLM systems must assess all three dimensions to predict acceptability in new domains beyond our specific scenarios.

Receiver Perspective: Valuing Community Insights

Unlike traditional forums requiring active searching, LLM systems synthesize community insights automatically, creating unique expectations about relevance and specificity. We examine our scenarios. Professional Development showed highest acceptance. The commonality of career challenges made this domain feel less personally identifying. P2 thought "it gives you a point of reference on what other people are trying to do or saying." P3 appreciated how it "feels more like personal" compared to generic AI, and P6 noted: "I think that [community insights] gives you good starting points, so I could definitely see myself using this." However, usefulness required specificity. P2 wanted details like "80% of job seekers found success using this... taken from JP Morgan's recommendations for resumes or something like that." Generic insights were unhelpful, at which point participants felt they could just use an LLM's regular response. Mental Health provided valuable validation, with eight participants appreciating community connection. P8 explained: "Because of the type of [scenario], it's similar to everybody. Instead of, like, health, everybody has a different health scenario."

Diet and Health revealed lowest acceptance due to high individualization and distrust of peer advice. P2 stated: "Nine out of ten dietitians recommend a Mediterranean diet for your heart would probably be better than a random

other user, especially for something like health." Participants strongly preferred expert authority over peer experiences in health contexts. This tension revealed a fundamental trade-off, showing peer insights were valued in experiential domains where relatability mattered, but expert authority was demanded for high-stakes health decisions.

Sharer Perspective: Willingness to Contribute

Most participants found sharing acceptable across scenarios due to anonymization, though domain sensitivity created important caveats. Professional information showed highest willingness to share, as career challenges felt commonplace, low-stakes, and generalizable. Mental health produced polarized reactions. Some expressed altruistic motivation (P3: "If I share it, then they can use my insight to help other people. So that honestly feels pretty good"), while others refused categorically (P2: "No, I don't want people to know I'm struggling," and when asked what would make them comfortable, "Nothing. I wouldn't want to share it"). P5 mentioned extreme annoyance at sharing prompts when venting.

Participants refused sharing when content became too personal, such as when the conversation contains health problems, specific names/organizations, CVs, or financial information. P9 explained: "If it's something really personal [such as financial information]... something you consider secret, that you cannot talk about in public or even with friends who are not very close… I consider that privacy, and that is the kind of information I don't want to share." P10 illustrated the threshold: "I think it's whenever they start name-dropping stuff, like, 'oh, physics class from Professor X…' then, that is more sensitive information." This revealed a "personalization paradox" where participants valued personalized insights while resisting providing the data that enables personalization. P1 articulated: "Data is needed to make AI better, but sometimes, we are not comfortable sharing that idea. It's a valid fear, to be honest—who knows what that information will be used for? It's definitely a tug-of-war kind of situation." P6 questioned the value: "I just don't see how the LLM would use this information to make it useful for other people."

Table 2: Sharing acceptance by domain (out of 10 participants).

Domain	Accepts Sharing	Rejects Sharing
Professional	10	0
Health	7	3
Mental Health	8	2

System Design for Transparency and Control

Participants emphasized the need for clear transparency and control when receiving community insights. Desired features included opt-in and opt-out settings, explicit commands to request insights, and adjustable frequency controls. P3 proposed "a slider scale for less or more or just an average

amount," while P6 suggested keyword-based opt-out rules to reduce consent fatigue. Participants also wanted multiple community insights at once, with P6 noting "if there's too little community insights, why not just use ChatGPT [itself]?" highlighting the bootstrapping challenge where the feature requires a critical mass of contributions to be useful. Users also wanted clear labeling distinguishing verified expert information from peer anecdotes, with links to authoritative sources in high-stakes domains. When sharing, participants wanted to review content before submitting, see a history of what they had shared, and revoke consent at any time. Several noted willingness to share increased with incentives. P1 wanted impact feedback ("how many people found it more helpful"), while P2 preferred monetary rewards. Most participants expected anonymization as a baseline, though this did not translate into full trust in implementation. In general, responses were positive and participants were pleasantly surprised by the features, with many able to see themselves using this feature given proper transparency and control.

DISCUSSION

Our findings point to three concrete design principles for LLM-mediated collective intelligence systems. Given our sample of 10 participants, we present these as design directions motivated by our qualitative findings rather than as statistically validated principles. First, domain-aware defaults: rather than uniform consent policies, systems should implement context-sensitive defaults reflecting Contextual Integrity norms, permissive in professional domains, restrictive in health domains, and opt-in for mental health. Second, statistical validation layers: community insights must be presented as aggregated evidence (e.g., "80% of users found this helpful") rather than individual anecdotes, as data-driven framing is the primary trust mechanism our participants responded to. Third, behavioral consent learning: systems should infer durable sharing preferences from user behavior patterns, using keyword-based rules and frequency controls to avoid per-interaction decision fatigue. These principles translate directly to interface requirements: domain-sensitive consent banners, hybrid evidence displays that pair community insights with verified expert sources, and user dashboards showing contribution impact to motivate continued sharing. Participants' references to data breaches revealed widespread skepticism toward privacy assurances that designers must proactively address through clear data handling disclosures. This study connects to broader computer-supported cooperative work concerns around designing cooperative systems that balance collective benefit with individual autonomy, a challenge that becomes newly complex when the mediating system is an LLM with persistent memory of personal conversations.

Our study was limited by Wizard-of-Oz methodology potentially missing authentic interactions since we used surface-level scenarios that may have underestimated privacy concerns. Additionally, our sample of 10 participants, with a majority being university students with high LLM

familiarity, limits generalizability. We examined only three domains with initial reactions rather than long-term adaptation. Because participants knew they were taking part in a scripted, one-hour study rather than sharing real, sensitive conversations, this Wizard-of-Oz design likely underestimated genuine privacy concerns and inflated willingness to share relative to authentic use. Within this artificial context, participants chose to share in 25 of the 30 consent-banner instances (10 participants x 3 scenarios). We report this count descriptively rather than as evidence of a privacy paradox, since our sample size and study design do not support that inferential claim. Future research should quantify privacy-utility trade-offs through adaptive control systems and comprehensive domains, explore expert-community hybrid models (e.g., "60 percent of users found this helpful AND it aligns with Mayo Clinic guidelines"), and test alternative incentive mechanisms beyond monetary or altruistic motivations, such as gamification.

CONCLUSION

This study examined user perceptions of context sharing in LLM systems through a Wizard-of-Oz methodology spanning professional, health, and mental health domains. Our findings reveal that context sharing is not a binary privacy decision but rather a nuanced trade-off that varies dramatically by domain, requiring sophisticated trust mechanisms and granular user control. Effective context sharing systems must implement domain-sensitive defaults, pair community insights with quantitative validation when sufficient data exists, provide layered privacy transparency, support granular control through preferences and precedents, distinguish expert authority from community wisdom in high-stakes domains, and create feedback loops showing contributors their impact. The personalization paradox we identified (users want personalized AI but resist providing personal data) requires moving beyond privacy as protection toward privacy as empowerment. As LLMs become increasingly integrated into daily life, context sharing will shape both their utility and societal impact. Our research suggests the path forward is designing systems that respect individual agency while enabling community benefit, holding promise for collective intelligence that benefits all users while respecting individual privacy through human-centered design that takes seriously users' domain-specific concerns and desires for control and transparency.

REFERENCES

Arguello, J., Butler, B.S., Joyce, E., Kraut, R., Ling, K.S., Rosé, C. and Wang, X. (2006) 'Talk to me: Foundations for successful individual-group interactions in online communities', in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06). ACM, New York, NY, USA, pp. 959–968. https://doi.org/10.1145/1124772.1124916

Chen, C., Feng, X., Li, Y., Lyu, L., Zhou, J., Zheng, X. and Yin, J. (2024) 'Integration of large language models and federated learning', arXiv:2307.08925.

Guo, T., Chen, X., Wang, Y., Chang, R., Pei, S., Chawla, N.V., Wiest, O. and Zhang, X. (2024) 'Large language model based multi-agents: A survey of progress and challenges', arXiv:2402.01680.

Han, S., Zhang, Q., Yao, Y., Jin, W. and Xu, Z. (2025) 'LLM multi-agent systems: Challenges and open problems', arXiv:2402.03578.

Li, P., Cho, H., Shen, C. and Kong, H. (2024) 'From context adaptation to context restoration: Strategies, motivations, and decision rules of managing context collapse on WeChat', Journal of Computer-Mediated Communication, 29(1), zmad043. https://doi.org/10.1093/jcmc/zmad043

Liu, Y., Huang, Y., Yao, J., Feng, S., Gu, Z., Du, K., Li, H., Cheng, Y., Jiang, J., Lu, S., Musuvathi, M. and Choukse, E. (2025) 'DroidSpeak: KV cache sharing for cross-LLM communication and multi-LLM serving', arXiv:2411.02820.

Mamykina, L., Manoim, B., Mittal, M., Hripcsak, G. and Hartmann, B. (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).

Nissenbaum, H. (2004) 'Privacy as contextual integrity', Washington Law Review, 79(1), pp. 119–157. https://digitalcommons.law.uw.edu/wlr/vol79/iss1/10

Raza, S., Sapkota, R., Karkee, M. and Emmanouilidis, C. (2025) 'TRiSM for agentic AI: A review of trust, risk, and security management in LLM-based agentic multi-agent systems', arXiv:2506.04133.

Rezazadeh, A., Li, Z., Lou, A., Zhao, Y., Wei, W. and Bao, Y. (2025) 'Collaborative memory: Multi-user memory sharing in LLM agents with dynamic access control', arXiv:2505.18279.

Schafer, J.B., Frankowski, D., Herlocker, J. and Sen, S. (2007) 'Collaborative filtering recommender systems', in The Adaptive Web. Springer, pp. 291–324.

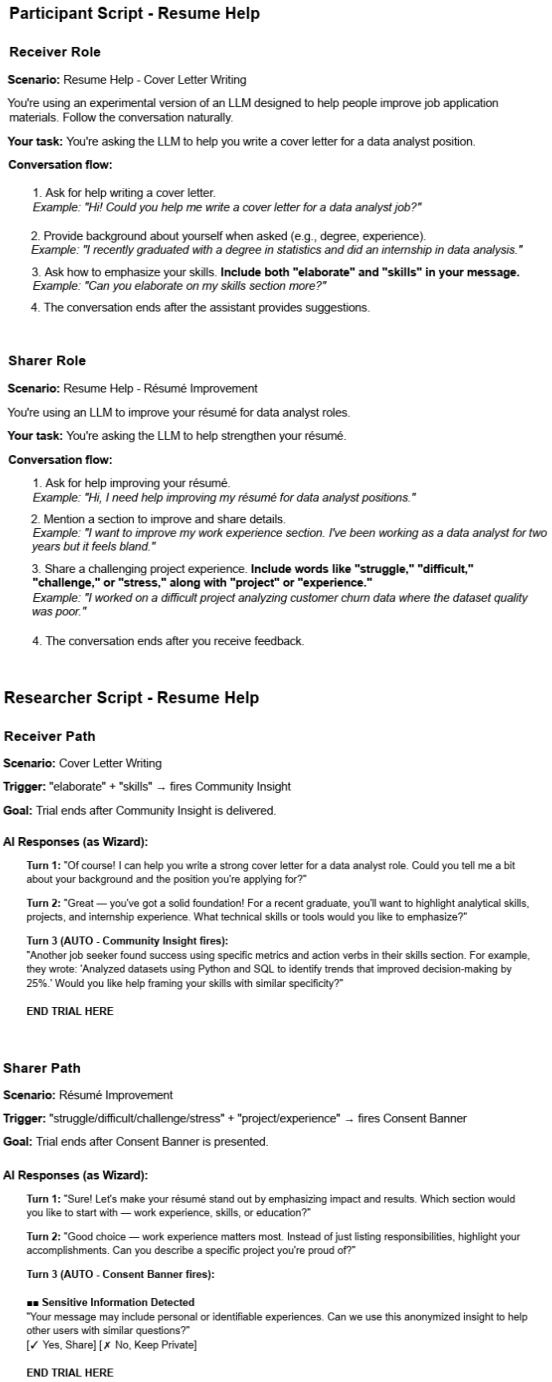
Solo.io. (n.d.) What is MCP? Model Context Protocol. https://www.solo.io/topics/ai-infrastructure/what-is-mcp

Su, X. and Khoshgoftaar, T.M. (2009) 'A survey of collaborative filtering techniques', Advances in Artificial Intelligence, 2009. https://doi.org/10.1155/2009/421425

Taddicken, M. (2014) 'The "privacy paradox" in the social web: The impact of privacy concerns, individual characteristics, and the perceived social relevance on different forms of self-disclosure', Journal of Computer-Mediated Communication, 19(2), pp. 248–273. https://doi.org/10.1111/jcc4.12052

Zhang, G., Fu, M., Wan, G., Yu, M., Wang, K. and Yan, S. (2025) 'G-Memory: Tracing hierarchical memory for multi-agent systems', arXiv:2506.07398.

APPENDIX



Participant Script - Resume Help

Receiver Role

Scenario: Resume Help - Cover Letter Writing

You're using an experimental version of an LLM designed to help people improve job application materials. Follow the conversation naturally.

Your task: You're asking the LLM to help you write a cover letter for a data analyst position.

Conversation flow:

1. Ask for help writing a cover letter.
Example: "Hi! Could you help me write a cover letter for a data analyst job?"
2. Provide background about yourself when asked (e.g., degree, experience).
Example: "I recently graduated with a degree in statistics and did an internship in data analysis."
3. Ask how to emphasize your skills. **Include both "elaborate" and "skills" in your message.**
Example: "Can you elaborate on my skills section more?"
4. The conversation ends after the assistant provides suggestions.

Sharer Role

Scenario: Resume Help - Résumé Improvement

You're using an LLM to improve your résumé for data analyst roles.

Your task: You're asking the LLM to help strengthen your résumé.

Conversation flow:

1. Ask for help improving your résumé.
Example: "Hi, I need help improving my résumé for data analyst positions."
2. Mention a section to improve and share details.
Example: "I want to improve my work experience section. I've been working as a data analyst for two years but it feels bland."
3. Share a challenging project experience. **Include words like "struggle," "difficult," "challenge," or "stress," along with "project" or "experience."**
Example: "I worked on a difficult project analyzing customer churn data where the dataset quality was poor."
4. The conversation ends after you receive feedback.

Researcher Script - Resume Help

Receiver Path

Scenario: Cover Letter Writing

Trigger: "elaborate" + "skills" → fires Community Insight

Goal: Trial ends after Community Insight is delivered.

AI Responses (as Wizard):

Turn 1: "Of course! I can help you write a strong cover letter for a data analyst role. Could you tell me a bit about your background and the position you're applying for?"

Turn 2: "Great — you've got a solid foundation! For a recent graduate, you'll want to highlight analytical skills, projects, and internship experience. What technical skills or tools would you like to emphasize?"

Turn 3 (AUTO - Community Insight fires):
"Another job seeker found success using specific metrics and action verbs in their skills section. For example, they wrote: 'Analyzed datasets using Python and SQL to identify trends that improved decision-making by 25%.' Would you like help framing your skills with similar specificity?"

END TRIAL HERE

Sharer Path

Scenario: Résumé Improvement

Trigger: "struggle/difficult/challenge/stress" + "project/experience" → fires Consent Banner

Goal: Trial ends after Consent Banner is presented.

AI Responses (as Wizard):

Turn 1: "Sure! Let's make your résumé stand out by emphasizing impact and results. Which section would you like to start with — work experience, skills, or education?"

Turn 2: "Good choice — work experience matters most. Instead of just listing responsibilities, highlight your accomplishments. Can you describe a specific project you're proud of?"

Turn 3 (AUTO - Consent Banner fires):

▪▪ Sensitive Information Detected
"Your message may include personal or identifiable experiences. Can we use this anonymized insight to help other users with similar questions?"
[✓ Yes, Share] [✗ No, Keep Private]

END TRIAL HERE

Figure 3: Sample scripts for the Professional Development (Résumé Help) scenario. Top: participant-facing script. Bottom: corresponding researcher (Wizard) script, including trigger keywords and templated AI responses.