Productive Frictions: Redesigning Social Computing Research in the Age of AI

iConference 2027 Workshop Proposal

Type: Workshop **Duration:** One 90-Minute Session

Yun Huang
yunhuang@illinois.edu
School of Information Sciences, University of Illinois Urbana-Champaign
Champaign, Illinois, USA

Yulin Yu
yulinyu@arizona.edu
College of Information Science, University of Arizona
Tucson, Arizona, USA

Jorge Goncalves
jorge.goncalves@unimelb.edu.au
School of Computing and Information Systems, University of Melbourne
Melbourne, Victoria, Australia

Chun Yu
chunyu@tsinghua.edu.cn
Department of Computer Science and Technology Tsinghua University
Beijing, China

Caterina Moruzzi
cmoruzzi@ed.ac.uk
Institute for Design Informatics, School of Design, University of Edinburgh
Edinburgh, Scotland, UK

Hangyue Zhang
hz85@illinois.edu
School of Information Sciences, University of Illinois Urbana-Champaign
Champaign, Illinois, USA

Sangho Suh
sanghos@allenai.org
Allen Institute for AI
Seattle, Washington, USA

Jiaxiong Hu
hujx@ust.hk
Hong Kong University of Science and Technology
Hong Kong, China

Yiren Liu
yirenl2@illinois.edu
Informatics, University of Illinois Urbana-Champaign
Champaign, Illinois, USA

Yang Wang
yvw@illinois.edu
School of Information Sciences, University of Illinois Urbana-Champaign
Champaign, Illinois, USA

Abstract

The growing adoption of large language models across scientific workflows—from literature synthesis and data exploration to manuscript drafting and peer review—has intensified debates over intellectual agency, research integrity, and scholarly deskilling. While automated systems offer rapid idea generation and pattern synthesis, researchers express widespread concern over methodological opacity, intellectual homogenization, and the erosion of foundational analytical skills. This workshop convenes scholars from human-computer interaction, information science, and computational social science to examine empirical tensions and design opportunities in AI-supported research. Rather than treating automation as an uncritical goal, we investigate the concept of *productive friction*: interface mechanisms, cognitive pauses, and procedural constraints that provoke critical reflection, preserve epistemic diversity, and sustain human judgment. We structure our agenda around four interconnected themes: knowledge discovery and inquiry; data combinations and methodological innovation; team collaboration and newcomer socialization; and evaluation, peer review, and institutional governance. The 90-minute workshop will combine lightning talks, breakout discussions, and collective synthesis to produce an actionable research agenda for human-centered scholarship in social computing.

Keywords

AI-supported research, social computing, large language models, human-centered AI

1 Description

1.1 Purpose and Intended Audience

Large language models are increasingly integrated into scientific workflows. Models such as ChatGPT, Gemini, and Asta [2] offer specialized capabilities that automate literature synthesis, hypothesis generation, and experimental design. Recent literature describes "Agentic Science" as an emerging framework for automated inquiry [24], with complementary work exploring co-evolving human-AI ecosystems [20]. While initial benchmarks show that models can generate novel research ideas [21], field-wide surveys

document concerns regarding methodological reliability, intellectual homogenization, and the erosion of foundational research training [13]. As automated systems take on complex analytical tasks, researchers face open questions about how human judgment, problem formulation, and research oversight are sustained.

Social computing is inherently interdisciplinary, demanding theoretical grounding, mixed-methods triangulation, and iterative empirical evaluations [22]. Recent HCI studies demonstrate that language models can augment creative research workflows beyond routine task automation [7, 10, 14–16, 18, 30]. In empirical deployments, models proved most useful when providing unfamiliar perspectives that prompt critical reflection [30], when artificial processing delays gave researchers time to deliberate over problem framings [14], or when multi-agent debates highlighted competing assumptions [15]. In each setting, the primary benefit arose from points of productive friction: intentional constraints and pauses that preserve researcher agency and promote deeper reflection [8, 15, 16]. As Ottino and Uzzi argue [17], scientific advancement relies not merely on transactional speed, but on redundancy, divergent exploration, and productive friction that safeguard epistemic diversity.

The workshop is intended for researchers, educators, practitioners, and graduate students across information science, HCI, communication, organizational science, AI, and adjacent fields. It welcomes participants at all career stages who study, design, use, teach, or govern AI-supported research practices.

1.2 Goals and Outcomes

This workshop provides an open forum to examine emerging opportunities and practical tensions in AI-supported social computing research. Building on our CSCW 2026 Special Interest Group on agency, collaboration, and governance [9], we extend this dialogue into a concrete research agenda tailored to the iSchool community. Our focus is on operationalizing productive friction across the research lifecycle by identifying interface mechanisms, methodological protocols, and community norms that encourage critical reflection instead of passive reliance on model outputs. By bringing together perspectives from human-computer interaction, information science, and computational social science, we aim to formulate actionable design principles and address open challenges across four interconnected themes:

- **Knowledge discovery, learning, and inquiry**: Conversational search tools and automated summaries reshape how scholars encounter prior work. By replacing exploratory query reformulation with direct synthesis, these tools risk narrowing information diversity and serendipitous discovery [5, 27]. Similarly, AI adoption on technical forums redistributes epistemic labor between knowledge seekers and contributors [11]. These patterns raise critical questions about methodological judgment and conceptual synthesis, especially for early-career researchers developing evaluative expertise. How can discovery interfaces surface unconventional perspectives, support durable comprehension, and invite critical skepticism rather than funneling users toward consensus answers?
- **Data, methods, and research innovation**: Social computing relies on diverse empirical datasets, yet the proliferation of synthetic data and automated extraction introduces challenges for data provenance, reproducibility, and consent. At the same time, prior work demonstrates that unusual combinations of datasets yield higher scientific impact [28], and cross-field data repurposing can drive scientific disruption [25]. How can AI help scholars discover productive pairings across datasets and traditions while maintaining rigor and ethical boundaries? What shared documentation, data literacies, and credit models are needed to support responsible data reuse across disciplines?
- **Collaboration, teams, and scholarly communities**: Integrating AI assistants into collaborative environments alters group communication, team formation, and community participation. Enterprise data indicate that AI adoption can shift knowledge work toward solitary production at the expense of communicative interaction [26], even while structured recommendations help newcomers establish valuable professional ties [29]. Concurrently, software agents, meeting summarizers, and facilitative bots reshape interpersonal trust, coordination, and team authority [1, 19]. How can these systems support spontaneous dialogue and cross-domain collaboration without replacing human exchange or reinforcing academic hierarchies?
- **Evaluation, authorship, and institutional governance**: Automated tools now participate in proposal drafting, manuscript preparation, and peer review. Although model feedback can help reviewers identify overlooked technical issues [23], automated evaluators frequently exhibit faulty reasoning and struggle with counterfactual analysis [3], even as unacknowledged AI-generated text appears across conference reviews [12]. Furthermore, benchmarking systems that evaluate and rank proposals risk encoding subjective disciplinary norms and flattening methodological diversity [6]. Given fragmented disclosure policies across venues [4], which evaluative responsibilities must remain human, how should authors disclose AI contributions, and what governance models can sustain institutional trust?

Together, these themes address the broader problem of structuring human-AI interaction in empirical inquiry. Rather than treating speed as the primary benchmark of technical progress, the workshop examines how friction, cognitive checkpoints, and deliberate redundancy safeguard scientific integrity. The workshop will produce a shared research agenda and an open-access report detailing research questions, design opportunities, and principles for responsible AI-supported scholarship, alongside a community repository to support continued exchange and future interdisciplinary collaborations.

1.3 Proposed Format

Our team brings complementary expertise in human–AI interaction (Yun Huang, Jorge Goncalves, Chun Yu, Hangyue Zhang, Jiaxiong Hu), computational social science (Yulin Yu), AI-supported creativity (Caterina Moruzz, Sangho Suh, Yiren Liu), and privacy and AI governance (Yang Wang). Our research spans research ideation and critical reflection through processing delays, simulated expert perspectives, and multi-agent deliberation; information diversity

and the influence of AI-generated summaries; data reuse and scientific innovation; workplace collaboration and newcomer networks; and conversational agent design. Building on our prior work on agency, collaboration, and governance in AI-mediated scholarship, we unite these perspectives to examine how productive friction can preserve human judgment, sustain intellectual diversity, and support responsible research practices.

At the conference, we plan a 90-minute session structured around lightning talks, theme-based discussions, and collective synthesis (Table 1).

- **Lightning Talks & Short Panel (35 min):** Invited experts will present 5-minute lightning talks on key tensions in AI-supported research, followed by an interactive panel and audience Q&A to provoke debate.
- **Theme-based Group Discussions (35 min):** In a "birds of a feather" format, participants will split into breakout tables corresponding to the four workshop themes: knowledge discovery and inquiry; data, methods, and research innovation; collaboration and scholarly communities; and evaluation, authorship, and governance. Organizers will facilitate each table to brainstorm challenges and design opportunities.
- **Report-backs & Closing (10 min):** Group representatives will summarize insights to outline a shared research agenda and action items.

Table 1: Schedule for the 90-minute workshop session.

Time	Event
0–5 min	Welcome & introduction
5–40 min	Lightning talks & short panel
40–45 min	Transition & preparation for group work
45–80 min	Theme-based group discussions
80–90 min	Report-backs, summarization, & closing remarks

1.4 Engagement

Before the workshop, a dedicated website and repository will allow participants to review background materials, suggest discussion topics, and begin networking; recruitment will leverage academic mailing lists and special interest groups across ASIS&T, ACM, and communication associations. During the session, audience Q&A and facilitated breakouts will give every attendee structured opportunities to contribute. Discussion prompts will guide groups in comparing disciplinary perspectives and translating practical tensions into concrete research questions. Session report-backs, the summary report, and the open repository will support ongoing collaborative initiatives. The ten organizers span information science, computing, human-centered design, and communication across five institutions in three countries.

1.5 Relevance to the iConference

This workshop directly engages a central iSchool priority: understanding how people, information, and sociotechnical systems interact within institutional contexts. As AI tools alter knowledge curation, information seeking, and data stewardship, information scholars are uniquely positioned to address emerging methodological and ethical challenges. The workshop connects central iConference topics—including scholarly communication, collaborative data practices, and research integrity—through structured cross-disciplinary inquiry.

1.6 Duration

We request one standard 1.5-hour iConference session (90 minutes). The complete schedule appears in Table 1; no extended session is requested. We anticipate 30–50 participants, from graduate students to senior scholars.

Acknowledgments

This material is based upon work supported by the Institute of Education Sciences, U.S. Department of Education, through Grant 2229873 (NSF). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

[1] Joseph A Allen. 2026. From human gatherings to human–agent meetings: Rethinking small group and team collaboration in the age of AI. *Small Group Research* 57, 3 (2026), 518–522.

[2] Allen Institute for AI. 2026. Asta: Scholarly Research Assistant. https://asta.allen.ai/. Accessed: 2026-06-22.

[3] Nils Dycke and Iryna Gurevych. 2026. Automatic reviewers fail to detect faulty reasoning in research papers: A new counterfactual evaluation framework. *Transactions of the Association for Computational Linguistics* 14 (2026), 465–488.

[4] Conner Ganjavi, Michael B Eppler, Asli Pekcan, Brett Biedermann, Andre Abreu, Gary S Collins, Inderbir S Gill, and Giovanni E Cacciamani. 2024. Publishers' and journals' instructions to authors on use of generative artificial intelligence in academic and scientific publishing: bibliometric analysis. *bmj* 384 (2024).

[5] Jarod Govers, Cherie Sew, Eduardo Velloso, Vassilis Kostakos, and Jorge Goncalves. 2026. Narratives and Perspectives: How AI Summaries Steer Users' Opinions and Engagement on Social Media. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems* (Barcelona, Spain) *(CHI '26)*. Association for Computing Machinery, New York, NY, USA, Article 1398, 19 pages. doi:10.1145/3772318.3790945

[6] Sikun Guo, Amir Hassan Shariatmadari, Guangzhi Xiong, Albert Huang, Myles Kim, Corey M. Williams, Stefan Bekiranov, and Aidong Zhang. 2025. IdeaBench: Benchmarking Large Language Models for Research Idea Generation. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2* (Toronto ON, Canada) *(KDD '25)*. Association for Computing Machinery, New York, NY, USA, 5888–5899. doi:10.1145/3711896.3737419

[7] Jiaxiong Hu, Xiwen Yao, Zeyu Huang, Danxuan Liang, Dongjie Yang, Dingdong Liu, Junze Li, Yuanhao Zhang, and Xiaojuan Ma. 2026. From Human Pragmatic Language Skills to Conversational Agent Design: A Systematic Review of Transfer Strategies. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems* (Barcelona, Spain) *(CHI '26)*. Association for Computing Machinery, New York, NY, USA, Article 911, 21 pages. doi:10.1145/3772318.3792787

[8] Yun Huang. 2026. "AI Scientists" with Productive Friction: Preserving Human Agency in AI-Mediated Research. In *ACM AI Leadership Summit*. ACM. https://aisummit.acm.org/accepted-papers Accepted.

[9] Yun Huang, Noshir Contractor, Tai-Quan Peng, Dashun Wang, Cuihua Shen, Manling Li, Yiren Liu, Hangyue Zhang, and Sangho Suh. 2026. Rethinking Social Computing Research in the Age of AI: Agency, Collaboration, and Governance. In *Proceedings of the 2026 Companion of the Computer-Supported Cooperative Work and Social Computing* (Salt Lake City, USA) *(CSCW Companion '26)*. Association for Computing Machinery, New York, NY, USA, 572–575. doi:10.1145/3785651.3831580

[10] Anna-Kaisa Kaila, André Holzapfel, and Petra Jääskeläinen. 2024. Gardening Frictions in Creative AI: Emerging Art Practices and Their Design Implications.. In *ICCC*. 186–195.

[11] Ji Eun Kim, Léa Vitale, Libby Hemphill, and Yulin Yu. 2026. Navigating the Changing Landscape of Online Knowledge Consumption and Production in the Age of Generative AI: Evidence from Stack Overflow. *arXiv preprint arXiv:2609.38563* (2026). arXiv:2609.38563 [cs.HC]

[12] Weixin Liang, Zachary Izzo, Yaohui Zhang, Haley Lepp, Hancheng Cao, Xuandong Zhao, Lingjiao Chen, Haotian Ye, Sheng Liu, Zhi Huang, Daniel A. McFarland, and James Y. Zou. 2024. Monitoring AI-modified content at scale: a case study on the impact of ChatGPT on AI conference peer reviews. In *Proceedings of the 41st International Conference on Machine Learning* (Vienna, Austria) *(ICML'24)*. JMLR.org, Article 1192, 46 pages.

[13] Zhehui Liao, Maria Antoniak, Inyoung Cheong, Evie Yu-Yen Cheng, Ai-Heng Lee, Kyle Lo, Joseph Chee Chang, and Amy X Zhang. 2025. LLMs as Research Tools: A Large Scale Survey of Researchers' Usage and Perceptions. In *Second Conference on Language Modeling*. https://openreview.net/forum?id=p0BwJk3R1p

[14] Yiren Liu, Si Chen, Haocong Cheng, Mengxia Yu, Xiao Ran, Andrew Mo, Yiliu Tang, and Yun Huang. 2024. How AI Processing Delays Foster Creativity: Exploring Research Question Co-Creation with an LLM-based Agent. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA). Association for Computing Machinery, New York, NY, USA, Article 17, 25 pages. doi:10.1145/3613904.3642698

[15] Yiren Liu, Viraj Nischal Shah, Sangho Suh, Pao Siangliulue, Tal August, and Yun Huang. 2026. Perspectra: Choosing Your Experts Enhances Critical Thinking in Multi-Agent Research Ideation. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. 1–27.

[16] Yiren Liu, Pranav Sharma, Mehul Oswal, Haijun Xia, and Yun Huang. 2025. Personaflow: Designing llm-simulated expert perspectives for enhanced research ideation. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference*. 506–534.

[17] Julio M Ottino and Brian Uzzi. 2026. Progression without progress. *Science* 392, 6800 (2026), 791–791.

[18] Kevin Pu, KJ Kevin Feng, Tovi Grossman, Tom Hope, Bhavana Dalvi Mishra, Matt Latzke, Jonathan Bragg, Joseph Chee Chang, and Pao Siangliulue. 2025. Ideasynth: Iterative research idea development through evolving and composing idea facets with literature-grounded feedback. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–31.

[19] Jan B Schmutz, Neal Outland, Sophie Kerstan, Eleni Georganta, and Anna-Sophie Ulfert. 2024. AI-teaming: Redefining collaboration in the digital era. *Current Opinion in Psychology* 58 (2024), 101837.

[20] Chenyang Shao, Dehao Huang, Yu Li, Keyu Zhao, Weiquan Lin, Yining Zhang, Qingbin Zeng, Zhiyu Chen, Tianxing Li, Yifei Huang, et al. 2025. OmniScientist: Toward a Co-evolving Ecosystem of Human and AI Scientists. *arXiv preprint arXiv:2511.16931* (2025).

[21] Chenglei Si, Diyi Yang, and Tatsunori Hashimoto. 2025. Can llms generate novel research ideas? a large-scale human study with 100+ nlp researchers. In *International Conference on Learning Representations*.

[22] Constantine Stephanidis, Gavriel Salvendy, Margherita Antona, Jessie YC Chen, Jianming Dong, Vincent G Duffy, Xiaowen Fang, Cali Fidopiastis, Gino Fragomeni, Limin Paul Fu, et al. 2019. Seven HCI grand challenges. *International Journal of Human–Computer Interaction* 35, 14 (2019), 1229–1269.

[23] Nitya Thakkar, Mert Yuksekgonul, Jake Silberg, Animesh Garg, Nanyun Peng, Fei Sha, Rose Yu, Carl Vondrick, and James Zou. 2025. Can LLM feedback enhance review quality? A randomized study of 20k reviews at ICLR 2025. *arXiv preprint arXiv:2504.09737* (2025).

[24] Jiaqi Wei, Yuejin Yang, Xiang Zhang, Yuhan Chen, Xiang Zhuang, Zhangyang Gao, Dongzhan Zhou, Guangshuai Wang, Zhiqiang Gao, Juntai Cao, et al. 2025. From ai for science to agentic science: A survey on autonomous scientific discovery. *arXiv preprint arXiv:2508.14111* (2025).

[25] Yulin Yu, Yong-Yeol Ahn, and Daniel M. Romero. 2026. Dataset repurposing and disruptive AI research. *arXiv preprint arXiv:2609.16736* (2026). arXiv:2609.16736 [cs.CY]

[26] Yulin Yu, Yan Chen, Rui Hu, Siddharth Suri, and Scott Counts. 2026. Adoption of Generative AI in the Workplace: Increasing and Shifting the Balance of Productivity and Communication Activity. *arXiv preprint arXiv:2608.15550* (2026). arXiv:2608.15550 [cs.HC]

[27] Yulin Yu, Yizhou Li, Siddharth Suri, and Scott Counts. 2026. From Searchable to Non-Searchable: Generative AI and Information Diversity in Online Information Seeking. In *Proceedings of the Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26)*. ACM, 1–6. doi:10.1145/3772363.3798802

[28] Yulin Yu and Daniel M. Romero. 2024. Does the use of unusual combinations of datasets contribute to greater scientific impact? *Proceedings of the National Academy of Sciences* 121, 41 (2024), e2402802121. doi:10.1073/pnas.2402802121

[29] Yulin Yu, Longqi Yang, Siân Lindley, and Mengting Wan. 2023. Large-Scale Analysis of New Employee Network Dynamics. In *Proceedings of the ACM Web Conference 2023 (WWW '23)*. ACM, 2719–2730. doi:10.1145/3543507.3583400

[30] Qingxiao Zheng, Minrui Chen, Pranav Sharma, Yiliu Tang, Mehul Oswal, Yiren Liu, and Yun Huang. 2025. EvAlignUX: Advancing UX Evaluation through LLM-Supported Metrics Exploration. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–25.