LLMs and the Infrastructure of CSCW

Chinmay Kulkarni
chinmay.kulkarni@emory.edu
Emory University
Atlanta, Georgia, USA

Tongshang Wu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Kenneth Holstein
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Q. Vera Liao
Microsoft Research
Montreal, Quebec, Canada

Min Kyung Lee
University of Texas at Austin
Austin, Texas, USA

Mina Lee
Microsoft Research, University of Chicago
Stanford, California, USA

Hariharan Subramonyam
Stanford University
Stanford, California, USA

ABSTRACT

Large Language Models have made many completing many previously difficult to achieve artificial intelligence tasks approachable to more programmers and non-programmers alike. More recently, open-source versions of large language models and the creation of new finetuning methods have been developed. These models and their decvelopment models lead this panel to discuss how the infrastructure of CSCW will influence LLM model development. It will also discuss how open source LLMs might influence CSCW research, and how they might allow the CSCW community to have new input into trust, safety, and responsibility in AI.

CCS CONCEPTS

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

KEYWORDS

large language models, collaboration, open source

ACM Reference Format:
Chinmay Kulkarni, Tongshang Wu, Kenneth Holstein, Q. Vera Liao, Min Kyung Lee, Mina Lee, and Hariharan Subramonyam. 2023. LLMs and the Infrastructure of CSCW. In *Computer Supported Cooperative Work and Social Computing (CSCW '23 Companion), October 14–18, 2023, Minneapolis, MN, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3584931.3608438

1 INTRODUCTION

Open-source projects have attracted the interest of the CSCW community not only as a social computing phenomenon, but a phenomenon of economic significance [4], that also achieves social outcomes of broader interest, such as online collective action, self-organization, and voluntary participation [2, 3, 6].

While early development in Large Language Models was dominated by large, well-funded private companies, highly capable open-source LLMs are increasingly becoming available (StableLM, sponsored by Stability AI; StarCoder, sponsored by HuggingFace; RedPajama by Together), to name just a few.

This panel will discuss the infrastructure of CSCW that makes open-source projects possible, and the effects of open-source large language models in turn on how we design collaborative systems. We hope to also have a balanced discussion contrasting these open-source development models with closed-sourced variants, such as how they are differently accessible to practitioners and end-users.

2 POTENTIAL TOPICS FOR DISCUSSION

We list a few topics for discussion, as suggested by panelists. We will engage with the audience to faciliate discussion on related questions, as discussed in Section 3.

2.1 Equity and access

Many have argued that LLMs can increase equity of participation in creating AI systems, by reducing technical barriers. In addition, LLM-augmented AI systems can also lower other barriers downstream (e.g., by expanding tutoring access, or helping non-native speakers to engage in programming AI systems). But as with other infrastructure, it is also possible that LLMs further increase existing inequities. In what ways should we design LLM-based systems towards more desirable outcomes?

A different thread of discussion will center around access: computing resources vary across institutions, and even countries. While open-source efforts may reduce the need for computing resources for using LLMs, developing such open-source models themselves may require large computational resources, limiting inclusion in such developer communities.

2.2 Explainability, Actionability, and Transparency

What are the design opportunities and challenges for helping people make informed decisions about LLM-infused products? Do existing methods for AI explainability suffice or we need new methods? A related strand of discussion centers around the specific problem of

CSCW '23 Companion, October 14–18, 2023, Minneapolis, MN, USA

ACM ISBN 979-8-4007-0129-0/23/10.
https://doi.org/10.1145/3584931.3608438

"hallucination" in LLMs – text generation that is not factual – and ways to handle it.

2.3 Using LLMs in CSCW

Several panelists have experience developing systems (for research and otherwise) using LLMs. One topic of discussion is around the value of prompt engineering (i.e. using in-context learning or instructing LLMs using a small number of examples.) Panelists will discuss whether this is a skill in much the same way as programming is, and whether it will remain valuable in the future.

Specific to CSCW, we will also discuss what roles can LLMs play in social interactions? For instance, what might LLM-infused social media or cooperative work applications look like? In addition to productive and unproductive use-cases of LLMs, we will also discuss possible new CCSW problems LLMs may create.

Finally, the application of LLMs for qualitative research in CSCW/HCI is increasingly common, along with skepticism on their applicability to thematic analyses and grounded theory approaches. What research problems might LLMs be applied to, and where might they be problematic?

2.4 Perceptions of LLM-generated materials

Because LLMs are pre-trained on predicting the next token, they may in some cases anchor everyone towards the same set of ideas [5]. This may be especially impactful when ideating with LLMs (or in general collaborating with them for completely novel ideas). How might interaction with LLMs lead to improving human creativity?

Similarly, as LLMs are increasingly used for work automation, they may change the value of the work they assist with. For instance, some already view AI-generated or assisted writing as inherently less valuable than than created by unassisted humans. At the same time, generative AI (such as text-to-image models) have also led to the creation of new kinds of art that are meaningful precisely because they were partly created by artificial intelligence [1]. How might the CSCW community engage in research that investigates these informal observations and improve our understanding of the value of work in the presence of AI?

3 PARTICIPATION PLAN

To facilitate a discussion-focused panel, we will primarily rely on the audience to suggest questions for discussion. Additionally, panelists will propose initial questions for each other, thereby seeding the discussion. To ensure ample time for discussion, panelists will provide brief introductions to their research as it relates to the panel, rather than presenting a comprehensive overview.

To gather questions, we will utilize an open and accessible online platform, with the link provided on the Panels page of the CSCW website and various social media platforms. For attendees at CSCW, we will also allow in-person questions during the panel. To promote equitable participation between online and in-person audience members, we will moderate the question stream. Due to the size of the panel, not all panelists will answer all questions. Instead, panelists will meet ahead of the panel, and informally "bid" on questions to which they would particularly like to contribute.

4 PANELISTS

Tongshang Wu is Assistant Professor at Carnegie Mellon University, where she conducts research on HCI and NLP. She studies in particular how humans (AI experts, lay users, domain experts) interact with (debug, audit, collaborate) AI systems. Her work seeks to design, evaluate, build, and interact with AI systems that are compatible with actual human goals.

Ken Holstein is an Assistant Professor in the Human-Computer Interaction Institute at Carnegie Mellon University, where he directs the CMU CoALA Lab. His research focuses broadly on AI-augmented work and improving the design and evaluation of AI systems for real-world use. He is deeply interested in understanding the gaps between human versus artificial intelligence across various real-world tasks and leveraging this understanding to develop systems that harness human expertise effectively. In his research, he integrates perspectives from human-computer interaction (HCI), AI, design, cognitive science, learning sciences, statistics, and machine learning.

Q. Vera Liao is a Principal Researcher at Microsoft Research Montréal, where she is a member of the FATE (Fairness, Accountability, Transparency, and Ethics of AI) group. Her work centers around human-AI interaction, with a particular focus on explainable AI and responsible AI. She is dedicated to bridging the gap between emerging AI technologies and human-centered design practices. To inform technology design, she employs both quantitative and qualitative methods to generate valuable recommendations. Prior to her role at MSR, she made significant research contributions at IBM T.J. Watson Research Center, which included advancements in IBM products such as AI Explainability 360, Uncertainty Quantification 360, and Watson Assistant.

Hari Subramonyam is an Assistant Professor (Research) at the Graduate School of Education and a Faculty Fellow at Stanford's Institute for Human-Centered AI. He is also a member of the HCI Group at Stanford. His research focuses on augmenting critical human tasks (such as learning, creativity, and sensemaking) with AI by incorporating principles from cognitive psychology. He also investigates support tools for multidisciplinary teams to co-design AI experiences.

Min Kyung Lee is an assistant professor in the School of Information at the University of Texas at Austin. She is affiliated with UT Austin Machine Learning Lab—one of the first NSF funded national AI research institutes, Good Systems—a UT Austin 8-year Grand Challenge to design responsible AI technologies, and Texas Robotics. Dr. Lee has conducted some of the first studies that empirically examine the social implications of algorithms' emerging roles in management and governance in society. She has extensive expertise in developing theories, methods and tools for human-centered AI and deploying them in practice through collaboration with real-world stakeholders and organizations. She developed a participatory framework that empowers community members to design matching algorithms that govern their own communities.

Mina Lee is a final-year Ph.D. student at Stanford University, specializing in Natural Language Processing and Human-Computer Interaction. Soon, she will be joining Microsoft Research as a post-doctoral researcher in the Computational Social Science group. Given her academic focus on AI and her love for writing, Mina

actively develops AI writing tools, explores the interaction between AI and writing, and investigates the impact of AI on the writing process.

Chinmay Kulkarni Chinmay Kulkarni is an Associate Professor in Computer Science at Emory University. Dr. Kulkarni's research interests contribute to the scientific fields of human-computer interaction, computer-supported cooperative work, and human-centered AI. While on a sabbatical at Google's PAIR lab (until July 2023), he helped develop Makersuite (Google's tool for writing LLM prompts), improved finetuning for the PaLM 2 model, and was a key contributor to Google's guides to help developers at Google and elsewhere prototype with LLMs. Dr. Kulkarni will moderate this panel.

5 DIVERSITY AND INCLUSION

While the panelists for this proposal are diverse in gender and ethnicity, their affiliation with US-based institutions may limit the inclusion of perspectives from other regions. Similarly, they have been trained largely in the learning sciences, computer science, and human-computer interaction, and may not represent other disciplinary perspectives adequately.

While we aim to be inclusive, we acknowledge potential biases in our perspectives and will draw upon the experiences of CSCW attendees when possible.

ACKNOWLEDGMENTS

ChatGPT was used to lightly edit the manuscript, fix syntax errors, and to shorten bios for panelists. Tthe authors accept responsibility for the veracity and correctness of all material in their work, including any computer-generated material.

Min Kyung Lee's participation is made possible by the National Science Foundation CNS-1952085, IIS-1939606, DGE-2125858, DASS-2217721 grants; Good Systems, a UT Austin Grand Challenge for developing responsible AI technologies.

REFERENCES

[1] Minsuk Chang, Stefania Druga, Alex Fiannaca, Pedro Vergani, Chinmay Kulkarni, Carrie Cai, and Michael Terry. 2023. The Prompt Artists. *arXiv preprint arXiv:2303.12253* (2023).

[2] Kevin Crowston, Qing Li, Kangning Wei, U Yeliz Eseryel, and James Howison. 2007. Self-organization of teams for free/libre open source software development. *Information and software technology* 49, 6 (2007), 564–575.

[3] Kevin Crowston, Kangning Wei, James Howison, and Andrea Wiggins. 2008. Free/Libre open-source software development: What we know and what we do not know. *ACM Computing Surveys (CSUR)* 44, 2 (2008), 1–35.

[4] Nadia Eghbal. 2016. Roads and bridges. *The Unseen labor behind our digital infrastructure* (2016).

[5] Sophie Jentzsch and Kristian Kersting. 2023. ChatGPT is fun, but it is not funny! Humor is still challenging Large Language Models. *arXiv preprint arXiv:2306.04563* (2023).

[6] Eric Raymond. 1999. The cathedral and the bazaar. *Knowledge, Technology & Policy* 12, 3 (1999), 23–49.