Human-AI Collaboration: How AIs Augment Human Teammates

Tongshuang Wu, Diyi Yang, Joseph Chang, Marti A. Hearst, and Kyle Lo

Website: https://acl25-hai-team.github.io/

The continuous, rapid development of general- purpose models like LLMs suggests the theoretical possibility of AI performing any human task. Yet, despite the potential and promise, these models are far from perfect, excelling at certain tasks while struggling with others. The tension between what is possible and a model's limitations raises the general research question that has attracted attention from various disciplines: What is the best way to use AI to maximize its benefits? In this tutorial, we will review recent developments related to human-AI teaming and collaboration. To the best of our knowledge, our tutorial will be the first to provide a more integrated view from NLP, HCI, Computational Social Science, and Learning Science, etc., and highlight how different communities have identified the goals and societal impacts of such collaborations, both positive and negative. We will further discuss how to operationalize these Human-AI collaboration goals, and reflect on how state-of-the-art AI models should be evaluated and scaffolded to make them most useful in collaborative contexts.

Sherry Wu, Assistant Professor, Carnegie Mellon University
Email: sherryw@cs.cmu.edu
Website: http://cs.cmu.edu/∼sherryw
Sherry Wu is an assistant professor at the Human-Computer Interaction Institute, Carnegie Mellon University. Her primary research investigates how humans (AI experts, lay users, domain experts) interact with (debug, audit, and collaborate) AI systems. Sherry has organized two workshops at NLP and HCI conferences: Shared Stories and Lessons Learned workshop at EMNLP 2022 and Trust and Reliance in AI-Human Teams at CHI 2022 and 2023. She has given two well-received tutorials relevant to Human-AI Interaction, one at EMNLP 2023 on Designing, Learning from, and Evaluating Human-AI Interactions, and another one on Human-AI Interactions in the Era of LLMs at NAACL 2024.

Diyi Yang, Assistant Professor, Stanford University
Email: diyiy@stanford.edu
Website: https://cs.stanford.edu/∼diyiy/
Diyi Yang is an assistant professor in the Computer Science Department at Stanford University. Her research focuses on human-centered natural language processing and computational social science. Diyi has organized four workshops at NLP conferences: Widening NLP Workshops at NAACL 2018 and ACL 2019, Casual Inference workshop at EMNLP 2021, NLG Evaluation workshop at EMNLP 2021, and Shared Stories and Lessons Learned workshop at EMNLP 2022. She also gave a tutorial at ACL 2022 on Learning with Limited Data, and a tutorial at EACL 2023 on Summarizing Conversations at Scale. Diyi and Sherry have co-developed a new course on Human-Centered NLP that has been offered at both Stanford and CMU.

Joseph Chee Chang, Research Scientist, AI2
Email: josephc@allenai.org
Website: https://joe.cat/
Joseph Chee Chang is a research scientist at the Allen Institute for AI, where he study and design novel Human-AI systems and user interfaces to facilitate sensemaking. His recent projects include interactive human-agent planning and execution, intelligent and interactive reading interfaces for scholarly documents, and other research support tools including systems

Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 5: Tutorial Abstracts), pages 5–6

for ideation and literature review.

Marti A. Hearst, Professor, UC Berkeley
Email: hearst@berkeley.edu
Website: https://people.ischool.berkeley.edu/~hearst/
Marti A. Hearst is a professor and the Interim Dean for the UC Berkeley School of Information. She is both an ACL Fellow and a SIGCHI Academy member, and former ACL President. Her research has long combined HCI and NLP; recent projects include adding interactivity to scholarly documents and creating interactive newspods. She recently gave invited keynote talks at the EACL NLP + HCI workshop, the KDD Workshop on Data Science with a Human in the Loop, and she advised the 2022 NAACL program chairs on the Human-Centered Natural Language Processing special theme. She has taught courses in NLP, HCI, and information visualization for 25 years.

Kyle Lo, Research Scientist, AI2
Email: kylel@allenai.org
Website: https://kyleclo.com/
Kyle Lo is a research scientist at the Allen Institute for AI working on natural language processing, machine learning and human-AI interaction, with emphasis on the impact of training data on model behavior, evaluation methodology, and intelligent reading interfaces. Kyle has organized four workshops on NLP for scholarly documents, including at NAACL 2021, COLING 2022 and AKBC 2020-2021, as well as three shared tasks on scientific information retrieval and fact checking at TREC 2021 and TAC 2020. Kyle also presented a tutorial on language model development at NeurIPS 2024.