

I am interested in Educational Technologies, Computer-Supported Collaboration, Computer-Mediated Communication, and Social Computing. My broad goal is to build systems that make it easier for people to learn, work, and access information. In order to inform the design of novel systems, I am keen on studying how people communicate and behave while engaging in these activities. As an undergraduate, I have pursued interdisciplinary research with a focus on Human-Computer Interaction (HCI) and Social Computing and have worked on information summarization, conversational systems, crowdsourcing, and behavioral studies. The Ph.D. program at Carnegie Mellon University's HCII will be the perfect opportunity for me to continue to further my knowledge in these areas and work on impactful and challenging problems in HCI. Being part of the large and interdisciplinary community of HCI researchers at CMU will expose me to diverse perspectives which, along with the experience and skills I can acquire during the program, will equip me well to pursue a career in HCI research.

Bridging the gap between information online and how users reason about it is crucial to help people assimilate and learn from information effectively. People consult reviews and blogs to make decisions such as where to eat and what to purchase, requiring comparison between options. Most blogs and articles online, however, detail a single entity only. To improve the organization of information available for decisions, I worked with Prof. Plaban Bhowmick at IIT Kharagpur on generating comparisons between pairs of entities using their Wikipedia articles. I experimented with several topic models and clustering algorithms, only to find that they failed to cluster comparable sentences when documents were longer than a few sentences and comparisons were sparse. To enable extraction of such sparse comparisons from large unstructured text, I built a system to construct a graph-based representation of the documents with nodes representing sentences and edge weights corresponding to semantic similarity. Using heuristics to retain edges and a semantic similarity classifier to filter the comparisons further, we were able to effectively extract comparisons. This project prompted my further interest in using language technologies to augment learning and work.

Since communication and discussions are central to how people collaborate and consume information, I began looking into methods to enhance communication in groups as well as between users and computing systems. I worked with Dr. Monojit Choudhury, at Microsoft Research, on making conversational systems more engaging by leveraging cultural context. Since multilingual users mix languages in conversations- a phenomenon called code-mixing - we studied whether they would prefer chatbots that code-mix. We devised two computational code-mixing policies, rooted in Linguistics and Communication theory, for chatbots to code-mix systematically and designed a user study for the Hindi-English language pair. I prototyped our chatbots as human-in-the-loop systems around an English chatbot and built the study platform. We trained 3 human wizards who paraphrased users' code-mixed messages to English inputs for the bot and introduced code-mixing to the bot's responses guided by our policies. Analyzing participants' responses, we found that they strongly preferred chatbots that code-mix and individual differences were a strong predictor of user preferences within the policies. We submitted a paper to CSCW detailing our findings and proposing design practices [1]. While my previous project explored ways to improve how information is communicated to users, this project exposed me to the complementary question of how to enable users to communicate with computing systems more intuitively. As a Ph.D. student I am interested in exploring these questions further to build novel interactive systems.

Designers of social computing systems often draw insights from interactions between people. However, as AI systems permeate virtual spaces and increasingly mediate communication and collaboration, I believe it is equally important to study interactions between people and these AI systems to improve the design of computing systems. I explored these ideas as an intern in Stanford's HCI Group with Prof. Michael Bernstein. I joined a team working on increasing volunteer participation in crowdsourcing by rephrasing requests using cues from social psychology. I helped set up and run experiments that used bots on Instagram to ask people questions about their photos in order to curate visual datasets. In our paper presented at HCOMP [3], we showed that rephrasing these questions to mimic how humans seek help, increases response rates. During this project, I became curious about what aspects of AI systems affect user attitudes towards them. Why was Tay- described using the metaphor "AI that's got no chill"- subject to anti-social behavior while Xiaoice-"the empathetic ear"- was a hit with its users, though these systems were largely similar? Understanding how users reason about systems can facilitate redesigning to afford more satisfying interactions.

Along these lines, I led a project, while at Stanford, to understand how such metaphors associated with AI systems affect user attitudes towards them. We hypothesized that metaphors implicitly shape expectations and thus affect user-attitudes. To test this, I sampled different metaphors varying along the axes of 'warmth' and 'competence' which are identified as the main factors of impression formation in social cognition literature. We deployed a study online where participants were introduced to an AI agent using one of the metaphors and then completed a task with a Wizard-of-Oz agent. I measured the effect of different metaphors on participants' pre-use expectations and post-use evaluations of the AI-agent. We found that users were more likely to adopt agents that projected low competence and were more likely to co-operate with agents that signaled low competence and high warmth. We compiled our findings and showed how they extend to existing systems in a submission to CHI [2].

My experience working on these projects has motivated me to pursue a PhD and subsequently a career in HCI research. At CMU, I am particularly interested in the work of **Prof. Amy Ogan**, **Prof. Carolyn Rosé**, and **Prof. Chinmay Kulkarni**. Prof Rosé's and Prof. Ogan's work on incorporating virtual agents into group work and learning environments to improve outcomes and provide more engaging learning experiences, aligns naturally with my interests. Prof. Ogan's work on designing culturally appropriate educational technology also meshes with my interests. Similarly, I am interested in working with Prof. Kulkarni on studying how different aspects of learning communities, platforms, and instructional resources affect student experiences and using these insights to build new systems and interventions to improve learning outcomes. His recent work on voice-based agents and their use in augmenting work, also interests me. Given my interests, I believe the Ph.D. program in HCI would be a great fit for me and I hope to continue growing as a researcher at CMU.

[1] Anshul Bawa, Pranav Khadpe, Pratik Joshi, Kalika Bali and Monojit Choudhury. Do Multilingual Users Prefer Chat-bots that Code-mix? Let's Nudge and Find Out! 2019. (*In submission to CSCW 2020*)

[2] Pranav Khadpe, Ranjay Krishna, Li Fei-Fei, Jeff Hancock and Michael Bernstein. Low Expectations Lead to Better Experiences: The Effect of Conceptual Metaphors on Human-AI Collaboration. 2019. (*In submission to CHI 2020*)

[3] Junwon Park, Ranjay Krishna, Pranav Khadpe, Li Fei-Fei and Michael Bernstein. AI-Based Request Augmentation to Increase Crowdsourcing Participation. 2019. *In Proceedings of the HCOMP 2019*