

The future of Social AR certainly will include exciting new systems that explicitly provide information about people we encounter and interact with, but one aspect of Social AR that I am particularly interested in exploring is the potential to transform the ambience of a social encounter with less explicit cues to mediate conversations in a less distracting and potentially less "creepy" way.

This concept might also be applicable to AR systems that augment what everyone in a room sees and hears, but I am especially curious about the implementation of "ambient AR" in a way that is only perceivable to a single person (accomplished using personal glasses-based devices, for instance). I believe that AR can be a valuable learning tool for individuals who want to become more aware of and address their own implicit biases but find it difficult to do so tactfully in social situations. An AR system might be used to help a user identify when they unintentionally use language that is offensive or insensitive to certain groups -- perhaps by creating overlays that might obscure the user's vision when problematic language is uttered or highlight particular people in the room who might feel marginalized. Such a system could also be used to help a user become more aware when others are using problematic language or gestures as well, highlighting those individuals when such situations arise.

There are certainly challenges and many open design questions when it comes to implementing such a system. How do we develop the vocabulary around what the system should react to? What is the most effective way of augmenting the user's "reality" to productively raise awareness and support self-improvement for the user? Is it even possible for such systems to be effective for increasing understanding of underlying complex biases? Answering such questions requires careful consideration of social, ethical, and privacy concerns. For instance, if we were to build a system that identifies particular people in a room who might be particularly affected by a user's remarks, ideally, the system would base such identification on a pre-populated list of "offensive terms" that each person provides, which is an unrealistic requirement. Alternatively, a system might automatically glean what topics, words, or actions might be construed as offensive to a person in the room solely from outward traits and behaviors alone, but of course, this inference strategy and its consequences can be extremely problematic as well. Still, I remain optimistic that Social AR has the potential to find applications in such a space and would be curious to discuss related possibilities with others in this workshop.

Katherine Song, PhD student, UC Berkeley (kwsong@berkeley.edu)