
Human-Agent Collaboration: Can an Agent be a Partner?




Figure 1: Robots hold great promise as expert advisor, but as collaboration partners they need to be *trusted* experts [1].

Rachel K. E. Bellamy

IBM T. J. Watson Research
PO Box 218, Yorktown Heights,
NY 10598, USA
rachel@us.ibm.com

Sean Andrist

Microsoft Research,
Redmond, WA, USA
sandrist@microsoft.com

Timothy Bickmore

Northeastern University
Boston, MA 02115, USA
bickmore@ccs.neu.edu

Elizabeth F. Churchill

Google Inc.
Mountain View, CA, USA
churchill@acm.org

Thomas Erickson

IBM T. J. Watson Research
PO Box 218, Yorktown Heights,
NY 10598, USA
snowfall@us.ibm.com

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).
CHI'17 Extended Abstracts, May 6–11, 2017, Denver, CO, USA.
ACM ISBN 978-1-4503-4656-6/17/05.
<http://dx.doi.org/10.1145/3027063.3051138>

Abstract

Human-Agent Interaction has been much studied and discussed in the last two decades. We have two starting points for this panel. First we observe that *interaction* is not the same as *collaboration*. *Collaboration* involves mutual goal understanding, preemptive task co-management and shared progress tracking. Second, that much of the on-going discourse around human-agent interaction implies that agents can be *trusted*, collaborative partners. Our position is that while virtual and embodied agents have the potential to be work partners, for this goal to be achieved we need to better understand what partnership in collaboration involves. In this panel we ask: Can this potential for trusted collaboration be realized? If so, what will it take? This panel will bring together HCI experts who work on collaboration, virtual agent design and human-robot-interaction. Panelists will engage the audience through discussion of their shared and diverging visions, and through suggestions for opportunities and challenges for the future of human-agent collaboration.

Author Keywords

Virtual agents, collaborative partnership, robots, AI.

ACM Classification Keywords

(HCI): Interaction design process and methods, Artificial intelligence




Figure 2: Web-based automated customer service agent [image attribution:10].




Figure 3: A comforting agent offers support in stressful situations [3].

Introduction

In “Man-Machine Symbiosis,” J. C. R. Licklider (1960) originated the concept of symbiotic computing [8]. Licklider stressed that this kind of computer-supported cooperation was important because many problems: *“... are very difficult to think through in advance. They would be easier to solve, and they could be solved faster, through an intuitively guided trial-and-error procedure in which the computer cooperated, turning up flaws in the reasoning or revealing unexpected turns in the solution”*

Licklider's vision was remarkably prescient. The notion of a *computer as a partner* is an increasingly common metaphor for interaction design. Some examples are: trusted expert advisors (Fig 1), customer service agents (Fig 2), robots that motivate people in therapeutic activities (Fig 3) or engage in shared tasks (Fig 4). Partners need not necessarily be embodied: smart rooms include agents that can help business users make decisions (Fig 5), and increasingly consumer gadgets for the home utilize voice-based, service “avatars” whose interactions and agency appear increasingly collaborative. The vision in all these instances is about a synergistic, mutually amplifying relationship, not simply mechanistic assistants that do our bidding, responding to simple commands with a limited set of pre-scripted actions. As partners, these agents will help us get tasks done faster and perhaps also allow us to do more than we could alone. To achieve this vision of human-agent collaboration, we need to better understand what is required in terms of collaboration and cooperation capabilities, and how we can build on the current somewhat limited state of the art.

In this panel, experts will discuss what it means to design computers that *partner* with us, rather than services that offer “assistance”. We will explore different perspectives on how autonomous agents/assistants could partner with us, and could thus change how we work and what it means to think together. We will use the theme of the conference—Explore, Innovate, Inspire—to structure the panel.

We will start by exploring the differences between assistant, assistance and partner. For example:

- What is the difference between a system being an assistant and being a partner?
- What aspects of human collaboration should be used as models for human-agent collaboration?
- Does the form of an agent make a difference (e.g., physical robot, text chatbot, spoken dialog, etc.). If so, how?
- What are appropriate roles for agents: leader? facilitator? participant? advisor? How do those roles affect the dynamics of group interaction?
- Should agents mimic human emotion and social conventions when collaborating?
- How aware will agents need to be of human emotions, conventions and mental models?
- Since trust is crucial for collaboration, how will agents affect trust within human-agent teams?
- What ethical considerations are there for agent collaborators? Should they be trained to withhold information, to be diplomatic or, if necessary for the social good, be deceptive?




Figure 4: An engaging robotic partner that matches its personality to that of its partner via its gaze behavior [2].

We will then address innovation:

- Which design processes work best for creating agents that collaborate with humans?
- Which metaphors are appropriate for portraying agents? How should agents be depicted in terms of team roles?
- What are the current development and user testing practices for agents that are already on the market?

Audience members will be invited to take part in an online conversation mediated by a virtual discussion agent before, during, and after the panel.

Panel Moderator

Rachel Bellamy manages the Human-Agent Collaboration group at IBM T. J. Watson Research Center. The group investigates how to design intelligent agents that support human collaboration. The group specializes in designing agents embodied in smart environments, and everyday objects. In these cases there are typically multiple users interacting with the system synchronously or asynchronously. The designs account for multi-user interaction, and awareness of other users and their activities.

Panelists

The panelists are top researchers from the CHI community, experts in collaboration and/or agents.¹

¹ We invited an agent to participate on the panel, but unfortunately it had already committed to another panel titled "Agent-Human Collaboration: Can a Human be a Partner?"

Panelists hold a variety of viewpoints about the intelligence and capabilities of digital agents. Some think digital agents are far from smart and never will be. Lack of smarts may have advantages however, for example talking to a 'stupid' agent rather than a sales person, may be preferable because the agent won't intentionally mislead or lie to you. Other panelists have spent many years designing agent interaction to make it acceptable to consumers, or to give agents capabilities that enhance emotional and social interactions. Our panelists also have differing viewpoints on the form agents should take. Some have worked on embodied agents that take human form either in virtual worlds or as robots. Others have created dis-embodied agents that are part of smart buildings and objects.

Brief biographical sketches and proposed panel statement topics for each panelist appear in alphabetical order below.

Sean Andrist, Microsoft Research, Adaptive Systems and Interaction group.

Sean Andrist is a researcher at Microsoft Research in the Adaptive Systems and Interaction group. He received his PhD in 2016 from the Department of Computer Sciences at the University of Wisconsin–Madison, where he conducted research on gaze mechanisms for the development of communicative characters, including both embodied virtual agents and social robots. His current research interests involve designing, building, and evaluating socially interactive technologies that are physically situated in the open world. To be truly situated, these technologies must be able to track, interpret, and respond to an array of




Figure 5: A smart room that facilitates decision makers [7].

multimodal signals from the human user or users, including speech, gesture, gaze, posture, and so on. Furthermore, these technologies must act in real-time while reasoning about uncertainty from noisy sensors and ambiguities inherent to human social interactions. By embodying these technologies as virtual agents on a screen or as physical robots, they will also gain the ability to produce human-understandable verbal and nonverbal signals in order to hold more natural, effective, and rewarding interactions.

Timothy Bickmore, Professor, College of Computer and Information Science, Northeastern University.

Dr. Bickmore's research focus is on the development of agents and robots designed to build long-term, social-emotional relationships with their users. These agents have been deployed within the context of behavior change interventions in which they are designed to establish working alliance relationships with patients in order to maximize health intervention outcomes. Prior to joining Northeastern, Dr. Bickmore was an Assistant Professor of Medicine at the Boston University School of Medicine. Dr. Bickmore received his PhD from the MIT Media Lab, where he studied social and emotional interactions between people and virtual agents.

His work is mostly centered in the medical domain, simulating face-to-face counseling sessions between health providers and patients, including the verbal and nonverbal behavior used in these interactions. Most of the work uses virtual agents (over 25 clinical trials on a wide range of platforms for a wide range of health problems), but is now exploring the use of anthropomorphic robots in health counseling. He is particularly interested in longitudinal interactions, and

behaviors that establish a sense of trust and bonding that lead to greater system use, intervention retention, adherence, and patient satisfaction.

Elizabeth Churchill, Director of User Experience, Google

Currently a Director of User Experience at Google, Dr. Elizabeth Churchill is an applied social scientist working in HCI, computer mediated communication, mobile/ubiquitous computing and social media. Originally a psychologist by training, throughout her career Elizabeth has focused on understanding people's social and collaborative interactions in their everyday digital and physical contexts. She has studied, designed and collaborated in creating online collaboration tools (e.g. virtual worlds, collaboration/chat spaces), applications and services for mobile and personal devices, and media installations in public spaces for distributed collaboration and communication.

Elizabeth has long been skeptical about most of the envisionments for digital agents over the years. Based on her own research experience in designing and evaluating "embodied conversation agents" in the late 1990's and early 2000's, her biggest concern has been with the "cliffs" of capability that necessarily exist with all agents who are "service avatars"—that is, front-ends to limited knowledge bases and the limited learning capabilities agents have displayed to date [4][9]. The key questions about which she remains curious have to do with agent identity, interrogability, allegiance and autonomy: this can be summed up as a critical and technical design focus on the agency of the agent.

Thomas Erickson, Social Computing, IBM Research

Thomas Erickson is an interaction designer and social scientist at IBM Research. He is concerned with designing systems that enable groups of people to interact coherently and productively. Tom's interest in agents and agency dates back two decades to the article *Designing Agents as If People Mattered* [5], where he pursued the question of how to *limit* users' expectations of agents. His basic premise was (and is) that digital agents are, to be blunt, stupid. They lack, and will continue to lack, common sense and the ability to reason, learn, remember and gracefully interact in ways that ordinary people assume is a concomitant of ordinary intelligence. To be sure, digital agents will be extraordinarily knowledgeable in very restricted domains — but this only exacerbates the problem of what happens when the user steps off the plateau of expertise and plummets into a vast swamp of incompetence. People are not accustomed to such discontinuities of competence. Tom is interested both in ways of depicting agents — he sees service animals as a possible metaphor — and in ways of structuring their interactions so that their discontinuities of competence are understandable. He does *not* believe that agents are, or will be, smart enough to serve as assistants (a role that requires considerable general intelligence and collaborative abilities [6]), or partners (an even more demanding role).

Panel Format and Audience Engagement

The panel will begin with brief (3-4 minute) position statements from each of the panelists focusing on what he/she sees as the key challenges and opportunities for robots as collaborative partners. To ground the conversation, panelists will show video examples of

agents and discuss design choices made with respect to how the agent partners with humans. There will also be a live Q&A session in which members of the audience can pose questions to the entire panel.

Before, during and after the event, the audience members and panelists will be able to participate in an on-line collaborative discussion facilitated by a virtual discussion assistant. We will invite the audience to use this discussion forum to inspire the community with their thoughts on the future of human-agent collaboration. This will also give the audience experience with a specific example of human-agent collaboration.

References

1. Andrist, S., Ziadee, M., Boukaram, H., Mutlu, B., and Sakr, M. (2015). Effects of Culture on the Credibility of Robot Speech: A Comparison between English and Arabic. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI '15)*. ACM, New York, NY, USA. 157-164.
2. Andrist, S., Mutlu, B. and Tapus, A. (2015). Look Like Me: Matching Robot Personality via Gaze to Increase Motivation. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 3603-3612.
3. Bickmore, T., Fernando, R., Ring, L., Schulman, D. (2010). Empathic Touch by Relational Agents. *IEEE Transactions on Affective Computing* 1(1), 60-71.
4. Churchill, Elizabeth F., et al. "May I help you?": designing embodied conversational agent allies." *Embodied conversational agents*. MIT Press, 2001.
5. Erickson, T. *Designing Agents as if People Mattered*. *Intelligent Agents* (ed. J Bradshaw). AAAI Press: 1997.

6. Erickson, T., Danis, C., Kellogg, W. A., Helander, M. E. Assistance: The Work Practices of Human Administrative Assistants and their Implications for IT and Organizations. Proc. CSCW '08. ACM Press, 2008.
7. Farrell, R., Lenchner, J., Kephjart, J.O., Webb, A.M., Muller, M.J., Erikson, T.D., Melville, D.O., Bellamy, R.K., Gruen, D.M., Connell, J.H. and Soroker, D., 2016. Symbiotic Cognitive Computing. *AI Magazine*, 37(3), pp.81-93.
8. Licklider, J. C. R. 1960. Man-Computer Symbiosis. *IRE Transactions on Human Factors in Electronics* Volume HFE-1(1): 4–11.
9. Prevost, S., Hodgson, P., Cook, L., and Churchill, E.F. "Face-to-face interfaces." *CHI'99 Extended Abstracts on Human Factors in Computing Systems*. ACM, 1999.
10. Creative Commons, image attribution: Bemidji State University