

Exploring the Implications of Virtual Human Research for Human-Robot Teams

Jonathan Gratch¹(✉), Susan Hill², Louis-Philippe Morency³,
David Pynadath¹, and David Traum¹

¹ University of Southern California Institute for Creative Technologies,
Los Angeles, USA

{gratch, pynadath, traum}@ict.usc.edu

² Army Research Laboratory, Adelphi, USA

susan.g.hill.civ@mail.mil

³ Carnegie Mellon University, Pittsburgh, USA

morency@cs.cmu.edu

Abstract. This article briefly explores potential synergies between the fields of virtual human and human-robot interaction research. We consider challenges in advancing the effectiveness of human-robot teams makes recommendations for enhancing this by facilitating synergies between robotics and virtual human research.

1 Introduction

Advances in autonomy raise the potential for rich partnerships between humans and machines. Human-robot teams are emerging across a range of high-stakes situations including military operations, first-responders and caring for vulnerable populations. To date, the preponderance of robotics research addresses the challenge of individual robots interacting with the physical environment, yet teamwork involves navigating a social environment. To address this gap, research on human-machine teams is increasingly turning to the social sciences to inform the design of automation. For example, the research into how to get users to trust automation increasingly builds on theories of how trust arises between people [1, 2]. Similarly, research into human-computer communication builds on theories of human verbal and nonverbal communication, often incorporating into automation analogs of facial expressions or bodily gestures [3, 4].

Whereas the robotics community is beginning to explore the role of social science theory and anthropomorphic techniques, these elements are the *raison d'être* for the field of virtual human research (e.g., see [5, 6] and the Intelligent Virtual Agent's conference series). Virtual humans are software artifacts that look like, act like and interact with humans but exist in virtual environments. To achieve this, the virtual humans must provide a sufficient illusion of human-like behavior that human users interpret, respond, and learn from such virtual interactions much as they would react in real-world social interactions. To this end, virtual humans must be responsive; that is, they must respond to the human user and to the events surrounding them. They must be

interpretable; the user must be able to interpret their response to situations, including their dynamic cognitive and emotional state, using the same verbal and non-verbal cues that people use to understand one another. Finally, they must evoke similar social effects as are expected to occur in face-to-face interactions (e.g., social anxiety, impression management, emotional contagion). Thus, the virtual humans cannot simply create an illusion of life through cleverly designed randomness in their behavior; they are successful to the extent that they evoke responses from humans indistinguishable from how people would respond to another person.

In this article, we highlight the need for greater collaboration between robotics and virtual human research. Virtual human research has placed a premium on how to understand, model and simulate spoken language, how to recognize and utilize non-verbal communication, and how to model and utilize social cognitive processes such as intention recognition, collaborative decision-making, and even the role of emotion in teams. Many of these capabilities are of relevance in human-robot interaction. Yet collaboration is required to translate these findings to the domain of robotics. Virtual human research is usually explored in “pristine” simulated or laboratory settings that finesse many of the challenges of operating in complex real-world environments. More fundamentally, the goal of much virtual human research is to literally replicate human appearance and behavior, yet this is less possible, and potentially less desired within the context of physical robots. Rather, a collaboration is required to understand which approaches are relevant and which are relevant in an analogous, if not literal form.

Here, we recommend several potentially fruitful points of interaction between virtual human and robotics research as it relates to the challenge of mixed human-machine teams. These include 1) research into the potential benefits but also pitfalls of incorporating anthropomorphic elements into robotic systems; 2) research aimed at transitioning natural language and 3) nonverbal communication techniques developed within the virtual human community into settings that involve human-machine teams; and 4) research into technology for enhancing trust in human-machine teams, including methods for automatically generating explanations of machine decisions and establishing shared mental models. Each of these recommendations are presented in the following sections of this article and are discussed in much greater detail in a technical report [7].

2 Cost/Benefit of Anthropomorphism in Human-Robot Teams

A growing trend within robotics has focused on endowing robots with more human-like characteristics, including human-like form, natural language and even emotions. This interest is fueled by the assumption that human-robot and human-computer interactions can be enhanced by bringing how we interact with machines closer to how we interact with other people, thereby leveraging the vast experience we have with human-human interaction. This assumption needs to be rigorously examined.

It is important to note that machines can be made more capable without necessarily making them more “natural” or human-like. Interaction with robots could be explicitly unnatural as naturalness might get in the way of efficiency (for example, communication

with air-traffic controllers is highly scripted to be efficient while avoiding ambiguity). Research on natural interfaces demonstrates that machines can be made more human-like, but less research has considered if this benefits or harms human-machine team performance. Indeed, a review of the literature illustrates several examples where incorporating human-like qualities results in unintended and disruptive consequences. Attempts to merely replicate human characteristics overlooks an opportunity to improve on human-human interaction: might machines be designed to interact in different but complementary ways that make them better than “natural” teammates?

Some research has emphasized the potential benefits of anthropomorphism. For example, Gratch and colleagues have shown that a computer agent that incorporates rapport-building behaviors can enhance feelings of engagement and lead to greater self-disclosure in spoken interviews [8]. People favor human-like machines in economic settings when they incorporate human-like features, such as offering more money in a variety of economic games [9], and donating more money when asked by a human-like robot [10]. People are more persuaded by machines that incorporate human-like gestures [11] or humor [12]. Students have been shown to learn better when automated tutors incorporate emotional feedback [13]. Other research has shown that adding that adding human-like mental capabilities, like theory of mind, can improve joint outcomes in social games [14]. Many of these findings have been replicated within the context of human-robot interaction (e.g., [2, 15]).

Yet, other research has emphasized the potential harms of anthropomorphism. People lie to human-like machines, they get emotional, they make “irrational” decisions, and they evoke moral principles that get in the way of maximizing material rewards. For example, in medicine, it is important for healthcare providers to solicit honest information from their patients. Yet patients are more honest when being interviewed by a computer compared with being interviewed by a person [16]. Anthropomorphism can undermine this benefit by evoking the social mechanism of socially-desirable responding. Lucas, et al., showed a depression screening agent elicited more truthful and more diagnostic information when its “computerness” was made salient compared with an agent that emphasized its “humanness” [17]. In economic settings, people often make financially disadvantageous decision with human teammates [18] or human-like computers [9] when compared to decisions with a machine. More generally, people engage in more emotion, moral and reactive decision-making with other people compared to their interactions with computers, leading to a host of negative outcomes, especially in conflict situations [19, 20].

Opportunities for Research: Incorporating natural and anthropomorphic characteristics into robotic systems can have a strong impact on human-robot team effectiveness. Unfortunately, these effects can be both beneficial and harmful depending on a variety of task, contextual and individual factors. More research is needed regarding when and how anthropomorphism benefits human-machine systems. Specifically, we recommend research directed at selectively evoking social effects: In that human-like traits unconsciously evoke human-like responses, and that some responses have benefits but others harms, research is needed in how to distinguish and differentially evoke specific social effects that lead to benefits, while avoiding the evocation of disruptive social effects.

3 NLP for Virtual Humans and Robots

Virtual humans and robots are both artificial, automated agents that can engage in complex behavior and complex interaction with humans. Natural language is one of the main ways that humans communicate with each other, particularly for abstract concepts, processes, or objects that are not immediately visible or manipulable. By engaging in natural language dialogue, automated agents can make it easier to communicate with people by making use of this same communication method. There is a large degree of overlap in the kinds of tasks that robots and virtual humans can talk to people about.

Many issues make natural language processing difficult for automation, including noisy input, vagueness, and the contextual meaning of utterances. For most tasks, there has been more progress with virtual humans than robots in overcoming these challenges – both because less effort must be spent in creating the basic interactive capability with animations rather than robots with complex physical components, and because language generation in the virtual world finesses challenges with real world-perception, using instead meta-data or virtual world databases for perceptual information.

Much of the work on virtual human natural language dialogue can be adapted for improving human-robot natural language dialogue. For example a key problem for both domains is navigating through a complex environment and giving and understanding directions. Some examples of virtual human work include virtual characters on a mobile device who gives tours of a museum exhibit [21], and the GRUVE challenge on generating instructions in an urban environment [22]. Another shared problem is the Grounding problem, which involves coordination of interlocutors using multi-modal dialogue interaction to increase confidence of shared understanding. The computational models developed in [23] have been implemented and used within a number of virtual human systems (e.g., [24]). Another point of intersection is runtime and support software tools, authoring tools, and the development process that can be applied across domains. For example, speech recognizers, parsers, statistical classifiers [25], dialogue managers [26], language generators [27], and speech synthesizers [28], many freely available through the virtual human toolkit [29].

Finally, the development process itself is an area where robotics can exploit work pioneered in Virtual Human efforts. Natural language components need considerable training data to achieve high performance, but gathering this data is challenging for dialogue interaction, where the things people say to an artificial agent are determined by what the agent says and does. Thus, in order to gather the appropriate data, one already needs the system. The way out of this conundrum is a phased approach to data collection: beginning with purely human interaction, next moving to “wizard of oz” collection (where an agent is controlled by a human behind the curtain). Finally, versions of an automated system can be deployed and improved. A number of virtual human projects have followed this development path (e.g., [30, 31]).

Several challenges confront the use of virtual human technology in robotic systems. One of the greatest strengths for natural language processing for virtual humans in the virtual world is the ability to simplify the non-linguistic issues, such as perception,

locomotion and manipulation. However, this simplification can also become a limitation, since it may not be straightforward to adapt algorithms tuned to the simplified environment to work in the real environment. Another challenge is that the roles for virtual humans and robots may tend to diverge, which, in turn, may tend to cause a divergence in the kinds of language used, and thus the best algorithms and tools. Virtual humans are generally meant to take the place of a real human in a social interaction, so communication is generally using the “agent as human” metaphor for communication. In cases where a robot is very non-human in appearance, perception and manipulation capabilities, and purpose, this metaphor may tend to break down when communication with some robots may be more like communication with animals than like people.

Opportunities for Research: From this review, we identify a number of opportunities to enhance the effectiveness of human-robot teams by adapting research capabilities already developed within the context of virtual human systems. This includes 1) adapting virtual human dialogue authoring and run-time tools for use with robotics applications; 2) using empirical methods for data collection and training of natural language processing components; 3) incorporating advanced dialogue management techniques, and 4) adapting virtual world efforts on object and route descriptions, particularly from the direction-giving challenges.

4 Nonverbal Communication

Face-to-face communication is a highly interactive process where participants mutually exchange and interpret linguistic and gestural signals. Communication dynamics represent the temporal relationship between these signals. Even when only one person speaks at a time, other participants exchange information continuously amongst themselves and with the speaker through gesture, gaze, posture and facial expressions. The transactional view of human communication shows an important dynamic between communicative behaviors where each person serves simultaneously as speaker and listener [32]. At the same time you send a message, you also receive messages from your own communications (individual dynamics) as well as from the reactions of the other person(s) (interpersonal dynamics).

Individual and interpersonal dynamics play a key role when a teacher automatically adjusts his/her explanations based on the student nonverbal behaviors, when a doctor diagnoses a disorder such as autism, or when a negotiator detects deception. An important challenge for artificial intelligence researchers is creating socially intelligent robots and computers, able to recognize, predict and analyze verbal and nonverbal dynamics during face-to-face communication. This will not only open up new avenues for human-computer interactions but create new computational tools for social and behavior researchers—software able to automatically analyze human social and nonverbal behaviors, and extract important interaction patterns.

Nonverbal communicative behavior analysis is a growing field with a large number of applications and especially within the field of virtual human research, where sensing is often simplified through the interaction of a seated person in a well-lit room where all interesting characters and environmental events exist with a fixed computer screen

(e.g., see [33]). Over the past two decades, a first generation of multimodal approaches have been applied in many areas, including audio-visual speech recognition, multimodal object tracking, biometrics, human-computer interaction and multimedia analysis. Also related to this line of research is the research done on audio-visual emotion analysis. Several researchers used prosody (i.e., pitch, speaking rate, etc.) for speech based emotion recognition [34]. Some studies analyzed visual cues, such as facial expressions and body movements [35].

More recently, challenges have been organized focusing on the recognition of emotions using audio and visual cues (e.g., [36]) and drew the participation of many teams from around the world. Note however that all the previous work on audio-visual emotion analysis and multimodal perception was performed on dataset recorded in the laboratory. Also, most of these analyses focus on a generalization of behaviors over a large population, ignoring the idiosyncratic and cultural-specific behaviors of the participants.

Several challenges confront the immediate adoption of this technology to human-robot teams. A robot needs to not only understand the facial expression, body gestures and voice patterns, but it needs to put them in the context of the interactions in the external world, taking into account the multiple human participants, their individuality in expressing personality and emotions and events in the real world. Much of the virtual human research has also focused on dyadic interactions with a very abstract environment (such as a game or simple computer tasks). Rather, human-machine teams demand a focus on more complex interactions involving possibly multiple parties and complex relationships between these entities and environmental events. This will likely require extensions to standardized perception frameworks developed within the multimodal perception community.

Opportunities for Research: From this review, we identify a number of opportunities to enhance the effectiveness of human-robot teams by adapting research capabilities already developed within the context of virtual human systems. This includes (1) learning from readily available data from online website such YouTube, Twitter and Facebook where people are posting a large array of videos with multimodal behaviors and emotions; (2) multimodal deep learning, building on recent achievements in deep neural network modeling to learn the complementarity and synchrony between communicative modalities, and (3) context-based multimodal dialogue, that explicitly models nonverbal behaviors in the shared environment.

5 Trust and Theory of Mind

The increasing capability of autonomous systems has rarely translated into a similar increase in the capability of the human-machine team unit [37]. Studies have identified many causes underlying this phenomenon, but have also shown that simply increasing the capability of the automation in isolation will not suffice [38]. We must instead improve the quality of the interaction between automation and its human operators.

A critical aspect of this interaction is trust [1]. If an autonomous system is better than the human operator at a certain task, then we want the operator to trust the system,

but if the system is worse, we want the operator to distrust it and perform the task manually. Failure to do so results in disuse of automation in the former case and misuse in the latter. Real-world case studies and laboratory experiments show that failures in both cases are common. To achieve proper use of automation, we must better understand why these trust failures occur and what steps we can take to avoid them.

An operator may be willing to trust an autonomous system that has never made a mistake, but it is also important that the operator not overreact to the mistakes that the system will inevitably make. Errors by an autonomous system often have a greater impact on trust than those made by human assistants [39]. Researcher has shown that human operators will more accurately trust an autonomous system if they have a more accurate understanding of its decision-making process and that explaining possible causes of errors can allow an autonomous system to maintain users' trust in the face of such errors [40].

It is thus clear that the transparency of the autonomous system is an important factor in earning appropriate trust. The need for such transparency has motivated researchers in artificial intelligence to develop autonomous agents capable of automatically explaining their decisions [41]. While such transparency certainly increases trust, it also generates a cost to human users in that they must divert attention to communication with an autonomous system. To best manage this cost/benefit tradeoff, the agent literature has framed the problem in terms of the impact of communication on team performance. Teammates communicate so that they can achieve a shared mental model that allows them to perform joint tasks in a coordinated fashion [42]. By weighing the cost of communication against its positive impact in achieving such shared models, agents can optimize their communication strategies to maximize team performance [43].

Transparency through team-oriented communication can help foster trust, but what an autonomous system says may not have as big an impact as what it does. It is thus also important that such systems make good decisions not just in communication, but also in choosing which tasks they do themselves, and which are better left to their teammates. Human-machine teams rely on this adjustable autonomy to flexibly assign different tasks to the most appropriate members, based on capability and situation [44]. Agent researchers have developed algorithms that can optimize the transfers of control that dynamically assign tasks among team members, both human and machine [45]. Combining these existing frameworks for both communication and adjustable autonomy allows researchers to model mixed teams of people, agents, and robots. More recently, we have extended this teamwork model into an agent-based representation of Theory of Mind reasoning [46], allowing agents to model the impact of their decisions on the mental models of their human teammates.

Like human-agent teams, human-robot teams also exhibit a need for trust, shared mental models, and adjustable autonomy. Unfortunately, there remains a sizeable gap between the human-subject studies that quantify human-robot team performance and existing agent-based coordination mechanisms. While the cited agent-based systems all derived better coordination with human users from their communication capability, there has been little quantitative evaluation of the effect (if any) this new capability had on their trust relationship with users. Furthermore, while there has been preliminary work on measuring this effect in virtual simulations of human-robot interaction [47],

none of these agent coordination algorithms have been evaluated in a mixed team combining both human users with physical robots. It thus remains an open question as to the degree that existing human-agent algorithms can benefit human-robot teams.

Opportunities for Research: We see a large opportunity for enhancing the effectiveness of robot-human teams through the use of technology that enhances trust. We have also identified a number of gaps between existing algorithms and HRI needs, as well as algorithmic refinement to close that the gaps found. Such a cycle can support the adaptation of existing agent algorithms to the specific needs of human-robot teams. Specifically, we recommend basic and applied research that addresses (1) automatic explanation algorithms for human-robot trust; (2) domain-independent frameworks for establishing shared mental models in human-robot teams; (3) transfer-of-control strategies for adjustable autonomy for robots to maximize the capabilities of both their human teammates and themselves, and (4) Theory of Mind for robots to adapt to the individual differences across their human teammates.

6 Summary

To conclude, this article identified several points of profitable interaction between research on virtual humans and research on human-robot interaction. These include a focus on core technology shared by both domains – i.e., natural language processing, nonverbal communication – as well as research on how to replicate human interpersonal processes – such as interpersonal trust – within the context of human-machine teams. Finally, we suggest the importance of not blindly assuming that more human-like machines will necessarily yield better teammates, and research is required on which set of interpersonal processes benefit, as opposed to undermine, effective human-machine teams. These recommendations are explained in greater detail in the following technical report: [7].

Acknowledgements. This work is supported by the Army Research Laboratory. Statements and opinions expressed do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred.

References

1. Lee, J.D., See, K.A.: Trust in automation: designing for appropriate reliance. *Hum. Factors J. Hum. Factors Ergon. Soc.* **46**(1), 50–80 (2004)
2. DeSteno, D., Breazeal, C., Frank, R.H., Pizarro, D., Baumann, J., Dickens, L., Lee, J.J.: Detecting the trustworthiness of novel partners in economic exchange. *Psychol. Sci.* **23**(12), 1549–1556 (2012)
3. Breazeal, C., Aryananda, L.: Recognition of affective communicative intent in robot-directed speech. *Auton. Rob.* **12**, 83–104 (2002)
4. Eyssel, F., et al.: If you sound like me, you must be more human: on the interplay of robot and user features on human-robot acceptance and anthropomorphism. In: *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*, ACM (2012)

5. Cassell, J., et al.: Embodiment in conversational interfaces: rea. In: Conference on Human Factors in Computing Systems, Pittsburgh, PA (1999)
6. Gratch, J., et al.: Creating interactive virtual humans: some assembly required In: IEEE Intelligent Systems, pp. 54–61 (2002)
7. Gratch, J., et al.: Exploring the implications of virtual human research for human-robot teams. University of Southern California Institute for Creative Technologies (2014)
8. Gratch, J., Kang, S.-H., Wang, N.: Using social agents to explore theories of rapport and emotional resonance. In: Gratch, J., Marsella, S. (eds.) *Social Emotions in Nature and Artifact*. Oxford University Press, Cambridge (2014)
9. de Melo, C., Carnevale, P.J., Gratch, J.: Humans vs. Computers: The Effect of Perceived Agency on People's Decision Making. University of Southern California (2013)
10. Siegel, M.: Persuasive robotics: how robots change our minds. In: *Media Arts and Sciences*. MIT, Cambridge, MA (2009)
11. Koda, T., Mori, Y.: Effects of an agent's displaying self-adaptors during a serious conversation. In: Bickmore, T., Marsella, S., Sidner, C. (eds.) *IVA 2014*. LNCS, vol. 8637, pp. 240–249. Springer, Heidelberg (2014)
12. Morkes, J., Kernal, H.K., Nass, C.: Effects of humor in task-oriented human-computer interaction and computer-mediated communication: a direct test of SRCT theory. *Hum. Comput. Interact.* **14**(4), 395–435 (1999)
13. Lester, J.C., Towns, S.G., FitzGerald, P.J.: Achieving affective impact: visual emotive communication in lifelike pedagogical agents. *Int. J. Artif. Intell. Educ.* **10**(3–4), 278–291 (1999)
14. Yoshida, W., et al.: Neural mechanisms of belief inference during cooperative games. *J. Neurosci.* **30**(32), 10744–10751 (2010)
15. Chidambaram, V., Chiang, Y.-H., Mutlu, B.: Designing persuasive robots: how robots might persuade people using vocal and nonverbal cues. In: *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM (2012)
16. Weisband, S., Kiesler, S.: Self disclosure on computer forms: meta-analysis and implications. In: *CHI* (1996)
17. Lucas, G.M., et al.: It's only a computer: virtual humans increase willingness to disclose. *Comput. Hum. Behav.* **37**, 94–100 (2014)
18. Sanfey, A.G., et al.: The neural basis of economic decision-making in the ultimatum game. *Science* **300**(5626), 1755–1758 (2003)
19. Lytle, A.L., Brett, J.M., Shapiro, D.L.: The strategic use of interests, rights, and power to resolve disputes. *Negot. J.* **15**(1), 31–51 (1999)
20. Dehghani, M., et al.: Sacred values and conflict over Iran's nuclear program. *Judgment Decis. Making* **5**(7), 540–546 (2010)
21. Damiano, R., et al.: A stroll with Carletto: adaptation in drama-based tours with virtual characters. *User Model. User-Adap. Inter.* **18**(5), 417–453 (2008)
22. Janarthanam, S., Lemon, O.: The GRUVE challenge: generating routes under uncertainty in virtual environments. In: *Proceedings of the 13th European Workshop on Natural Language Generation*. Association for Computational Linguistics (2011)
23. Traum, D.: A computational theory of grounding in natural language conversation. In: *Computer Science Ph.D. thesis*. University of Rochester, Rochester, NY (1994)
24. Traum, D., et al.: Negotiation over tasks in hybrid human-agent teams for simulation-based training. In: *International Conference on Autonomous Agents and Multiagent Systems*. Melbourne, Australia (2003)

25. Sagae, K., et al.: Towards natural language understanding of partial speech recognition results in dialogue systems. In: *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Short Papers*. Association for Computational Linguistics (2009)
26. Morbini, F., et al.: FLoReS: a forward looking, reward seeking, dialogue manager. In: Mariani, J., Rosset, S., Garnier-Rizet, M., Devillers, L. (eds.) *Natural Interaction with Robots, Knowbots and Smartphones*, pp. 313–325. Springer, New York (2014)
27. DeVault, D., Traum, D., Artstein, R.: Making grammar-based generation easier to deploy in dialogue systems. In: *Proceedings of the 9th SIGdial Workshop on Discourse and Dialogue*. Association for Computational Linguistics (2008)
28. Georgila, K., et al.: Practical evaluation of human and synthesized speech for virtual human dialogue systems. In: *LREC* (2012)
29. Hartholt, A., Traum, D., Marsella, S.C., Shapiro, A., Stratou, G., Leuski, A., Morency, L.-P., Gratch, J.: All together now. In: Aylett, R., Krenn, B., Pelachaud, C., Shimodaira, H. (eds.) *IVA 2013*. LNCS, vol. 8108, pp. 368–381. Springer, Heidelberg (2013)
30. Traum, D., Marsella, S.C., Gratch, J., Lee, J., Hartholt, A.: Multi-party, multi-issue, multi-strategy negotiation for multi-modal virtual agents. In: Prendinger, H., Lester, J., Ishizuka, M. (eds.) *IVA 2008*. LNCS (LNAI), vol. 5208, pp. 117–130. Springer, Heidelberg (2008)
31. DeVault, D., Mell, J., Gratch, J.: Toward natural turn-taking in a virtual human negotiation agent. In: *AAAI Spring Symposium on Turn-taking and Coordination in Human-Machine Interaction*. AAAI Press, Stanford, CA (2015)
32. Watzlawick, P., et al.: *Pragmatics of Human Communication. A Study of Interactional Patterns, Pathologies, and Paradoxes*, Paul Watzlawick, Janet Helmick Beavin, Don D. Jackson. WW Norton. (1967)
33. Vinciarelli, A., et al.: Social signals, their function, and automatic analysis: a survey. In: *Proceedings of the 10th International Conference on Multimodal Interfaces*, pp. 61–68. ACM, Chania, Crete, Greece (2008)
34. Tato, R., et al.: Emotional space improves emotion recognition. In: *INTERSPEECH* (2002)
35. Calder, A.J., et al.: Reading the mind from eye gaze. *Neuropsychologia* **40**(8), 1129–1138 (2002)
36. Schuller, B., Valstar, M., Eyben, F., McKeown, G., Cowie, R., Pantic, M.: AVEC 2011—the first international audio/visual emotion challenge. In: D’Mello, S., Graesser, A., Schuller, B., Martin, J.-C. (eds.) *ACII 2011, Part II*. LNCS, vol. 6975, pp. 415–424. Springer, Heidelberg (2011)
37. Parasuraman, R., Riley, V.: Humans and automation: use, misuse, disuse, abuse. *Hum. Factors J. Hum. Factors Ergon. Soc.* **39**(2), 230–253 (1997)
38. Sorkin, R.D., Woods, D.D.: Systems with human monitors: a signal detection analysis. *Hum. Comput. Interact.* **1**(1), 49–75 (1985)
39. Dzindolet, M.T., et al.: The perceived utility of human and automated aids in a visual detection task. *Hum. Factors J. Hum. Factors Ergon. Soc.* **44**(1), 79–94 (2002)
40. Dzindolet, M.T., et al.: The role of trust in automation reliance. *Int. J. Hum Comput Stud.* **58**(6), 697–718 (2003)
41. Myers, K., et al.: An intelligent personal assistant for task and time management. *AI Mag.* **28**(2), 47 (2007)
42. Grosz, B., Kraus, S.: Collaborative plans for complex group action. *Artif. Intell.* **86**(2), 269–357 (1996)

43. Pynadath, D.V., Tambe, M.: Multiagent teamwork: analyzing the optimality and complexity of key theories and models. In: *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems: Part 2*. ACM (2002)
44. Dorais, G., et al.: Adjustable autonomy for human-centered autonomous systems. In: *Working notes of the Sixteenth International Joint Conference on Artificial Intelligence Workshop on Adjustable Autonomy Systems* (1999)
45. Tambe, M., Scerri, P., Pynadath, D.V.: Adjustable autonomy for the real world. *J. Artif. Intell. Res.* **17**(1), 171–228 (2002)
46. Pynadath, D.V., Marsella, S.C.: PsychSim: modeling theory of mind with decision-theoretic agents. In: *IJCAI* (2005)