Embodied social robots trigger gaze following in real-time

Eva Wiese, Patrick P. Weis, and Daniel M. Lofaro

George Mason University, USA, VA

ABSTRACT

In human-human interaction, we use information from gestures, facial expressions and gaze direction to make inferences about what interaction partners think, feel or intend to do next. Observing changes in gaze direction triggers shifts of attention to gazed-at locations and helps establish *shared attention* between gazer and observer - a pre-requisite for more complex social skills like mentalizing, action understanding and joint action. The ability to follow others' gaze develops early in life and being able to process gaze signals is a crucial milestone in human development. While human gaze signals are so essential for social interactions that we automatically follow them, it is unclear whether robot gaze cues are followed to similar degrees, and whether they have the ability to establish shared attention between human and robot. Furthermore, most studies on social attention in human-robot-interaction (HRI) use robot images and videos in controlled laboratory settings, which makes it necessary to determine whether gaze following can also be observed in social interactions with embodied robot platforms in real-time. In the current experiment, we use the humanoid robot Meka to examine whether gaze following can be induced in realistic interactions with social robots. The results indicate that Meka's gaze cues were reliably followed, and that they were able to establish shared attention in HRI. Implications of this finding for social robotics are discussed.

Author Keywords

Social Robotics; Intentions in HRI; Gaze Following.

ACM Classification Keywords

H.1.2., H.5.2., I.2.9., I.2.11.

INTRODUCTION

Bodily signals like gaze direction or gestures are important cues during social interactions that inform us about the social relevance of observed behavior and influence how we react towards others. A crucial part of this process concerns the degree to which interaction partners are believed to be

intentional agents with internal states like beliefs, emotions and action goals [12]. Perceiving intentionality in behaviors of others is a prerequisite for developing a Theory-of-Mind and allows us to make inferences about their internal states [2]. For instance, perceiving a fearful facial expression triggers the inference that someone is anxious or stressed, and observing someone look at an apple or banana triggers the thought that he/she must be hungry and intends to make an arm movement in order to grasp the food and eat it.

While reasoning about the internal states of others happens automatically in human-human interactions, research suggests that robots trigger processes of social inference only when believed to be intentional beings with the abilities to plan, act and feel [25,30,32]. To be perceived as intentional agent, it was shown that robots need to appear as similar to humans, which is accomplished by giving them human-like physical features and/or behaviors [1,29]. Once robots are perceived as intentional beings, social relevance is ascribed to their actions, which positively affects one's attitudes and performance in HRI [1,25,30,32]. For example, robots that exhibit human gestures like nodding or shrugging, trigger more positive emotional reactions in human observers and are trusted more than robots that show mechanistic gestures [5]. Other studies show that robots being perceived as intentional induce social facilitation effects in human observers [26,31], and positively affect joint action performance [4,6]. One agent feature that particularly seems to trigger attribution of intentionality is emulating human gaze behaviors in artificial agents such as social robots [23], or virtual avatars [7]. For instance, robots that move their gaze during social interactions as opposed to robots whose eyes do not move are perceived as more enjoyable [15], and robots that conjointly attend to where human interaction partners are looking are rated as more competent than robots that fail to establish shared attention [13]. Likewise, robots that use gaze signals to communicate improve recollection performance during memory tasks [22], and simplify the communication between human and robot partners by enabling turn-taking processes [23,33]. Furthermore, robots reacting to human input by shifting gaze in a coherent fashion positively affect the reported physical and emotional closeness between human and robot agents [21], and their gaze behavior is more likely to be judged as intentional rather than random [35].

Despite increasing the perceived human-likeness of a robotic agent, implementing social gaze behaviors in robots also has a positive effect on social attention processes: observing

changes in others' gaze direction has been shown to trigger shifts of the observers attention to the gazed-at location, and allows two social partners to conjointly attend to the same location, object or event in visual space (*shared attention*, [10]). Shared attention is necessary to coordinate actions of multiple interaction partners in time and space, is a prerequisite for developing a functioning Theory-of-Mind, and is necessary for making inferences about the internal states of others [11]. In laboratory experiments, shared attention processes are examined by presenting face stimuli on a screen that first gaze straight ahead (i.e., *mutual gaze*), and then change gaze direction to the left or right side of the screen (i.e., *averted gaze*), which triggers shifts of the participant's attention to the gazed-at location [10]. The change in gaze direction, or *gaze cue*, is then followed by the presentation of a target item (dot or letter), which either appears at the cued location (valid trial) or at an uncued location (invalid trial). Participants are instructed to respond as fast and accurately as possible to the location or identity of the target by pressing keys on a standard keyboard. Since observing changes in gaze direction shifts attention to gazed-at locations, reaction times to targets presented at the cued location are usually faster than reaction times to targets presented at uncued locations (i.e., *gaze-cueing effect*, [10]).

For the most part, gaze cues from robot agents induce similar effects on attentional orienting in observers as gaze cues from human agents [8,28]. For instance, robots that establish mutual gaze with human partners receive more favorable evaluations [34], and participants spend more time on interactions with them [14] compared to robots that do not establish mutual gaze. Robot gaze can also be interpreted with the same spatial accuracy as human gaze, although this effect seems to depend on the morphology of the robot eye (extent to which robot eye is covered by eyelids; [30]). Additionally, robot gaze cues seem to be so natural to human observers that adults readily follow the gaze of robot images in highly controlled computer experiments [1,30,32], and even ten-month old infants follow the line of sight of robots to gazed-at locations [20]. Despite these similarities, several studies have shown that the cognitive mechanisms underlying human versus robot gaze processing seem to differ quite substantially. Yu and colleagues, for instance, looked at the dynamics of attentional orienting in response to gaze cues, and found that participants spent more time looking at robot faces than human faces in a gaze-based object naming task, indicating that participants were concerned if the robot was able to attend to relevant objects [36]. Admoni and Scassellati further showed that while humans automatically follow the gaze of human agents, attending to where robot agents are looking can more easily be suppressed [1], which suggests that human gaze might be processed in different neural pathways than robot gaze. Similarly, eye tracking in 12-month-old infants revealed anticipatory gaze shifts and enhanced processing of looked-at objects in response to human but not robot gaze cues [24].

While there is evidence that robots can induce gaze following effects in controlled laboratory settings, it has not been systematically examined whether embodied, physical robots can induce gaze following effects in real-time human-robot interaction. This, however, is an essential question since the realism of the interaction and the embodiment of the robot could potentially have positive and negative effects. On the one hand, one could hypothesize that gaze following might be stronger in real-time interactions with robots since beneficial effects of physical embodiment over virtual environments have been vastly reported in human-robot interaction [3,16]. Physical embodiment should also increase the social relevance of robot gaze with the consequence that gaze signals are followed more strongly than in laboratory settings. On the other hand, one could speculate that real-time interactions with robots might also have negative effects on gaze following due to the fact that eye movements are mechanistic (i.e., lack biological motion) and that their temporal dynamics are unnatural (i.e., delayed gaze shifts in response to objects in the environment, [3,16]). To start addressing these questions, we examine in the current experiment whether gaze following can be induced in real-time interactions with embodied social robots. Future studies then need to systematically compare gaze following in laboratory settings versus realistic environments.

METHODS AND MATERIALS

Participants
Seventeen participants were recruited for the experiment using the participant management website SONA Systems. Two participants were excluded due to technical problems with the setup and one because of prolonged reaction times leading to a final sample size of fourteen participants (six females; mean age: 24.6; one left handed). All participants reported normal or to corrected-to-normal vision and gave informed consent. Data was stored and analyzed anonymously. Testing time was about twenty minutes.

Apparatus and Stimuli
The test setup consists of the S2 humanoid robot head Meka, two smart bulbs, a custom-made user touch interface, and a sub-millisecond accuracy timer.

Meka Head and Meka-Ach-Middleware
The Meka head has 8 degrees of freedom: four degrees of freedom in the neck in a serial formation (pitch, yaw, roll, pitch), and two degrees of freedom (pitch, yaw) for each eye (left, right). The eyes are attached to the face which is at the end of the neck's kinematic chain. The middleware used to control Meka is an extension of the Hubo-Ach system used for the Hubo (KHR-4), Hubo 2+, DRC-Hubo, and DRC-Hubo+ models [9,19]. Accordingly the middleware is called Meka-Ach, which is a multi-process based real-time control system that uses high-speed and low-latency shared memory channels with networking capabilities.

Straight Gaze | **Gaze Cue** | **Target** | **Inter-Trial Interval**

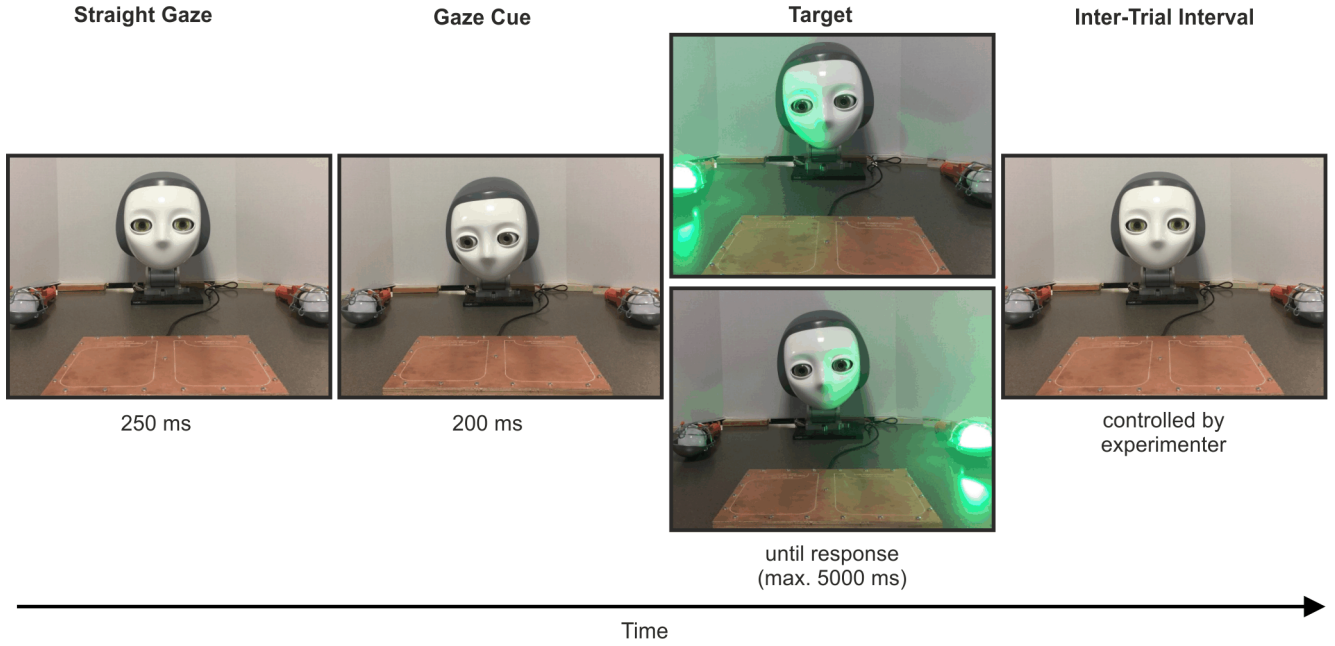250 ms | 200 ms | until response (max. 5000 ms) | controlled by experimenter

Time

Figure 1. Trial Sequence: Meka first established direct eye contact with the participant and then shifted her gaze to either the light on the left or the right side of the table. Afterwards, either the light that was looked at by Meka or the other light flashed. The task was to respond to the flash as quickly and as accurately as possible. At the end of the trial, Meka would return back to neutral position and establish eye contact once more to start a new trial.

Target Stimuli

The experiment requires lights that can be turned on and off with millisecond precision. For that purpose, two compatible smart lights were used. The smart lights are connected on a dedicated 2.4 *ghz* wireless network. An on/off message is created using a *Node.js* implementation and sent over TCP/IP. To allow for consistent communications latency each light is placed 1.5 *m* from the antenna in identical orientations. The resulting latency (denoted by t_l) was recorded to be 1.032 *ms* with a 0.217 *ms* standard deviation. This latency is taken into account when recording the reaction time of the participant.

Head Motion

As a gaze cue, Meka looks at either the left or right bulb based on pre-programmed joint space values before one of the lights turns on. These values are pre-determined by the use of the head's sparse reachable map and Inverse Jacobian Inverse Kinematics method [18]. The input to the inverse kinematic solver is the location of the object desired to be looked at by the robot. The location of the object is in reference to the robot head. The resulting joint space values are recorded and used for each given motion. Linear interpolation between joint space values is used when a work-space step input is given. The time of the linear interpolation is 0.25 *sec* and is updated at a rate of 500 *hz*.

Reaction Time Recording

In order to capture participants' responses, a self-calibrating capacitive touch interface with two large buttons was used (see *Figure 1*). System clock time is recorded twice every trial, first when the smart light turns on (t_{c0}) and second when a message from the capacitive touch interface is re-

ceived (t_{c1}). The test subject reaction time t_c additionally takes into account the light activation latency t_l and is calculated by: $t_c = (t_{c1} - t_{c0}) - t_l$.

Design and Procedure

At the beginning of the experiment, participants received written instructions, and gave informed consent. They were instructed to perform a gaze following task that required them to respond as fast and accurately as possible to the change in color of one of two lights in front of them. Responses to changes in color had to be given by pressing a button on their left for the light on the left and a button on their right for the light on the right (i.e., localization task). Participants were also told that before the color change occurred, Meka would perform an eye movement from initially looking at them to looking at one of the two lights. Participants were instructed to respond as soon as they noticed the change in color, and the time it took participants to react to the change in color was measured as dependent variable. As soon as participants had given their response by key press, Meka would move back to her original position and participants could initiate a new trial by establishing direct eye contact with Meka.

Figure 1 illustrates the sequence of events on a given trial. At the beginning of the trial, Meka established mutual gaze with the participant to signal her readiness to start the trial. 250 ms afterwards, Meka would change her gaze direction to either look at the left or right light on the table in front of her (including a slight head movement). After a stimulus-onset asynchrony (SOA) of 200 ms, either the light positioned at the gazed-cued location or the light positioned at the uncued location changed its color, and participants re-

sponded to the location of the lit-up light. Meka's head and gaze posture, as well as the color of the light remained unchanged until a response was given or a time-out of 5000 ms was reached, whichever appeared first. At the end of the trial, Meka moved back to her neutral position, and participants prepared for the beginning of the next trial.

The experiment was composed of 80 trials: 40 valid trials, and 40 invalid trials. Gaze direction (left, right), and target side (left, right) were selected pseudo-randomly and every combination appeared with equal frequency. Gaze validity was calculated based on the combination of gaze direction and target location: on valid trials, targets appeared where Meka was looking (i.e., gaze to the left, target on the left), while on invalid trials, targets appeared opposite of where Meka was looking (i.e., gaze to the left, target on the right). No information about the reliability of Meka's gaze behavior was disclosed to the participants at any time during the experiment.

Analysis
Reaction time data was analyzed using R 3.2.4. Misses and incorrect responses, as well as reaction times deviating by more than 2.5 SD from the individual participants means were removed prior to analyses, totaling 2.67% of all trials. The remaining data was analyzed in two steps: First, average reaction times for valid and invalid trials were calculated for each participant. Second, average reaction times for valid and invalid trials were compared using a t-test, with a significant difference in reaction times between valid and invalid trials being evidence for the presence of a gaze following effect.

RESULTS
Results of the analysis of the reaction time data are shown in *Figure* 2. The t-test revealed a significant difference between valid and invalid trials (t (13) = 4.00, p = .002, η^2_p = .533, d = 1.069), with shorter reaction times for valid (M = 500 *ms*) than invalid trials (M = 525 *ms*), which provides evidence for the presence of a gaze following effect.

DISCUSSION
The goal of the experiment was to examine whether gaze following in human-robot interaction is limited to laboratory settings or whether it can also be observed in real-time interactions with embodied robots. To address this issue, we adapted a computer-based gaze following protocol [10] to real-time interactions with the embodied humanoid robot head Meka. Participants had to perform a localization task, where they had to indicate by key press as fast and accurately as possible whether a light on the left or the right side of the table changed its color. Crucially, the light changing its color was either cued or uncued by Meka's gaze.

Reaction time to the target was measured as dependent variable and differences in reaction times between valid and invalid trials were calculated to determine whether significant gaze following effects were evident. Reaction times on valid trials were significantly shorter than reaction times on

invalid trials, showing that the observation of gaze following effects is not specific to controlled laboratory settings but generalizes to real-time interactions with physically embodied robots. With 25 ms, the observed effect is slightly larger than gaze following effects observed in laboratory experiments (i.e., 15 *ms*, [25,30,32]). This enhanced gaze following effect could be due Meka's physical embodiment or the fact that in addition to eye movements, Meka also performed head movements, potentially providing a second directional cue. The results are in line with previous studies showing that gaze following is not limited to human-human interaction, but can also be observed in real-time human-robot interaction [8,17,28].

Our results are also in line with previous studies that implemented non-verbal social cues to improve the quality of human-robot-interaction [4,27]. For example, it was found that participants were able to respond to a robot in a conversation-like setting more quickly when the robot indicated the participant's turn with flashing in its eyes [27]. On a descriptive level, it was also found that cues about internal states of a robot (i.e., eye gaze conveying focus of attention or shrugging gestures conveying confusion), increased efficiency on tasks that required participants to teach the names of different buttons to a robot and to subsequently ask the robot to press all of those buttons [4]. While those studies were either using non-human-like social cues (eye flashes) or a mixture of different human-like cues (eye gaze, shrugging, and more) to show the benefit of social cues in general, our study investigated whether natural eye gaze alone is sufficient to establish joint attention in social interactions between human and robot agents.
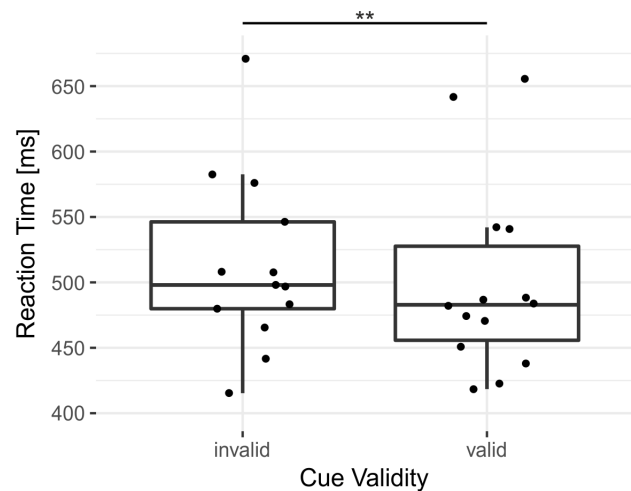
Figure 2. Reaction times in invalid and valid conditions: Horizontal boxplot lines represent 25th, 50th, and 75th percentiles. Whiskers mark data points that fall into the 1.5 interquartile range. Individual data points are plotted on top of the boxplot. ** p < .01.

There are two limitations to the current study that should be addressed in future research on real-time gaze following. First, the experiment cannot quantify the degree to which robot gaze triggered shifts of attention in human observers since gaze and head cues were presented at the same time. Because of that, it cannot be determined if gaze following is stronger in real-time interactions with embodied robots than in laboratory settings. Future studies need to address this limitation by manipulating head and gaze cues separately and comparing the effect they induce to gaze-cueing effects observed in laboratory experiments. Second, due to practical reasons, data was collected in a relatively noisy environment with a high amount of people passing by, which could potentially increase the noise in our data. Gaze cueing has mostly been investigated in well controlled laboratory settings [1,25,30,32] and while our current setting should have increased external validity compared to those studies, the different environments make it hard to compare effect sizes between studies.

CONCLUSION

The results show that humans ascribe social relevance to robot gaze and readily establish joint attention with them in interactive scenarios. In consequence, social robots should be designed to use their gaze when working with humans on joint tasks to resolve ambiguity and shift their interaction partner's attention to relevant locations and objects. Furthermore, since gaze direction is automatically and effortlessly used by humans to draw inferences about the internal states of others, robot gaze can be used to implicitly guide our expectations when interacting with the robot. To sum up, implementing gaze in robots allows us to predict robot behavior and make human-robot interaction more efficient and productive.

REFERENCES

1. Henny Admoni and Brian Scassellati. 2012. Robot Gaze Is Different From Human Gaze: Evidence that robot gaze does not cue reflexive attention. In *Proceedings of the "Gaze in Human-Robot Interaction" Workshop*.

2. Ralph Adolphs. 2003. Cognitive neuroscience: Cognitive neuroscience of human social behaviour. *Nature Reviews Neuroscience* 4, 3: 165–178. https://doi.org/10.1038/nrn1056

3. Wilma A. Bainbridge, Justin W. Hart, Elizabeth S. Kim, and Brian Scassellati. 2011. The benefits of interactions with physically present robots over video-displayed agents. *International Journal of Social Robotics* 3, 1: 41–52.

4. C. Breazeal, C.D. Kidd, A.L. Thomaz, G. Hoffman, and M. Berlin. 2005. Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 708–713. https://doi.org/10.1109/IROS.2005.1545011

5. Elizabeth J. Carter, Michael N. Mistry, G. Peter K. Carr, Brooke A. Kelly, and Jessica K. Hodgins. 2014. Playing catch with robots: Incorporating social gestures into physical interactions. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, 231–236. https://doi.org/10.1109/ROMAN.2014.6926258

6. Nathan Caruana, Genevieve McArthur, Alexandra Woolgar, and Jon Brock. 2016. Simulating social interactions for the experimental investigation of joint attention. *Neuroscience & Biobehavioral Reviews*, 74: 115–125. https://doi.org/10.1016/j.neubiorev.2016.12.022

7. Justine Cassell. 2000. Embodied conversational interface agents. *Communications of the ACM* 43, 4: 70–78.

8. Thierry Chaminade and Maria M. Okka. 2013. Comparing the effect of humanoid and human face for the spatial orientation of attention. *Frontiers in Neurorobotics* 7, 12. https://doi.org/10.3389/fnbot.2013.00012

9. Neil T. Dantam, Daniel M. Lofaro, Ayonga Hereid, Paul Y. Oh, Aaron D. Ames, and Mike Stilman. 2015. The ach library: a new framework for real-time communication. *IEEE Robotics & Automation Magazine* 22, 1: 76–85.

10. Chris Kelland Friesen and Alan Kingstone. 1998. The eyes have it! Reflexive orienting is triggered by nonpredictive gaze. *Psychonomic bulletin & review* 5, 3: 490–495.

11. Chris D. Frith and Uta Frith. 2006. How we predict what other people are going to do. *Brain Research* 1079, 1: 36–46. https://doi.org/10.1016/j.brainres.2005.12.126

12. Heather M. Gray, Kurt Gray, and Daniel M. Wegner. 2007. Dimensions of mind perception. *Science* 315, 5812: 619–619.

13. Chien-Ming Huang and Andrea L. Thomaz. 2011. Effects of responding to, initiating and ensuring joint attention in human-robot interaction. In *20th IEEE international symposium on robot and human interactive communication*, 65–71. https://doi.org/10.1109/ROMAN.2011.6005230

14. A. Ito, S. Hayakawa, and T. Terada. 2004. Why robots need body for mind communication - an attempt of eye-contact between human and robot. 473–478. https://doi.org/10.1109/ROMAN.2004.1374806

15. Takayuki Kanda, Hiroshi Ishiguro, and Toru Ishida. 2001. Psychological analysis on human-robot interaction. In *Robotics and Automation, 2001. Proceedings 2001 ICRA. IEEE International Conference on*, 4166–4173.

16. James Kennedy, Paul Baxter, and Tony Belpaeme. 2015. Comparing robot embodiments in a guided discovery learning interaction with children. *International Journal of Social Robotics* 7, 2: 293–308.

17. Hideki Kozima and Hiroyuki Yano. 2001. A robot that learns to communicate with human caregivers. In *Proceedings of the First International Workshop on Epigenetic Robotics*, 47–52.

18. Daniel M. Lofaro, Robert Ellenberg, Paul Oh, and Jun-Ho Oh. 2012. Humanoid throwing: Design of collision-free trajectories with sparse reachable maps. In *International Conference on Intelligent Robots and Systems*, 1519–1524.

19. Daniel Marc Lofaro. 2013. *Unified algorithmic framework for high degree of freedom complex systems and humanoid robots*. Drexel University.

20. Javier R. Movellan and John S. Watson. 2002. The development of gaze following as a Bayesian systems identification problem. In *Development and Learning, 2002. Proceedings. The 2nd International Conference on*, 34–40.

21. Jonathan Mumm and Bilge Mutlu. 2011. Human-robot proxemics: physical and psychological distancing in human-robot interaction. In *Proceedings of the 6th international conference on Human-robot interaction*, 331–338.

22. Bilge Mutlu, Jodi Forlizzi, and Jessica Hodgins. 2006. A Storytelling Robot: Modeling and Evaluation of Human-like Gaze Behavior. In *6th IEEE-RAS International Conference on Humanoid Robots*, 518–523. https://doi.org/10.1109/ICHR.2006.321322

23. Bilge Mutlu, Takayuki Kanda, Jodi Forlizzi, Jessica Hodgins, and Hiroshi Ishiguro. 2012. Conversational gaze mechanisms for humanlike robots. *ACM Transactions on Interactive Intelligent Systems* 1, 2: 1–33. https://doi.org/10.1145/2070719.2070725

24. Yuko Okumura, Yasuhiro Kanakogi, Takayuki Kanda, Hiroshi Ishiguro, and Shoji Itakura. 2013. The power of human gaze on infant learning. *Cognition* 128, 2: 127–133.

25. Ceylan Özdem, Eva Wiese, Agnieszka Wykowska, Hermann Müller, Marcel Brass, and Frank Van Overwalle. 2016. Believing androids–fMRI activation in the right temporo-parietal junction is modulated by ascribing intentions to non-human agents. *Social Neuroscience*: 1–12. https://doi.org/10.1080/17470919.2016.1207702

26. Nina Riether, Frank Hegel, Britta Wrede, and Gernot Horstmann. 2012. Social facilitation with social robots? In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*, 41–48.

27. Jef A. van Schendel and Raymond H. Cuijpers. 2015. Turn-yielding cues in robot-human conversation. *New Frontiers in Human-Robot Interaction*: 85.

28. Maria Staudte and Matthew W. Crocker. 2009. Visual attention in spoken human-robot interaction. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, 77–84.

29. Adam Waytz, Kurt Gray, Nicholas Epley, and Daniel M. Wegner. 2010. Causes and consequences of mind perception. *Trends in cognitive sciences* 14, 8: 383–388.

30. Eva Wiese, Agnieszka Wykowska, Jan Zwickel, and Hermann J. Müller. 2012. I See What You Mean: How Attentional Selection Is Shaped by Ascribing Intentions to Others. *PLoS ONE* 7, 9: e45391. https://doi.org/10.1371/journal.pone.0045391

31. Sarah Woods, Kerstin Dautenhahn, and Christina Kaouri. 2005. Is someone watching me?-consideration of social facilitation effects in human-robot interaction experiments. In *Computational Intelligence in Robotics and Automation, 2005. CIRA 2005. Proceedings. 2005 IEEE International Symposium on*, 53–60.

32. Agnieszka Wykowska, Eva Wiese, Aaron Prosser, and Hermann J. Müller. 2014. Beliefs about the minds of others influence how we process sensory information. *PLoS One* 9, 4: e94339.

33. Akiko Yamazaki, Keiichi Yamazaki, Matthew Burdelski, Yoshinori Kuno, and Mihoko Fukushima. 2010. Coordination of verbal and non-verbal actions in human–robot interaction at museums and exhibitions. *Journal of Pragmatics* 42, 9: 2398–2414. https://doi.org/10.1016/j.pragma.2009.12.023

34. Tomoko Yonezawa, Hirotake Yamazoe, Akira Utsumi, and Shinji Abe. 2007. Gaze-communicative behavior of stuffed-toy robot with joint attention and eye contact based on ambient gaze-tracking. In *Proceedings of the 9th international conference on Multimodal interfaces*, 140–145.

35. Yuichiro Yoshikawa, Kazuhiko Shinozawa, Hiroshi Ishiguro, Norihiro Hagita, and Takanori Miyamoto. 2006. The effects of responsive eye movement and blinking behavior in a communication robot. In *International Conference on Intelligent Robots and Systems*, 4564–4569.

36. Chen Yu, Paul Schermerhorn, and Matthias Scheutz. 2012. Adaptive eye gaze patterns in interactions with human and artificial agents. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 1, 2: 13.