Author note: 28 November, 2018 – This preprint is the accepted version of our manuscript.

In natural interaction with embodied robots we prefer it when they follow our gaze: a gaze-
contingent mobile eyetracking study

Cesco Willemse* and Agnieszka Wykowska

Istituto Italiano di Tecnologia, Genoa, Italy

* Corresponding author:

Cesco Willemse

Istituto Italiano di Tecnologia

Social Cognition in Human-Robot Interaction

Via Enrico Melen 83

16040 Genova

Italy

email: cesco.willemse@iit.it

Abstract

Initiating joint attention by leading someone's gaze is a rewarding experience which facilitates social interaction. Here, we investigate this experience of leading an agent's gaze whilst applying a more realistic paradigm than traditional screen-based experiments. We used an embodied robot as our main "stimulus" and recorded participants' eye-movements. Participants sat opposite a robot which had one of two 'identities' – "Jimmy" or "Dylan". Participants were asked to look at one of two objects presented on screens to the left and the right of the robot. Jimmy then looked at the same object in 80% of the trials and at the other object in the remaining 20%. For Dylan, this proportion was reversed. Upon fixating on the object of choice, participants were asked to look back at the robot's face. We found that return-to-face saccades were conducted earlier toward Jimmy *when* he followed the gaze compared with when he did not. For Dylan, there was no such effect. Additional measures indicated that our participants also preferred Jimmy and liked him better. This study demonstrates a) the potential of technological advances to examine joint attention where ecological validity meets experimental control, and b) that social reorienting is enhanced when we initiate joint attention.

Keywords: joint attention, gaze leading, mobile eyetracking, social robots, gaze contingency.

In natural interaction with embodied robots we prefer it when they follow our gaze: a gaze-contingent mobile eyetracking study

Humans are efficient in processing social information from others' gaze (1,2). The gaze-cueing effect, in which spatial orienting is facilitated by receiving the redirection of an agent's gaze (3,4) is one example. On the other hand, shifts of gaze-direction can also be *signalled* to an agent in a bid to establish joint attention. In typical gaze-leading paradigms, participants redirect their gaze toward a stimulus in the peripheral view, after which an avatar on the screen looks at either the same object or at a different one (5,6). Successful initiation of joint attention in this manner has been found to increase brain activation related to hedonistic reward (6) and to increase subjective experiences of liking and preference of the responding avatar (5,7,8).

Moreover, people frequently take less time to re-orient their gaze back from the object toward the avatar when this agent followed the gaze, compared to when it did not (5,7). This effect is also modulated by the typically experienced response of the avatar: the average onset latency of these return-to-face saccades is shorter toward avatars who usually follow the participant's gaze toward an object relative to those who typically looked at the other object (8).

Whereas past gaze-leading studies had their stimuli presented on a computer screen with the use of a desktop-mounted eyetracker, our current paper takes advantage of two technical advances to mimic realistic scenarios more closely. Firstly, we used a humanoid robot as an embodied agent, introducing social presence, to examine ecologically valid aspects of social cognition (9,10). Secondly, we used mobile eye-tracking technology to allow for a set up that would be more ecologically valid, and that would not rely on using a (single) screen, as that may

not be fully representative of natural gaze allocation and attention in environments outside of the laboratory (11).

Another clear advantage of using embodied agents as stimuli besides studying whether screen-based findings are analogous to human social cognition in more realistic scenarios, is the potential to examine whether human-robot interaction is analogous to human-human interaction. That is to say, examining this possible analogy here will give us a better understanding under which conditions robots evoke mechanisms of social cognition in the humans, whether robots are perceived as animate agents, or when they are anthropomorphised.

Goals of the current study

Technological advances in eyetracking form a welcome invitation to study joint attention, established by gaze-leading, with a humanoid robot. This poses a number of questions and challenges such as whether speeded return-to-face saccades for joint attention episodes and the associated increase in preference and likeability replicate in setups that go away from the screen. Another such question is whether robots are perceived as more human-like and in possession of mental states when they respond congruently to gaze-leading. Finally, a challenge lies in establishing whether a gaze-contingent experimental paradigm can be feasibly conducted with mobile eyetrackers. Specifically, the study presented here aims to shed light on the following questions:

1. *Do past screen-based findings replicate in a more naturalistic scenario with an embodied agent?* This speeded return-to-face saccade has been explained as a sensitive marker of attentional engagement with the agent. We previously also found that besides the gaze being followed on an ad-hoc basis, people are also sensitive to whether the agent *usually* follows their gaze or not (8). However, on the side-lines of the current reproducibility debate in psychology

(e.g. (12, 13, 14) in which one-one-one replication is encouraged, even less is currently known about whether previous findings in social cognition replicate in more naturalistic adaptations. One might argue that they would, as traditional experiments are often well-controlled and contain few confounding variables. On the other hand, they may not necessarily replicate. For example it has been found that people look more at a person on a screen compared to the same person when she is present and whom could thus be actually interacted with (15). This begs an invitation to exploring past findings in gaze leading with added social presence. Therefore, we designed an experiment with a more naturalistic social interaction scenario, where participants were seated in front of an *embodied* robot of human-like size.

The two robot ‘identities’ were manipulated as within-subjects factor: one identity who usually followed the gaze and one who usually did not. We expected our participants to be sensitive to these identities as well as the ad-hoc contingency, in line with previous findings. Note, however, that, as mentioned above, replication of previous screen-based findings would not only be a “mere” replication, but it would show that findings obtained in controlled (but artificial) setups generalise to more ecologically valid scenarios, with embodied agents that can manipulate the environment (in contrast to stimuli only presented on the screen). Moreover, it would serve as a proof of concept that traditional paradigms of cognitive experimental psychology can be successfully implemented in a human-robot-interaction scenario, with replicable results and adequate scientific rigor – a task that is not trivial, given the technological challenges of integrating various components of the setup (the humanoid robot, an eye tracker, stimulus presentation and response collection software) with excellent temporal synchronisation that is required.

2. *Does establishing joint-attention with gaze-leading influence subjectively-reported likability or preference toward other agents?* It has been reported that successful gaze-leading affects subjective attributions toward an agent, such as likability and preference, relative to non-following agents (5,7,8,12). It could also be argued that people will anthropomorphise a robot more after a higher degree of joint-attention (8). Therefore, we subjected participants to questionnaire items assessing anthropomorphising, likeability, and preference. We expected to find a positive relationship between the amount of successfully initiating joint attention and subjective ratings of likeability, human-likeness, and preference between the two agents.

Additionally, we explored whether establishing joint attention with the robot increased the adoption of the intentional stance. Dennett (17) proposed the idea that humans adopt various stances or strategies when predicting or explaining observed behaviours of other systems. Various strategies are best suited to various systems. For explaining/predicting behaviour of other humans the intentional strategy, which refers to mental states, works best. It might therefore be that humans also adopt the intentional stance towards humanoid robots too, and that a certain type of social interactions with them might influence the likelihood of adopting the intentional stance. In order to examine this, we presented a theory of mind test and an intentional stance questionnaire recently developed (18) directly after the participants completed the experimental session with each robot identity.

Method

Participants

We set a sample size of $N = 32$ before commencing this study, based on our previous study (8). We estimated to suffer approximately 10-20% data-loss, even though this was difficult to predict, and therefore deliberately overshot recruitment, maximising lab availability. In total,

37 participants (21 females, $M_{age} = 24.4$ years, $S.D. = 3.89$) took part in the study for a payment of €15. This study was conducted in accordance with the ethical approval from the local ethical committee (Comitato Etico Regione Liguria) and participants provided written informed consent to participate.

Materials

Participants sat opposite the iCub robot (19) at a distance of approximately 125 cm in a lit room, with a table in-between the robot and the participant. The iCub was mounted at a height such that his eyes were 122 cm from the floor, which we estimated to be roughly at eye-level with most participants. Object images (720 pixels height, variable width, average 513 pixels, except one object which was 720 pixels wide and 318 pixels high) were presented on two screens (27", 2560 x 1440 pixels resolution, 144 Hz refresh rate) placed on the table. These screens were positioned on both sides of the table, so that the iCub and the participant were in-between the screens (see Figure 1 for the set-up from the participant's perspective). The screens were tilted back by 15° from the vertical position, rotated 75° laterally, and positioned 42 cm apart measured between the closest corners.

We used a pair of Tobii Pro Glasses 2 to record gaze data at a sampling rate of 100 Hz. Additionally, these data were streamed live to iCub via a Python controller (https://github.com/ddetommaso/TobiiProGlasses2_PyCtrl). Specifically, we divided the front-mounted camera image into three zones (left, centre, and right; 30%, 40%, and 30% respectively). This information was sent to the presentation software (OpenSesame v3.2; (20)) to control where the robot looked.

Robot-related questionnaires and individual differences. A series of questionnaires was presented throughout the experiment. Firstly, we used the "InStance" questionnaire devised

to assess whether people adopt intentional stance towards the iCub robot (18). This questionnaire presents 34 three-image sequences of iCub interacting with objects and/or people, and participants are asked to give an explanation for the displayed behaviour by sliding a slider on a rating scale either to a provided answer that has a mentalistic description, or to one that is mechanistic.

Figure 1. Trial sequence. Starting top-left: **A-** The participant looked at the robot until they heard a beep. **B-** They looked to the left or right object as quickly as possible. **C-** iCub looked at an object (gaze-following example provided). **D-** In their own time, the participant looked back at the robot's face (return-to-face saccade onset-time) upon which the robot looked at the participant again.

Additionally, we adapted the Yoni Test of cognitive and affective Theory of Mind (21) and replaced the central original character with an iCub drawing to measure theory of mind toward the iCub. The Yoni Test is comprised of 98 trials with first order and second order scenarios which assesses both cognitive and affective theory of mind.

Further, we used the likeability and anthropomorphism sub-scales of the Godspeed questionnaire (22). These subscales are comprised of five items each with 1-5 Likert scale scoring between two antonymous adjectives (e.g. unpleasant – pleasant), so both subscales had a possible range of 5-25 with higher scores reflecting the attribution of more positive traits.

Finally, we assessed robot-preference with additional questions (which robot they preferred and why, which one they would think was meant to be the mentalistic/mechanistic one, and which one they would think was faster than the other). Other items served as a manipulation check, namely to assess whether participants noticed the difference between conditions, whether they were aware of the nature of the experiment and whether they had participated in studies with the iCub before. All questionnaires were presented in Italian.

Procedure

After receiving the task instructions, participants were calibrated on the eye-tracker, and a practice session started. The iCub looked up from a more downward position to mimic mutual gaze and participants were asked to look at the iCub's face, even when objects appeared on the screens, until they heard a beep (750 Hz, 100 ms, random onset between 750 ms and 1250 ms after stimulus onset). This beep acted as their cue to look at either the left or the right object as quickly as they could, using their eyes, but not their head, as much as possible. Looking direction was free choice, but they were asked not to look constantly toward the same direction. Live gaze samples were transmitted from the eyetracker to the experimental software via ethernet, and as soon as 10 samples in either the left or right zone were collected, iCub also turned his head - according to the trial condition- 20° horizontally and 5° vertically (relative to the robot's frame of reference, mean movement onset time was 127 ms) to give the impression that he was looking at one of the objects. Participants could look at the object for as long as they wanted, and were

asked to then look back at the robot's face in their own time. The robot then turned his head forward again, after which participants pressed the spacebar to initiate the next trial. See Figure 1 for a trial example, and the video demonstration in the supplementary materials:

<https://osf.io/zxkwn> (or at <https://youtu.be/rRZ9KdYnCes>).

Whereas in the practice session the robot was anonymous (and introduced as such) and followed or unfollowed the participants gaze at random (50/50), participants did the task with a robot presented to them as either “Jimmy” or “Dylan” in the experimental session. The identity introduced as Jimmy followed the participant's gaze to the same object in 80% of the trials, and Dylan in 20% of the trials, otherwise, they looked at the other object. Thus, Jimmy was an identity with whom joint-attention was more often established than with Dylan.

The entire procedure was as follows. Participants did 16 practice trials, after which they started with either Jimmy or with Dylan (order counterbalanced between participants, two blocks of 40 trials for each). At the beginning of each block the robot introduced himself (“Hi, I am Jimmy/Dylan”). After two blocks participants were taken to a computer to complete 17 items of the InStance questionnaire, which were randomly selected for each participant, the whole Yoni Test and Godspeed items in that order. Each task was presented in such a way that they specifically referred to the robot whom the participant had just done the task with. In the meantime, the experimenter sat in the experimental room to “change the robot identity” and kept up appearances by typing vigorously.¹ Next, the participant did the task with the other identity and then filled out the questionnaires about the other identity. Finally, they completed the additional preference questions / manipulation checks, and were debriefed. Participants wore the

¹ As a result, the experimenter was not blind towards the different robot identities. The experimenter acted neutral and the same during both introductions however.

eyetracking glasses for 10-15 minutes per identity. Taken together, the experimental session took approximately 90 minutes.

Data processing

We used the Binocular-Individual Threshold (BIT) algorithm (23) to classify fixations per individual per block. The BIT algorithm bases velocity thresholds on inter-individual and between-task variability in fixations and thus offers an objective method for eye-tracking data classification. Next, we intended to use the eye-tracker's proprietary software to map the relevant areas of interest (AOIs) automatically, but it failed to classify the screens. Therefore, for each participant, the three relevant AOIs (left object, right object, iCub) were assigned by means of an offline K-means cluster analysis, using the SciPy Python library. This is a technique, often used in machine learning, that classifies each data-point (in our case fixation coordinates) based on the smallest distance to the gravitational centre of each group (in our case three AOIs). After each iteration, the algorithm then repositions the gravitational centre so that the sum of distances between each centre point and its data-points are minimal. This is repeated until repositions no longer change the classification and an optimum is reached (see Figure 2 for an example K-means cluster output).

Figure 2. Example output of a K-means cluster classification of one participant's fixation locations in one block. Three distinct AOIs are clearly visible: left screen, iCub, right screen. All cluster outputs can be found in the supplementary materials (<https://osf.io/zxkwn>)

For each trial, we calculated the onset latency of return-to-face saccades as the time between the first fixation on the selected object and the subsequent first fixation on iCub's face. We discarded trials in which no fixation on one of the objects was detected, trials in which no gaze samples were detected prior to the fixations of interest, trials with anticipatory fixations on the object (i.e. before the cue), and trials with clear saccade undershoots or overshoots back on the robot's face. We also discarded trials in which object-fixations were too short. In other words, at times participants would look back at the robot before the robot completed its

trajectory and the program was not able to register, thus act on, the fixation back on the face.

Trials with a delay of larger than 500 milliseconds between the fixation back on the face and the trigger that prompted iCub to return its head to the centre were discarded accordingly.

We trimmed the means of the onset latencies of these return-to-face saccades to fall within 2 standard deviations of the mean per participant per condition (average 1.55% outliers, S.D. = 0.79). To preserve power we excluded participant data if there were eight or fewer valid trials in the infrequent conditions; that is to say the condition in which Jimmy did not follow the gaze and the condition in which Dylan did follow the gaze (removed $n = 10$, mean invalid trials = 65.6% across all blocks). Additionally, we discarded two more participant whose means were outliers ($z > 2.5$) in any of the four conditions. The values were now normally distributed; all Kolmogorov-Smirnov $ps > .28$, final $n = 25$; thus 12 participants dropped in total.² These data were subjected to a 2 (Identity: Jimmy, Dylan) x 2 (Contingency: followed, unfollowed) ANOVA.

Results

Return-to-face saccades

There were no main effects of identity ($p = .37$) or contingency ($p = .52$). However, there was an interaction effect between identity and contingency; $F(1, 24) = 12.4$, $p = .002$, $\eta^2 = .34$ (see Figure 3). Follow-up paired-samples t-tests showed that for Jimmy (the robot with 80% of joint attention trials), the onset latencies of participants' return-to-face saccades were shorter ($M = 1753$ ms, $SD = 401$) when the robot followed than when it looked at the other object ($M = 1829$, $SD = 444$); $t(24) = 2.60$, $p = .016$, $d = .52$. The other pairwise comparisons (respectively

² As this may appear a high exclusion rate, we report here that when we employed a less stringent, yet perhaps more subjective inclusion criteria resulting in a final N of 32 participants, we found the same effects with similar effect sizes.

Jimmy-Followed – Dylan-Followed and Dylan-Followed – Dylan-Unfollowed) were not significant; all $ps > .055$ (Dylan-Followed $M = 1749$ ms, $SD = 299$; Dylan-Unfollowed $M = 1699$, $SD = 319$).³

Figure 3. Mean onset latencies for the return-to-face saccades for each identity and contingency in milliseconds. Error bars: ± 1 SEM.

³ These effects were replicated when we entered median return-to-face saccade onset-times into these models.

Questionnaires

Intentional Stance. Higher scores reflect more mentalistic as opposed to mechanistic explanations for the iCub's behaviour in the scenarios. Scores on the InStance questionnaire related to Jimmy ($M = 40.4\%$ mentalistic, $S.D. = 18.1$) did not differ significantly from the scores related to Dylan ($M = 37.9\%$ mentalistic, $S.D. = 18.9$); $p = .40$. 12 participants rated Jimmy as more mentalistic and 13 participants Dylan. Adding these two fairly even groups as a between-subjects variable in our model showed there was no relationship between these scores and the return-to-face saccade onset-times.

Yoni. The 2 (Identity: Jimmy, Dylan) $\times$ 2 (Order: first-order theory of mind, second-order theory of mind) $\times$ 2 (Theory of mind-type: affective, cognitive) repeated-measures ANOVA revealed a three-way interaction effect; $F(1, 24) = 4.6$, $p = .04$, $\eta_p^2 = .16$, but follow-up analyses did not yield further results. We found no other statistically significant differences for accuracy between the Yoni test relating to Jimmy and the Yoni test relating to Dylan; all $ps > .07$; all η_p^2 s $< .14$.⁴

Godspeed. Participants did not anthropomorphise Jimmy more than Dylan (Jimmy: $M = 16.7$, $S.D. = 3.6$; Dylan: $M = 16.2$, $S.D. = 4.1$; $p = .36$), but attributed greater likeability toward the former, $t(24) = 2.7$, $p = .013$, $d = .54$ (Jimmy: $M = 20.5$, $S.D. = 3.5$; Dylan: $M = 18.7$, $S.D. = 3.8$).⁵

Additional questions. Out of the total 25 participants whose data were analysed, 19 (76%; 78% for the entire sample) expressed a preference for Jimmy and 6 for Dylan. 21/25

⁴ There was a main effect of order; $F(1, 24) = 14.2$, $p = .001$, $\eta_p^2 = .37$. Participants made fewer errors on trials with first-order ToM scenarios (accuracy $M = .94$, $SE = .027$) than with second-order ToM scenarios ($M = .83$, $SE = .026$). A 2 (Identity: Jimmy; Dylan) $\times$ 2 (Eyes: leading, misleading) ANOVA yielded no significant results either, indicating that sensitivity to the gaze-direction of the Yoni-avatars was not affected by the interaction with the iCub identities.

⁵ Similar statistics were found when we conducted the Intentional Stance, Yoni, and Godspeed analyses on the entire sample of 37 participants.

participants indicated that they noticed differences between the two identities related to choice (6), gaze-following (4), imitation (3), or other miscellaneous reasons (8). 15 participants reported Jimmy to be the mentalistic robot and 10 as mechanistic.

Discussion

The interaction effect between identity and contingency for the return-to-face saccades indicates that attentional engagement with the robot was facilitated after initiating joint attention, but only if the robot typically followed the participant's gaze. This partly replicates previous findings that people are not merely sensitive to establishing ad-hoc joint attention, but that this sensitivity depends on implicit expectations set by previous interaction (8, 12, 24, 25). Whereas a similar -but screen-based- study found a main-effect in which return-to-face saccades were faster towards the joint-attention robot avatar than towards the disjoint-attention one (8), the current study suggests that an agent's joint-attention disposition drives further interaction. This implies that interaction with an embodied partner can evoke slightly different (social-)cognitive processes, relative to screen-based experimental paradigms and we speculate that this interaction effect reported here is a direct effect of deeper social engagement with embodied robots compared with 2-dimensional avatars. In other words, our participants might have been more engaged in the interaction, thus they were more sensitive to when the robot followed their gaze or not. This sensitivity to online behaviour was strengthened when they formed particular expectations of positive social interaction with the typically following robot. This is analogous with studies that found similar results in human-robot interaction compared with screen studies (26) but also with studies that found different gaze behaviour in real life settings with human presence relative to humans on a screen (11,15,27) and extends the invitation to carefully examine if social cognition in realistic scenario follows the same principles that were reported in

perhaps less ecologically valid experiments (e.g., 15,28). Furthermore, our findings provide evidence that re-establishing eye-contact after episodes of joint-attention (“closing the loop”) is facilitated when these episodes occur frequently as people appear to be sensitive toward this frequency.

Secondly, we hypothesised that Jimmy, the joint-attention identity, would evoke more favourable subjective attributions than Dylan, the disjoint-identity robot. Seeing that participants gave significantly higher likeability ratings and indicated a greater preference for “Jimmy” than for “Dylan”, we confirm this hypothesis. However, participants did not anthropomorphise one identity more than the other. This suggests that contingent gaze behaviour in robots is perhaps too subtle to be perceived as more human-like, and other factors such as appearance and movement kinematics might play a bigger role required for anthropomorphism (29). On the other hand, one could argue that likeability and preference are typically human-like attributions and in that respect we replicate findings that gaze behaviour is a strong predictor of positive personal attributions (5,7,8).

Whereas personal attributions were generally more favourable toward Jimmy, the joint-attention identity, we found no such differences in adopting the intentional stance, which can be explained in three ways. Firstly, in the literature Theory-of-Mind is typically described as a characteristic trait (30). In other words, individuals have a high or low de-facto ability to attribute mind to other agents regardless of subtle behavioural differences between the agents. Even if the degree of adopting the intentional stance can vary as a function of identity (17), the difference in attributed ‘identities’ may have been too implicit between the two robots which only differed on gaze-contingency but which were highly similar otherwise. Secondly, perhaps the photographic (InStance questionnaire) and schematic (Yoni test) representations of the

robot were semantically and temporally too distinct from the embodied agents that the participants interacted with, to detect any such differences in adopting the intentional stance. Finally, it is worth noting that at the time of the experiment, the Intentional Stance questionnaire was still in development. This questionnaire may prove to be a useful tool in detecting how readily individuals adopt the intentional stance toward iCub. However, how this adoption can be manipulated is still to be examined.

Potential limitations of our study include the fact that “changing the robot identity” may not be high in believability. We attempted to increase this by referring to the identities by their names from the beginning of the experimental sessions, by trying to give the impression that the experimenter was working hard to achieve this change whilst the participants were doing their first round of questionnaires, and by having the robots introduce themselves with their voice before each block. Perhaps a future study could use two actual different robots in a counterbalanced design, even though this will bring forward other practical challenges and may introduce new confound variables.

Finally, our study is a proof-of-concept of the idea that traditional paradigms of experimental cognitive psychology can be implemented in a more naturalistic human-robot interaction setup, without compromising experimental control. Although this work has been the first to meet the challenge of integrating an online eye tracker with stimulus presentation software as well as with an embodied robot simultaneously (which exhibited eye movement behaviour contingent on the eye movements of the participants), we managed to successfully meet the challenge both in experimental design/control and data collection. Furthermore, we also overcame the challenges involved in the off-line data processing. Most notably, as the proprietary automatic mapping of AOIs did not perform as desired, we came up with the solution of using a

machine learning technique to detect the three AOIs in each participant's gaze data. Naturally this was not as fine-grained as static, pixel-defined AOIs typically used in desktop eyetracking, but from checking against the eyetracker recordings, k-means clustering proved to be an efficient and accurate method. Further, there was a spatial and temporal heterogeneity in eye movements between participants. This made it difficult to specify an overarching catch-all fixation filter with the ideal signal-to-noise ratio, which is why we opted for individual thresholds with the BIT-algorithm. Notwithstanding the potential noise in the data however, to our knowledge this is the first time that a gaze-contingent paradigm has been successfully implemented with mobile eyetracking and an embodied robot platform. We therefore provided evidence that the sampling rate and spatial resolution of the latest mobile eye-trackers are adequate for studying the subtle attentional mechanisms previously thought quantifiable only with stringent screen-based paradigms.

In conclusion, our results show that attention towards those with whom we typically establish joint attention is facilitated when our gaze is followed, and that we have a preference for agents whose gaze following we successfully initiate. However, it is not yet clear whether we are more likely to adopt intentional stance towards those who display more contingent joint attention than those who do not. Furthermore, we demonstrated the feasibility of mobile eyetracking as a tool to carry out advanced gaze behaviour studies in more naturalistic settings which use embodied robots as “stimuli”, thereby forming a viable opportunity to increase ecological validity whilst maintaining excellent experimental control within the walls of the laboratory.

Data and code

The scripts and data used and reported in this study can be accessed online via <https://osf.io/zxkwn>.

Competing interests

We have no competing interests.

Author's contributions

AW and CW conceptualised the study. CW collected, processed, and analysed the data. AW and CW produced the manuscript and both gave final approval for publication.

Acknowledgements

The authors thank Davide De Tommaso (engineering the software) and Serena Marchesi (providing translations) for their significant contributions to this experiment.

Funding

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant awarded to AW, titled "InStance: Intentional Stance for Social Attunement." G.A. No: 715058).

The content of this publication is the sole responsibility of the authors. The European Commission or its services cannot be held responsible for any use that may be made of the information it contains.

References

1. Emery NJ. The eyes have it: the neuroethology, function and evolution of social gaze. *Neurosci Biobehav Rev.* 2000;24(6):581–604.
2. Langton SR, Watt RJ, Bruce V. Do the eyes have it? Cues to the direction of social attention. *Trends Cogn Sci.* 2000;4(2):50–9.
3. Driver IV J, Davis G, Ricciardelli P, Kidd P, Maxwell E, Baron-Cohen S. Gaze perception triggers reflexive visuospatial orienting. *Vis Cogn.* 1999;6(5):509–40.
4. Friesen CK, Kingstone A. The eyes have it! Reflexive orienting is triggered by nonpredictive gaze. *Psychon Bull Rev.* 1998;5(3):490–5.
5. Bayliss AP, Murphy E, Naughtin CK, Kritikos A, Schilbach L, Becker SI. ‘Gaze leading’: Initiating simulated joint attention influences eye movements and choice behavior. *J Exp Psychol Gen.* 2013;142(1):76–92.
6. Schilbach L, Wilms M, Eickhoff SB, Romanzetti S, Tepest R, Bente G, et al. Minds made for sharing: initiating joint attention recruits reward-related neurocircuitry. *J Cogn Neurosci.* 2010;22(12):2702–15.
7. Grynszpan O, Martin J-C, Fossati P. Gaze leading is associated with liking. *Acta Psychol (Amst).* 2017;173:66–72.
8. Willemse C, Marchesi S, Wykowska A. Robot Faces that Follow Gaze Facilitate Attentional Engagement and Increase Their Likeability. *Front Psychol.* 2018;9:70.
9. Wiese E, Metta G, Wykowska A. Robots as intentional agents: using neuroscientific methods to make robots appear more social. *Front Psychol.* 2017;8:1663.
10. Wykowska A, Chaminade T, Cheng G. Embodied artificial agents for understanding human social cognition. *Phil Trans R Soc B.* 2016;371(1693):20150375.

11. Foulsham T, Walker E, Kingstone A. The where, what and when of gaze allocation in the lab and the natural environment. *Vision Res.* 2011;51(17):1920–31.
12. Estimating the reproducibility of psychological science. *Science* [Internet]. 2015 Aug 28;349(6251). Available from: <http://science.sciencemag.org/content/349/6251/aac4716.abstract>
13. Van Bavel JJ, Mende-Siedlecki P, Brady WJ, Reinero DA. Contextual sensitivity in scientific reproducibility. *Proc Natl Acad Sci.* 2016 Jun 7;113(23):6454.
14. Nosek BA, Alter G, Banks GC, Borsboom D, Bowman SD, Breckler SJ, et al. Promoting an open research culture. *Science.* 2015 Jun 26;348(6242):1422.
15. Laidlaw KEW, Foulsham T, Kuhn G, Kingstone A. Potential social interactions are important to social attention. *Proc Natl Acad Sci.* 2011 Apr 5;108(14):5548.
16. Dalmaso M, Gareth S, Bayliss AP. Re-encountering individuals who previously engaged in joint gaze modulates subsequent gaze cueing. *J Exp Psychol Learn Mem Cogn.* 2016;42(2):271–84.
17. Dennett D. 1987. *The intentional stance*. London: MIT Press; 1987.
18. Marchesi S, Ghiglini D, Ciardo F, Baykara E, Wykowska A. Do we adopt the Intentional Stance towards humanoid robots? *PsyArXiv* [Internet]. 2018 Jul 6 [cited 2018 Jul 16]; Available from: <https://psyarxiv.com/6smkq/>
19. Metta G, Natale L, Nori F, Sandini G, Vernon D, Fadiga L, et al. The iCub humanoid robot: An open-systems platform for research in cognitive development. *Neural Netw.* 2010 Oct 1;23(8):1125–34.
20. Mathôt S, Schreij D, Theeuwes J. OpenSesame: An open-source, graphical experiment builder for the social sciences. *Behav Res Methods.* 2012 Jun 1;44(2):314–24.

21. Shamay-Tsoory SG, Aharon-Peretz J. Dissociable prefrontal networks for cognitive and affective theory of mind: A lesion study. *Neuropsychologia*. 2007 Jan 1;45(13):3054–67.
22. Bartneck C, Kulić D, Croft E, Zoghbi S. Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots. *Int J Soc Robot*. 2009 Jan 1;1(1):71–81.
23. van der Lans R, Wedel M, Pieters R. Defining eye-fixation sequences across individuals and tasks: the Binocular-Individual Threshold (BIT) algorithm. *Behav Res Methods*. 2011 Mar 1;43(1):239–57.
24. Stephenson LJ, Edwards SG, Howard EE, Bayliss AP. Eyes that bind us: Gaze leading induces an implicit sense of agency. *Cognition*. 2018 Mar 1;172:124–33.
25. Kompatsiari K, Ciardo F, Tikhonoff V, Metta G, Wykowska A. Bidding for joint attention: On the role of eye contact in gaze cueing. *PsyArXiv* [Internet]. 2018 Apr 3 [cited 2018 Jul 16]; Available from: <https://psyarxiv.com/mx28g/>
26. Kompatsiari K, Pérez-Osorio J, Metta G, Wykowska A. Neuroscientifically-grounded research for improved human-robot interaction. 2018;
27. Jung Y, Lee KM. Effects of physical embodiment on social presence of social robots. *Proc PRESENCE*. 2004;80–7.
28. Pan X, Hamilton AF de C. Why and how to use virtual reality to study human social interaction: The challenges of exploring a new research landscape. *Br J Psychol*. 2018;
29. Dalibard S, Magnenat-Talmann N, Thalmann D. Anthropomorphism of Artificial Agents: A Comparative Survey of Expressive Design and Motion of Virtual Characters and Social Robots. In: *Workshop on Autonomous Social Robots and Virtual Humans at the 25th*

Annual Conference on Computer Animation and Social Agents (CASA 2012) [Internet]. Singapore, Singapore; 2012 [cited 2018 Jul 16]. Available from: <https://hal.archives-ouvertes.fr/hal-00732763>

30. Frith C, Frith U. Theory of mind. *Curr Biol*. 2005 Sep 6;15(17):R644–5.