This is an accepted manuscript version of the paper published as:

Willemse, C., Abubshait, A., & Wykowska, A. (2021). Motor behaviour mimics the gaze response in establishing joint attention, but is moderated by individual differences in adopting the intentional stance towards a robot avatar. *Visual Cognition*, 1-12. DOI: 10.1080/13506285.2021.1994494

Motor Behaviour Mimics the Gaze Response in Establishing Joint Attention, but is Moderated by
Individual Differences in Adopting the Intentional Stance Towards a Robot Avatar

Cesco Willemse, Abdulaziz Abubshait, and Agnieszka Wykowska*

Social Cognition in Human-Robot Interaction, Istituto Italiano di Tecnologia, Genoa, Italy

Author Note

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

The data and scripts related to this manuscript can be retrieved on <https://osf.io/brjqd/>

* Corresponding author:

Agnieszka Wykowska

Istituto Italiano di Tecnologia

Social Cognition in Human-Robot Interaction

via Enrico Melen 83

16152 Genova, Italy

Phone: +39 010 8172242

Email: agnieszka.wykowska@iit.it

Abstract

Leading another person's gaze to establish joint attention facilitates social interaction. Previously it was found that we look back at agents who engage in joint attention frequently and more quickly than agents who display this behaviour less often. However, it is unclear if motor behaviour mimics this following behaviour. This paper serves to fill in two knowledge gaps on the topic. Firstly, we examine whether this looking-back behaviour is replicated by a manual response. In an online, eyetracker-free task in which participants were asked to select one of two objects, one robot identity on the screen followed the selection most of the time, whilst the other looked at the other object most of the time. Participants moved back to the robot more quickly if the robot which most of the time followed their movement looked at the same object relative to when it did not. We found no such difference for the robot which most of the time did not follow participants. Secondly, we examined how individual differences in autistic traits and readiness to adopt the intentional stance toward artificial agents affect how participants' behaviour changed over time during the experiment. The results showed that autistic traits influenced participants' motor responses over time, independent of condition, whereas this response is only sensitive to the intentional stance when the robot usually did not follow the participant's object selection. Taken together, this indicates that whereas individual differences may not fully predicate reflexive social behaviour, there is a potential overlap between evoked gaze behaviour and manual motor actions.

Keywords: joint attention, gaze-leading, social robotics, motor response, growth-curve modelling

Motor Behaviour Mimics the Gaze Response in Establishing Joint Attention, but is Moderated by Individual Differences in Adopting the Intentional Stance Towards a Robot Avatar

The gaze-leading paradigm

Gaze is an important signal in social interaction as it allows us to orient another person's attention to a locus of interest (Emery, 2000). While gaze-cueing paradigms examine how strongly the gaze direction of others shifts our focus of attention (Friesen & Kingstone, 1998; Frischen, Bayliss, & Tipper, 2007), gaze-leading paradigms have been used to examine the cognitive processes involved in initiating joint attention, which is akin to how we interact with others in daily life. Previous gaze-leading paradigms have shown that when agents follow a person's gaze, people subjectively prefer them to agents that do not follow their gaze (Bayliss et al., 2013; Grynszpan, Martin, & Fossati, 2017; Willemse, Marchesi, & Wykowska, 2018; Willemse & Wykowska, 2019), perceive them as more trustworthy (Dalmaso, Edwards, & Bayliss, 2016), have a greater implicit sense of agency (Stephenson, Edwards, Howard, & Bayliss, 2018), and show higher brain activity in regions related to social reward (Schilbach et al., 2010).

More directly, people tend to be quicker to reorient their attention to the agent after it follows their gaze, relative to when the agent does not establish joint attention, especially when the agent frequently follows the person's gaze (Dalmaso et al, 2016; Willemse et al., 2018, Willemse & Wykowska, 2019). This could indicate that establishing joint attention facilitates interaction between two persons, or between a person and a robot (Willemse et al., 2018; Willemse & Wykowska, 2019). The latter is particularly informative for robot-designers, as they can implement the mechanisms and behaviours that evoke attentional engagement (Willemse &

Wykowska, 2019) and perception of a robot as an intentional agent (Bossi et al., 2020; Marchesi et al., 2019).

The study presented here builds on two of our previous gaze-leading studies: Willemse et al. (2018), and Willemse and Wykowska (2019). In both of these studies, we asked participants to select their preferred object of a laterally presented pair using their gaze. A robot then looked at either the same object or at the other one. The participants were asked to look back at the robot at their own pace. How frequently the robot responded in a joint-attention fashion by looking at the same object depended on the disposition of the robot: in one condition the robot had a joint disposition and looked at the same object 80% of the time, versus 20% for the robot with a disjoint disposition. In the remaining trials, the robot looked at the other object of the pair. In the study where the robot appeared on a computer screen, we found that the return-to-face saccade onset time was quicker for the robot who typically responded to initiating joint attention. In the study that utilised an embodied robot, we found that participants were slower with returning their gaze to the robot's face when it did not follow their gaze relative to when it did, but only for the joint disposition.

The manual-motor response in initiating joint attention

While other previous work has utilised gaze-leading paradigms as well (e.g., Schilbach et al., 2010), to our knowledge, no previous studies have adapted this paradigm to examine if similar joint attention processes are displayed by motor responses. Toddlers learn to point at things in a social manner by the age of 12 months (Leung & Rheingold, 1981). In adults, pointing is an effective tool to establish joint attention (Bangerter, 2004). Understanding how motor-responses mimic/differ from gaze behaviour would be helpful for research as, on a practical level, mouse-

tracking is implemented with greater ease than eyetracking and lends itself well for online testing when necessary.

Whereas the primary motivation of examining whether the manual motor response show concordance with oculomotor response in initiating joint attention was driven by this need for online testing, we also carefully considered the theoretical implications of this potential link. Finding an overlap between manual motor mechanisms and oculomotor mechanisms would be informative about the embodied oculomotor underpinnings of social attention. Neuroscience has demonstrated that directional vectoring of both arm movements and saccades are similar (Georgopoulos, 1990). Therefore, they may both be controlled and inhibited by similar attentional mechanisms. Indeed, attentional selection activates saccades as well as reaching trajectory in a similar manner (Tipper, Howard, & Paul, 2001) and it has been shown that eye movements as well as manual reaching movements are coupled with attentional focus (Baldauf, Rolf, & Deubel, 2006; Deubel & Schneider, 1996).

Moreover, there is an overlap in the brain regions involved in cognitive control of both eye movements and hand movements (Leung & Cai, 2017). Further evidence from this overlap comes from the areas of development and learning, and it is no longer thought that hand-eye coordination is visually guided, but occurs in parallel with eye-movements in a multi-modal fashion and is embodied (Corbetta, Wiener, Thurman, & McMahon, 2018). Specifically related to joint attention, it was found that toddlers with good hand-eye coordination, i.e. looking at the object they are manipulating, frequently establish joint attention with their parent who they were doing the task with (Yu & Smith, 2017). In sum, given that the motor cortex and the visual cortex respond similarly to attentional execution and that these action executions develop in parallel, adds to the notion the hand-movement response could be similar to saccades in our paradigm

where attention shifts between an object and a face. As this has not been studied in an initiating joint attention task before, to our knowledge, it could open up future research questions about the exact link between manual-motor movements and oculomotor movements in joint attention, for example with the addition of simultaneous eye-tracking and covert/overt attentional shift manipulations.

Aims of the study

Examining the relationship between oculomotor behaviour and manual responses

The first aim of the current study was to investigate whether the speeded return “saccade” after initiating joint attention is also displayed by the manual motor response. If motor responses mirror similar cognitive mechanisms to those exhibited by saccadic responses, the return motion to the robot’s face should be faster when the robot follows the gaze relative to when it does not. This effect should be enhanced when the robot typically follows – in the joint condition, compared with when it typically looks at the non-selected object – in the disjoint condition; a pattern paralleling our previous results.

Individual differences in the manual motor response to initiating joint attention

Additionally, in our previous studies we examined how the mechanisms of interest were affected by individual differences in autistic traits and adopting the intentional stance. Individuals higher in autistic traits often have difficulties in social communication and interaction (Baron-Cohen, Wheelwright, Skinner, Martin, & Clubley, 2001). A number of gaze-cueing studies have reported a modulation of social gaze behaviour by autistic-like traits (Morgan, Smith, & Freeth, 2021; Wiese, Müller, & Wykowska, 2014; Wykowska, Kajopoulos, Ramirez-Amaro, & Cheng, 2015), less is clear about how autistic traits affect the cognitive processes described above (Willemse et al., 2018).

Similarly, individual tendency to adopt the intentional stance toward a robot may modulate the extent to which its eyes are seen as a social signal (Willemse & Wykowska, 2019). The intentional stance (Dennet, 1987) is a strategy to explain a system's actions and behaviour as being driven by mental states, such as intentions or desires – we do this naturally with other persons. A stopping car on the other hand is more appropriately explained by the mechanistic strategy of adopting the design stance: the mechanical processes that were activated when the brake pedal was engaged. Intuitively, humanoids may be perceived as either more human-like, which would evoke the intentional stance, or as machine-like, which would evoke the design stance (Bossi et al, 2020; Marchesi et al., 2019). In previous studies (Bossi et al., 2020; Marchesi et al., 2019; Spatola, Marchesi, & Wykowska, 2021), it was suggested that individuals might differ with respect to their tendency to adopt the intentional stance towards robot. Therefore, it remains to be answered whether individual tendency of adopting the intentional stance towards a robot facilitates social interaction with the agent.

Including these individual differences in this study was driven by the prior studies outlined above. However, in addition to this main motivation, we aimed at examining how these traits affect the effects of interest over time. One point of criticism with the gaze-leading paradigm is that gazers learn the likelihood of an agent following their gaze, which influences how fast a gazer re-establishes joint attention with the agent (i.e., faster saccadic returns are due to a non-social congruency signal). While the prior criticism assumes that this statistical learning occurs in a non-social fashion, we argue that social dispositions can influence how we learn about others' social behaviour (i.e., while this effect is due to *statistical* learning, it could nevertheless be due to *social* learning).

One way to address this criticism is to examine the temporal dynamics of gazers' saccadic reactions on a trial-by-trial basis using growth-curve modelling procedures. Traditionally used for longitudinal data, growth-curve modelling is a method that captures within-subject changes over time (i.e. reaction times) whilst accounting for individual differences between participants (between-subject reaction times and questionnaire scores) in repeated measures data (Curran, Obeidat, & Losardo, 2010).

Therefore, our secondary aim was to apply growth curve modelling to autistic traits and tendency for adopting the intentional stance as predictors of the learning rate in the task. This is an exploratory research question for which we tentatively predict that social dispositions (i.e., individual differences in social traits) influence the rate to which participants learn the task. Specifically, we hypothesize that those higher in autistic traits (i.e., a social disposition that influences how well people process social information) will learn at a slower rate and adapt slower to the task than those lower on the spectrum. Similarly, those who attribute more intentional stance toward robotic agents at baseline level (i.e., a social disposition that is related to the likelihood that people treat an agent as possessing beliefs, intentions, and feelings), will see the agent as more intentional rather than a mechanistic artefact, which would speed up the learning rate compared with those with less attribution of intention.

Method

Participants

We recruited 100 participants (66 male, 33 female, one transgender female, Age $M = 22.7$ years, $S.D. = 4.37$) via the online recruitment platform Prolific. Participants received a payment of

£ 2.20 for completion. All participants provided informed consent, adhering to the Declaration of Helsinki and as approved by the local Ethics Committee (Comitato Etico Regione Liguria).

We defined our sample size of 100 based on the following criteria. Both our lab based studies had a prospected sample size of 32 participants and the effects of main interest had effect sizes of $\eta^2 = .27$ and $\eta^2 = .34$ respectively. However, there are a few key differences between the study reported here and the previous ones. Firstly, the current design was fully-between subjects: half of the participants would do the task with a robot who typically followed the participants object selection, and the other half would interact with a robot who typically looked at the non-selected object. Additionally, online studies do not warrant a well-controlled environment (i.e., more noise in the data) and growth curve modelling needs to have larger sample sizes when researchers expect more noise in the data (Vallejo et al., 2019). Thus, in order to aim for good engagement with the task, we decided to decrease the duration of the experimental procedure and presented 80 trials rather than 160 in the previous studies. With these caveats in mind we deducted that 40 participants per condition would be sufficient (80 in total). As we expected a data rejection between 15-25% based on online quality, we recruited 100 participants in total (50 in each condition).

Materials

For the choice-leading task we adopted the stimuli from Willemse et al. (2018). We presented an image of the humanoid robot iCub (Metta et al., 2010) at 300 x 300 pixels and the object pairs at 240 pixels in their longest dimension. We collected responses using a mouse or trackpad. We programmed the task in PsychoPy3 v 2020.1.2 (Peirce et al., 2019) and presented it via Pavlovia on the participant's web browser.

We presented the 34-item Intentional Stance Questionnaire (ISQ; Marchesi et al., 2019) in two 17-item parts, one presented before the task and one afterwards. The ISQ is developed to examine tendencies in adopting intentional stance towards robots (i.e., predicting and/or explaining robots' behaviour with reference to mental states). Each item shows a scenario of the iCub in three images. Below the images, two sentences are displayed: one explaining the scenario in mechanistic terms, and the other in intentional terms. A slider response between the sentences is used to indicate which sentence is most representative in the participant's opinion. Total scores range from 0-100 with higher scores indicating more attribution of intentional stance. We split the items into the two parts according to a factor analysis on previous data, which indicated the presence of three factors. We assigned items to the two parts of the ISQ in a way the number of items per each factor would be balanced across the two parts. The order of the parts was counterbalanced between participants.

Furthermore, we used the AQ-10 (Allison, Auyeung, & Baron-Cohen, 2012) to measure autistic traits. This is a brief version of the Autism Quotient (Baron-Cohen et al., 2001) and is reported to have similar sensitivity (Allison et al., 2012; Booth et al. 2013). This questionnaire was developed as a screening tool for adult Autism Spectrum Disorder, and can be used by clinicians to indicate further diagnostic assessment (Booth et al., 2013). More recently the AQ-10 was suggested to have adequate validity for measuring autistic traits (Lundin, Kosidou, & Dalman, 2018) and has been used as a tool to examine individual differences in non-clinical samples in other studies (e.g. Bertrams & Schlegel, 2020; Gollwitzer, Martel, McPartland, & Bargh, 2019). Items are scored with a 4-point Likert scale, ranging from "definitely disagree" to "definitely agree". Both of the questionnaires were presented with SoSci Survey (Leiner, 2019)

Procedure

After signing up for the study, participants were directed to the SoSci Survey platform for the first part of the ISQ. For each item, they were instructed to move the slider toward the position between the two sentences that they felt was most representative of their interpretation of the scenario. They were then redirected to Pavlovica for the main leading task.

Before the experimental task commenced, participants were requested to use a mouse and to indicate whether they were using a mouse or a trackpad¹. They were instructed that they would see two objects next to a robot and to take their time to think about which object they “liked best” and then click on its centre. If they were ready to see the next objects, they could click on the centre of the robot’s face. This was followed by a practice session of four trials with objects that were not used in the task.

Each trial started with a fixation cross presented for 500 milliseconds, followed by a screen with the iCub gazing forward flanked by two objects. Three hundred milliseconds after a mouse click on one of the objects, the robot redirected its gaze toward the clicked object or toward the other object. In the joint condition, the robot looked at the clicked-on object in 80% of the trials and at the other object in 20% of the trials. In the disjoint condition, this proportion was reversed. The trial was ended by a mouse click on an invisible circular area (160-pixel diameter) on the robot’s face and was followed by a 500-millisecond blank screen. See Figure 1 for an example trial. There were 80 trials in total, presented in random order, with a self-paced break after 40 trials. Each object pair was thus presented ten times in random order and left-right location of each object-pair was counterbalanced between trials.

¹ Based on a similar pilot study, we did not expect this to be a relevant factor. 52 out of our final $n = 80$ participants self-reported using the mouse. No manipulation checks involving this variable were statistically significant.

After the task, participants were redirected to the SoSci Survey platform to complete the second part of the ISQ, followed by the AQ-10. Finally, we asked them to write down which of the items they remembered, as well as whether they noticed anything about the robot's behaviour, as a manipulation check. The experiment took approximately twenty minutes to complete.

Figure 1. Trial procedure. **A:** A fixation cross appears for 500 ms. **B:** The robot face and a pair of objects appear. The participant decides which of the laterally presented objects they like best. **C:** 300 ms after a mouse click on one of the objects, the robot redirects its gaze. This example is of an ‘unfollowed’ trial. **D:** To start the next trial, the participant clicks on the centre area of the robot’s face. We defined the return-to-face onset as the time passed between the mouse click on the object in panel C and the mouse click on the robot’s face in panel D.

Data processing

Our dependent measure was the time between the mouse click on one of the objects and the subsequent (returning) mouse click on the robot. As the eye-movement of the robot took place 300 milliseconds after the object was selected, response times lower than this were discarded.

Furthermore, reaction times were trimmed for each participant per following condition (± 2.5

S.D.). Based on these criteria, we excluded participants if 20% of the trials were discarded; final $n = 80$. Trimmed means were entered into a 2 x 2 mixed ANCOVA with ‘disposition’ (joint, disjoint) as between-subjects factor, and ‘following’ (followed, unfollowed) as within-subjects factor. Participants’ total scores on the AQ-10 were mean-centred and added as a covariate. For the effect on intentional stance, we ran a 2 x 2 between-subjects ANOVA on the ISQ scores with ‘time’ (pre, post) and ‘disposition’ (joint, disjoint) as factors.

Additionally, growth curve modelling was carried out to examine participants’ return-to-face onset times on a trial-by-trial basis. Growth curve modelling was done by fitting six different models that incrementally increased in their degree of complexity to the data and comparing them using a nested model comparison. The nested model comparison allows for comparing the models and determining which model fits the data best without overfitting the data (i.e., a parsimonious model that explains sufficient amounts of variance). Models that contain more predictors are considered more complex. The variations in the models are outlined below. The nested model comparison was carried out separately for each individual difference.

The nested model comparison allows us to infer if the models fit significantly different from the reference model (i.e., a parsimonious model that predicts reaction time from the grand mean of reaction time). As a second step, we examined the models that fit significantly different from the reference model and compare each model’s Bayesian Information Criterion (BIC). The BIC is a model fitting index that allows us to determine which of the models fit best. Generally, lower BIC indices suggest a better fit (i.e., explaining the data without overfitting). We chose the BIC as the model index of interest as it penalizes models for their complexity to ensure that we are not overfitting the data with nuisance variables.

The models that we fit included the intercept-only model (i.e., the reference model: M1), a model that only contained the growth predictor (i.e., a linear function that tracks changes in reaction time on a trial-by-trial basis regardless of condition: M2), a model that contained the growth predictor as a function of disposition (i.e., joint vs. disjoint) and gaze-following (Following vs. Unfollowing: M3), a model that contained the growth term as a function of disposition, gaze-following and individual difference score (i.e., Autism Quotient score or the ISQ score: M4), a model that contained the growth term as a quadratic function as a function of disposition, gaze-following and individual difference score (M5) and finally a model that contained the growth term as a cubic function as a function of disposition, gaze-following and individual difference score (M6). All the models varied the intercept for each participant. After finding the model of best fit, we examined the individual predictors to see how individual predictors influenced participant's reaction time. All variables were mean-centred prior to including them in the models. To conclude that individual differences differed overtime, we looked at significant interactions between any of the growth terms and individual differences scores.

Results

Manual-motor return-to-face onset times

We found a significant main effect of following behaviour; $F(1, 77) = 20.9, p < .001, \eta^2 = .009$. Participants' return clicks were faster after the robot followed their gaze ($M = 877$ ms, $S.D. = 318$) than when it looked at the other object ($M = 937$ ms, $S.D. = 331$). There was no significant between-subjects main effect of disposition; $p = .43$. The interaction effect between following

behaviour and disposition was significant; $F(1, 77) = 17.4, p < .001, \eta^2 = .007$.² Post-hoc analysis showed that for the joint disposition, return-clicks slowed down after the robot did not follow its gaze toward the selected object ($M = 992$ ms, $S.D. =$) compared with when it looked at the same object ($M = 878$ ms, $S.D. = 302$); $p < .001$, Tukey corrected, *Cohen's d* = 0.36. In contrast, for the disjoint disposition return-clicks had a similar onset across both following behaviour conditions (followed $M = 876$ ms, $SE = 337$; unfollowed $M = 881$ ms, $SE = 320$; $p = .993$, Tukey corrected), see Figure 2.

Figure 2. Return-click onset times for the joint and disjoint dispositions when the gaze followed to the same object or unfollowed. When the robot's gaze typically followed (joint disposition),

² Given the unbalanced number of trials in each condition (i.e. 80 / 20%), we corroborated our main trimmed reaction time analyses with the median reaction times as the dependent variable to justify reporting mean analyses subsequently. Main effect of following behaviour: $F(1, 77) = 7.35, p = .008, \eta^2 = .003$; main effect disposition: $F(1, 77) = 0.43, p = .51$; interaction effect $F(1, 77) = 12.7, p < .001, \eta^2 = .005$.

the return-to-face clicks were faster in trials where the gaze followed relative to when the agent looked at the other object. Gaze-following did not significantly affect return-to-face mouse movements in the disjoint disposition condition. *Note.* error bars = within subjects 95% CI.

Individual differences as covariate

There was a main effect on return-click onset times for the AQ-10 score ($M = 3.41$, $S.D. = 1.77$); $F(1, 77) = 7.24$, $p = .009$, $\eta^2 = .081$. AQ-10 scores correlated positively with overall return-to-face mouse click onset times; $r = .29$, $p = .010$. However, the AQ-10 did not interact significantly with following behaviour of the robot ($p = .43$).³

For the ISQ, we found no significant main effects for disposition ($p = .80$) or pre/post ISQ ($p = .17$). There was a marginally significant interaction effect; $F(1, 78) = 3.74$, $p = .057$, $\eta^2 = .009$. In exploration, post-hoc analyses showed that the post-ISQ score ($M = 39.7$, $S.D. = 17.5$) did not change significantly relative to the pre-ISQ ($M = 38.8$, $S.D. = 16.1$) for the participants who interacted with the robot in the disjoint disposition ($p = 1$; Tukey corrected), or for those who interacted with the robot in the joint (post $M = 35.6$, $S.D. = 19.9$; pre $M = 41.1$, $S.D. = 16.2$); $p = .093$; Tukey corrected.

Individual differences – growth curve modelling

The nested model comparison examining the moderating effect of the AQ-10 revealed that the model of best fit is the growth model that contained growth term as a quadratic function as moderated by disposition, gaze-following and the ISQ. See Table 1 for the results of the nested model comparison. The model showed a significant effect of the linear growth function, the quadratic growth function and the interaction between the linear growth function and AQ-10, suggesting

³ In exploration, we examined these analyses with accuracy error of landing coordinates relative to the center of the face image. We found a main effect of following behaviour in which followed trials yielded greater accuracy than unfollowed trials, $F(1, 77) = 10.1$, $p = .002$, $\eta^2 = .022$. However, there were no other statistically significant main effects or interaction effects. Further, the trade-off between accuracy and reaction times is not straight-forward, as illustrated by the lack of correlation between reaction time and accuracy ($r = .002$, $p = .99$).

that autistic traits influenced participants' performance over time, however, this effect was unrelated to the disposition or the robots' gaze-following, see Figure 3.

Table 1.

Summary of the nested model comparisons

	BIC	AQ			BIC	ISQ		
		χ^2	df	p		χ^2	df	p
M1	103417				103417			
M2	102235	1191.2	1	<.001	102235	1191.2	1	<.001
M3	102257	30.78	6	<.001	102257	30.78	6	<.001
M4	102305	22.61	8	.003	102316	11.91	8	.15
M5	102082	294.65	8	<.001	102090	297.14	8	<.001
M6	102116	36.51	8	<.001	102123	37.41	8	<.001

Note. The reference model for all nested model comparisons was the intercept model.

Figure 3. Predicted values from the study. The graph illustrates the participants' performance over time based on AQ-10 score. Although all participants responded faster across the experiment, those who scored higher on the AQ-10 had a steeper slope compared to those who scored lower. Y axes: Return-to-face saccade onset in milliseconds. X axis: Trial number. Legend: Blue- 1 SD above the mean, Yellow-1SD below the mean of the AQ-10. The shaded regions illustrate the asymptotic 95% confidence intervals of the prediction value.

Similarly, the nested model comparison examining the moderating effect of ISQ revealed that the model of best fit is the growth model that contained growth term as a quadratic function as moderated by disposition, gaze-following and the ISQ. The model showed a significant interaction effect between the quadratic function, Disposition, Following and ISQ score ($b = .006$, $SE = t(7047) = 2.07$, $p = .03$). The interaction suggests slope differences over time for those who scored high vs. those who scored low on the ISQ when the robot followed their gaze vs. when it did not follow their gaze. This effect was only evident for the disjoint disposition, but not the joint disposition condition. The interaction can be found in the predicted values of the model, which is illustrated in Figure 4. Specifically, those who scored high on the ISQ were more likely to slow their responses over time for the disjoint robot when it followed their gaze (or in this specific study, their mouse click) compared to those who scored low on the ISQ performed similarly across the entire experiment. This relationship is reversed when the robot did not follow the participants' gaze. In other words, those who scored high on the ISQ were more likely to slow their responses over time, while this relationship was weaker for those who scored low on the ISQ. This effect was not evident for the joint disposition condition. The model summary can be found in Table 2.

Table 2
Summary of the growth model from Study 3 for the AQ-10 and ISQ

	AQ-10				ISQ			
	b	SE	t	p	b	SE	t	p
Intercept	1003	46.89	21.39	< .001	1004	46.69	21.5	< .001
Growth	-7.26	.73	-9.84	< .001	-7.18	.73	-9.71	< .001
Following	2.2	20.92	1.05	.29	21.85	20.83	1.04	.29
Disposition	2.35	63.98	.36	.71	17.79	63.66	.28	.78
ID	5.16	25.3	.2	.83	-3.42	28.99	-1.18	.23
Growth ²	.15	.03	4.45	< .001	.15	.03	4.36	< .001
Growth X Following	-9.57	.79	-1.2	.22	-1.11	.79	-1.4	.15
Growth X Disposition	.48	.79	.61	.53	.41	.79	.5	.6
Following X Disposition	-29.22	29	-1	.31	-28.58	28.91	-.98	.32
Growth X ID	.87	.38	-2.25	.02	-.09	.04	-1.99	.04

Following X ID	.43	11.1	.03	.96	1.58	1.13	1.2	.22
Disposition X ID	24.77	36.66	.67	.5	5.93	3.95	1.5	.13
Growth ^2 X Following	-.003	.03	-.09	.92	.001	.03	.03	.97
Growth ^2 Disposition	-.005	.03	-.13	.89	.001	.03	.03	.96
Growth ^2 X ID	.003	.01	.17	.85	.001	.002	1.92	.05
Growth X Following X Disposition	.82	1.12	.72	.46	1.27	1.12	1.13	.25
Growth X Following X ID	.38	.41	.91	.35	.06	.04	1.4	.15
Growth X Disposition X ID	.75	.42	1.78	.07	.06	.04	1.44	.15
Following X Disposition X ID	3.53	16.62	.21	.83	-.93	1.78	-.52	.6
Growth ^2 X Disposition X Following	2.67	.05	.5	.61	.01	.05	.25	.8
Growth ^2 X Following X ID	-.01	.01	-.68	.49	-.005	.002	-2.33	.01
Growth ^2 X Disposition X ID	-.006	.01	-.32	.74	-.005	.002	-2.18	.02
Growth X Following X Disposition X ID	-.93	.63	-1.46	.14	.08	.07	-1.15	.25
Growth ^2 X Following X Disposition X ID	.02	.02	.9	.36	.006	.003	2.07	.03

Note. All significance testing was based on 7047 degrees of freedom. ID indicates the respective individual difference.

Figure 4. Predicted values. The panels illustrate the interaction between disposition, ISQ score, and Following overtime. The model shows that the interaction is such that when the robot follows the participants gaze, and is disjoint, return-to-face mouse clicks became much slower over time for those who scored high on the ISQ, while no large changes occurred over time for those

who scored low on the ISQ. Y axes: Return-to-face mouse click onset in milliseconds. X axis: Trial number. Legend: Blue-1 SD below the mean, Yellow-1SD above the mean of ISQ. The shaded regions illustrate the asymptotic 95% confidence intervals of the prediction value.

Discussion

The manual response overlaps with the oculomotor response

We implemented a gaze-leading, or rather a choice-leading paradigm using mouse responses instead of eyetracking. This enabled us to examine whether manual behaviour in social interaction is analogous to previous findings with eye movements. Given the corroboration with findings from our previous studies, this is likely the case. The return-clicks showed that our participants interacted more swiftly with the agent when it followed their gaze to the same object. This was driven by the disposition of the agent. Establishing joint attention did not affect the manual response toward the disjoint robot. However, when the robot typically followed the gaze, participants slowed down when it infrequently looked at the non-selected object. This indicates that the violation of expectations in this type of interaction-based task are only disruptive when joint attention is expected, which matches our hypothesis.

Interestingly, these results closely mimic our previous eye-tracking findings. This demonstrates that manual movements are potentially activated by similar attentional mechanisms as eye movements. Our paradigm does not lend itself to closely tease apart the overlap between the mechanisms that regulate eye and hand movements conclusively, but fits well with what we currently understand about this link. Vision plays a key role in (manual) motor planning (Sober & Sabes, 2005), and eye movements actively guide hand movements (Land & Hayhoe, 2001). Attentional focus is strongly linked with both eye movements and manual reaching movements (Baldauf, Rolf, Deubel, 2006; Deubel & Schneider, 1996). On a neural level, both modalities utilise overlapping networks of the brain (Vesla & Crawford, 2012). Furthermore, the motor

execution is vectorised in a similar way (Georgopoulos, 1990), they tap into similar mechanisms of cognitive control (Tipper et al., 2001), and they develop interactively with one another (Corbetta et al., 2018). However, how exactly eye and hand movements and their trajectories overlap, whether they are causally involved in sensorimotor learning, and how they co-develop is a research area still in progress (De Brouwer, Flanagan, & Spering, 2021). We note that in our study there was no requirement of cognitive inhibition, manipulation of task demands, or distinct overt and covert mechanisms, but we rather relied on the assumption that participants clicked where they looked, and future studies should examine the exact link between manual movements and attention more closely. However, together with the ongoing technological advances in research methods, only designing a study that utilises both motor types, perhaps with an embodied robot, will be necessary to examine these questions in the current context of social gaze and hand movements. As a proof-of-concept, the applicability of similar paradigms is promising so far, when eye tracking may not be available and/or when experiments are conducted online.

The role of individual differences as covariate remains unclear

Additionally, we found that attentional selection was not affected by autistic traits, as it was not the case in Willemsse et al. (2018), or Willemsse & Wykowska (2019). Similarly, whether the intentional stance is affected by joint attention is unclear. It appeared that responding to a robot who typically followed the participant's gaze attenuated intentional attribution, whereas it remained equal when the robot typically looked at the other object, but this trend did not reach significance. We did not find this effect Willemsse and Wykowska (2019), but therein we only presented the ISQ after the task. Having a t0/baseline measurement in the pre-test in the current study, allowed us to observe how establishing joint attention with humanoid agents drives the

sort of attributions we make toward them. Whilst being cognizant that this interaction was limited to an image on a screen, if we were to interpret this, perhaps the attribution of intentional stance lessens with a ‘compliant’ robot, as this is what we expect from designed artefacts, and we are thus more likely to make mechanistic attributions rather than intentional ones, but this effect needs further confirmation in future studies, as our post-hoc comparisons did not reach significance.

Modelling: higher autistic traits indicate faster overall learning rate

To thoroughly examine the inconclusive effects of individual differences such as autistic traits and the likelihood of adopting intentional stance, we looked at how these characteristics affected the task performance over time using growth curve modelling. The current study shows differences in performance over time depending on how much participants exhibit autistic traits. In general, the findings illustrate that those with higher autistic traits tended to perform slower in the initial trials, and then compensated with a higher learning rate (i.e., a steeper slope). However, this effect was not related to the robot’s disposition (i.e., joint vs. disjoint), nor its respective gaze behaviour (i.e., followed vs. unfollowed). This effect could indicate that those who exhibited more autistic traits were more efficient in statistically learning the task compared to those who exhibited less autistic traits. Similar to the overarching discussion posited above, one question that remains unanswered is whether the effect of exhibiting autistic traits on performance over time is evident in manual responses as well as in social eye movements. Furthermore, shortly after carrying out this study, it has been suggested that whilst the AQ-10 might be a useful diagnostic tool, it serves poorly as a non-clinical measure (Taylor, Livingston, Clutterbuck, & Shah, 2020). With a narrower range than the standard AQ, it is well possible that

the heterogeneity and subtle effects of autistic traits were not sufficiently captured by our current design.

Modelling: high intentional stance slows down responses when the disjoint robot follows

More interestingly, individual differences that are related to adopting the intentional stance towards robots influenced participants' performance over time as it related to the robots' disposition (i.e., joint: whether it mostly followed the participants' gaze or disjoint: mostly did not follow the participants' gaze) and whether the robot followed or did not follow the participant's gaze. The findings showed that generally, participants were slower as the experiment progressed, which could be due to experimental fatigue or potentially confounded by regression to the mean. However, those who were more likely to adopt the intentional stance towards the robot became relatively slower over time for the robot with the disjoint disposition when it followed their object selection compared to those who were less likely to adopt the intentional stance towards the robot. This could be due to the fact that those who were more likely to adopt the intentional stance towards robots and did the task with the disjoint robot are more sensitive to violations of expectations (i.e., they are expecting the robot to look in the opposite direction of their choice) and as such, it would influence their responses. Those who were less likely to adopt the intentional stance towards the robot, however, showed no differences for when the robot followed or did not follow their choice as their performance over time was not driven by the social factors of the paradigm.

On the other hand, the attribution of intentional stance to robotic agents was an unsuccessful *predictor* of the learning rate, as we found no effect of the post-ISQ in the current study. Interestingly however, we are able to compare it with a baseline measurement, as we presented ISQ items before the behavioural task as. Here, we found that participants sped up

over time, but this is not contingent on ISQ baseline scores. Coupling this with our results above, using ISQ as a dependent variable rather than a predictor, indicates that establishing joint attention may temporarily shape the attribution of intentional stance toward agents, but as an aspect of trait personality it is less predictive of how people choose to interact.

Conclusion

In conclusion, non-clinical individual differences in autistic traits, as well as differences in tendency to adopt the intentional stance, remain inconclusive as underpinning choice-leading behaviour. On the one hand, this could indicate that the reward systems related to initiating joint attention are innate and/or reflexive. On the other hand, this could be explained by methodological issues which should be teased apart in future studies. However, we found that the overlap between eye movements and hand movements may be demonstrated in joint attention tasks as well, as indicated by the fact that the current findings related to manual responses mimic gaze behaviour in previous eyetracking paradigms.

References

- Allison, C., Auyeung, B., & Baron-Cohen, S. (2012). Toward brief “red flags” for autism screening: the short autism spectrum quotient and the short quantitative checklist in 1,000 cases and 3,000 controls. *Journal of the American Academy of Child & Adolescent Psychiatry*, 51(2), 202-212.
- Baldauf, D., Rolf, M., Deubel, H. (2006). Deployment of visual attention before sequences of goal-directed hand movements. *Vision research*, 4355-4374.
- Bangerter, A. (2004). Using pointing and describing to achieve joint focus of attention in dialogue. *Psychological science*, 15(6), 415-419.

- Baron-Cohen, S., Wheelwright, S., Skinner, R., Martin, J., & Clubley, E. (2001). The autism-spectrum quotient (AQ): Evidence from Asperger syndrome/high-functioning autism, males and females, scientists and mathematicians. *Journal of Autism and Developmental Disorders*, 31, 5–17.
- Bayliss, A. P., Murphy, E., Naughtin, C. K., Kritikos, A., Schilbach, L., & Becker, S. I. (2013). “Gaze leading”: Initiating simulated joint attention influences eye movements and choice behavior. *Journal of Experimental Psychology: General*, 142(1), 76.
- Bertrams, A., & Schlegel, K. (2020). Speeded reasoning moderates the inverse relationship between autistic traits and emotion recognition. *Autism*, 24(8), 2304-2309.
- Booth, T., Murray, A. L., McKenzie, K., Kuenssberg, R., O'Donnell, M., & Burnett, H. (2013). Brief report: An evaluation of the AQ-10 as a brief screening instrument for ASD in adults. *Journal of Autism and Developmental Disorders*, 43(12), 2997-3000.
- Bossi, F., Willemse, C., Cavazza, J., Marchesi, S., Murino, V., & Wykowska, A. (2020). The human brain reveals resting state activity patterns that are predictive of biases in attitudes toward robots. *Science robotics*, 5(46).
- Corbetta, D., Wiener, R. F., Thurman, S. L., & McMahon, E. (2018). The embodied origins of infant reaching: Implications for the emergence of eye-hand coordination. *Kinesiology Review*, 7(1), 10-17.
- Curran, P. J., Obeidat, K., & Losardo, D. (2010). Twelve frequently asked questions about growth curve modeling. *Journal of cognition and development*, 11(2), 121-136.
- Dalmaso, M., Edwards, S. G., & Bayliss, A. P. (2016). Re-encountering individuals who previously engaged in joint gaze modulates subsequent gaze cueing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(2), 271.

- De Brouwer, A. J., Flanagan, J. R., & Spering, M. (2021). Functional use of eye movements for an acting system. *Trends in Cognitive Sciences*, 25(3), 252-263.
- Dennett, D. C. (1987). *The Intentional Stance*. MIT Press.
- Deubel, H., and Schneider, W. X. (1996). Saccade target selection and object recognition: evidence for a common attentional mechanism. *Vision Research*, 36, 1827–1837.
- Emery, N. J. (2000). The eyes have it: the neuroethology, function and evolution of social gaze. *Neuroscience & biobehavioral reviews*, 24(6), 581-604.
- Friesen, C. K., & Kingstone, A. (1998). The eyes have it! Reflexive orienting is triggered by a nonpredictive gaze. *Psychonomic bulletin & review*, 5(3), 490-495.
- Frischen, A., Bayliss, A. P., & Tipper, S. P. (2007). Gaze cueing of attention: visual attention, social cognition, and individual differences. *Psychological bulletin*, 133(4), 694.
- Georgopoulos, A. P. (1990, January). Neural coding of the direction of reaching and a comparison with saccadic eye movements. In *Cold Spring Harbor symposia on quantitative biology* (Vol. 55, pp. 849-859). Cold Spring Harbor Laboratory Press.
- Gollwitzer, A., Martel, C., McPartland, J. C., & Bargh, J. A. (2019). Autism spectrum traits predict higher social psychological skill. *Proceedings of the National Academy of Sciences*, 116(39), 19245-19247.
- Grynszpan, O., Martin, J. C., & Fossati, P. (2017). Gaze leading is associated with liking. *Acta Psychologica*, 173, 66-72.
- Land, M. F., & Hayhoe, M. (2001). In what ways do eye movements contribute to everyday activities?. *Vision research*, 41(25-26), 3559-3565.
- Leiner, D. J. (2019). SoSci Survey (Version 3.1.06) [Computer software]. Available at <https://www.soscisurvey.de>

- Leung, H. C., & Cai, W. (2007). Common and differential ventrolateral prefrontal activity during inhibition of hand and eye movements. *Journal of Neuroscience*, 27(37), 9893-9900.
- Leung, E. H., & Rheingold, H. L. (1981). Development of pointing as a social gesture. *Developmental psychology*, 17(2), 215.
- Lundin, A., Kosidou, K., & Dalman, C. (2019). Measuring autism traits in the adult general population with the brief autism-spectrum quotient, AQ-10: Findings from the Stockholm public health cohort. *Journal of autism and developmental disorders*, 49(2), 773-780.
- Marchesi, S., Ghiglino, D., Ciardo, F., Perez-Osorio, J., Baykara, E., & Wykowska, A. (2019). Do we adopt the intentional stance toward humanoid robots?. *Frontiers in psychology*, 10, 450.
- Metta, G., Natale, L., Nori, F., Sandini, G., Vernon, D., Fadiga, L., ... & Montesano, L. (2010). The iCub humanoid robot: An open-systems platform for research in cognitive development. *Neural networks*, 23(8-9), 1125-1134.
- Morgan, E. J., Smith, D. T., & Freeth, M. (2021). Gaze cueing, mental States, and the effect of autistic traits. *Attention, Perception, & Psychophysics*, 1-9.
- Peirce, J. W., Gray, J. R., Simpson, S., MacAskill, M. R., Höchenberger, R., Sogo, H., Kastman, E., Lindeløv, J. (2019). PsychoPy2: experiments in behavior made easy. *Behavior Research Methods*. 10.3758/s13428-018-01193-y
- Schilbach, L., Wilms, M., Eickhoff, S. B., Romanzetti, S., Tepest, R., Bente, G., ... & Vogeley, K. (2010). Minds made for sharing: initiating joint attention recruits reward-related neurocircuitry. *Journal of cognitive neuroscience*, 22(12), 2702-2715.
- Sober, S. J., & Sabes, P. N. (2005). Flexible strategies for sensory integration during motor planning. *Nature neuroscience*, 8(4), 490-497.

- Spatola, N., Marchesi, S., & Wykowska, A. (2021, May 28). The Instance Task: how to measure the mentalistic bias in human-robot interaction. <https://doi.org/10.31234/osf.io/b3wtq>
- Stephenson, L. J., Edwards, S. G., Howard, E. E., & Bayliss, A. P. (2018). Eyes that bind us: Gaze leading induces an implicit sense of agency. *Cognition*, 172, 124-133.
- Taylor, E., Livingston, L., Clutterbuck, R., & Shah, P. (2020). Psychometric concerns with the 10-item Autism-Spectrum Quotient (AQ10) as a measure of trait autism in the general population. *Experimental Results*, 1, E3. doi:10.1017/exp.2019.3
- Tipper, S. P., Howard, L. A., & Paul, M. A. (2001). Reaching affects saccade trajectories. *Experimental brain research*, 136(2), 241-249.
- Vallejo, G., Ato, M., Fernández, M. P., & Livacic-Rojas, P. E. (2019). Sample size estimation for heterogeneous growth curve models with attrition. *Behavior Research Methods*, 51(3), 1216-1243.
- Vesia, M., & Crawford, J. D. (2012). Specialization of reach function in human posterior parietal cortex. *Experimental brain research*, 221(1), 1-18.
- Wiese, E., Müller, H. J., & Wykowska, A. (2014, October). Using a gaze-cueing paradigm to examine social cognitive mechanisms of individuals with autism observing robot and human faces. In *International Conference on Social Robotics* (pp. 370-379). Springer, Cham.
- Willemse, C., Marchesi, S., & Wykowska, A. (2018). Robot faces that follow gaze facilitate attentional engagement and increase their likeability. *Frontiers in psychology*, 9, 70.
- Willemse, C., & Wykowska, A. (2019). In natural interaction with embodied robots, we prefer it when they follow our gaze: a gaze-contingent mobile eyetracking study. *Philosophical Transactions of the Royal Society B*, 374(1771), 20180036.

- Wykowska, A., Kajopoulos, J., Ramirez-Amaro, K., & Cheng, G. (2015). Autistic traits and sensitivity to human-like features of robot behavior. *Interaction Studies*, 16(2), 219-248.
- Yu, C., & Smith, L. B. (2017). Hand–eye coordination predicts joint attention. *Child development*, 88(6), 2060-2078.

Author's Preprint