



OPEN

The influence of human agency beliefs on ascribing gaze-signalled communicative intent

Friederike Charlotte Hechler^{1,2}✉, Emmanuele Tidoni^{5,6}, Emily S. Cross³ & Nathan Caruana⁴

Communication with artificial agents, such as virtual characters and social robots, is becoming more prevalent, making it crucial to understand how their behaviours can best support social interaction. Eye gaze is a key communicative behaviour, as it signals attention and intentions. Prior research shows that perceiving an agent as sentient affects how its gaze is interpreted. This study examined how such beliefs affect the interpretation of gaze as a signal of communicative intent. In a semi-interactive online task, 160 participants viewed a virtual agent exhibiting dynamic gaze sequences. Each trial varied whether eye contact occurred and whether the agent looked at the same object twice. Participants judged whether the agent was requesting help or merely inspecting the object. Beliefs about the agent's sentience (human- or AI-controlled) were also manipulated. Results showed that when gaze cues were ambiguous, participants were more likely to ascribe communicative intent if they believed the agent was human-controlled compared to when they believed the agent was AI-controlled. Subjective ratings also indicated a general preference for human-controlled agents. These findings underscore the influence of user expectations on interpreting gaze in artificial agents.

Keywords Eye contact, Communicative intent, Human–AI interaction, Artificial intelligence, Joint attention, Intentional stance

Virtual and artificial agents, including animated characters in virtual reality (VR) and social robots, are becoming increasingly prevalent across various sectors, such as healthcare¹, education², the service industry³, and marketing⁴. Specifically, artificial intelligence (AI) is transforming daily life, enhancing human health, safety, and productivity^{5–8}. Interactions with these agents are increasingly resembling human–human interactions⁹, as many of these agents now display human-like features and multimodal behaviours such as social gaze displays, facial expressions, and gestures^{10,11}. This perceived humanlikeness can lead the user to engage in 'anthropomorphism'—the attribution of human-like traits and intentions^{12–16}. However, people vary in their tendency to engage in anthropomorphism¹⁷, and this variation is further amplified by the differences in degree of humanlikeness exhibited by virtual and artificial agents¹⁸. Human-like agent features (e.g., appearance, behaviours) and user predispositions (e.g., anthropomorphism tendencies, beliefs about agents) can influence human–machine interactions in a bottom-up (agent-driven) and top-down (user-driven) manner respectively^{10,19–24}. This influence can be positive or negative, depending on the context. For instance, while human-like appearance can foster trust and acceptance of *social* robots^{20,21,25–27}, it may negatively impact the perceived reliability and attention allocation for *industrial* robots²⁸. Understanding how the influence of an agent's appearance and behaviour interacts with or is overwritten by our beliefs about agents (e.g., that they are artificially controlled), is critical to informing how we design and position agents in a way that promotes effective communication with, and acceptance by, humans.

A key behavioural feature of human-like artificial agents is social gaze, which also plays a crucial role in *human–human* interactions. Eye gaze in human–human interactions is a particularly rich source of social information²⁹, providing constant and rapid information about others' visual perspectives and intentions^{29–31}. As such, eye gaze information is often used to signal and identify opportunities for 'joint attention'—the coordination of attention and actions with others³². Joint attention plays a crucial role in the development of language and social cognition in children, particularly, in understanding another person's unobservable mental states, such as

¹Macquarie University, Balaclava Rd, Macquarie Park, NSW 2109, Australia. ²Potsdam University, Building 14, Karl-Liebknecht-Straße 24–25, 14476 Potsdam, Germany. ³ETH Zurich, Stampfenbachstrasse 69, 8092 Zurich, Switzerland. ⁴Flinders University Institute for Mental Health and Wellbeing, College of Education, Psychology and Social Work, Flinders University, Sturt Rd, Bedford Park, SA 5042, Australia. ⁵University of Leeds, University Rd, Woodhouse, Leeds LS2 9JT, UK. ⁶University of Hull, Hull, UK. ✉email: fhechler@uni-potsdam.de

intentions, beliefs, desires^{33,34}, and remains essential for collaborative interactions throughout life³⁵. Successfully achieving joint attention activates regions in the social brain network associated with attributing mental states to another person ('mentalising')³⁶.

In *human-machine* interactions, eye gaze also plays a crucial role^{37–40}. Recent studies by Caruana and colleagues have explored how gaze behaviour of artificial agents, including the frequency, contextual sequence, and timing of eye contact within dynamic eye movements, influences both the likelihood and certainty (i.e., speed) with which people interpret an agent as signalling communicative intent⁴¹. In a paradigm by Hechler and colleagues⁴², participants engaged in a semi-interactive task with an on-screen agent that shifted its gaze three times in each trial, inspecting three different objects (see their OSF page for details on the design, data, and analysis: <https://osf.io/6uyhr/>). The task was to decide whether the agent's gaze behaviour signalled a request for one of the objects (signalling communicative intent) or whether the agent was privately inspecting the objects (without communicative intent). The study manipulated across trials whether the agent displayed eye contact and made repeated averted gaze shifts at the same object. Results showed that eye contact and repeated averted gaze shifts to the same object independently and additively increased ascribed communicative intent (henceforth '**Eye Contact**' and '**Repeated Gaze**' effects). Eye contact was the more potent signal of the two, but the strongest effect occurred when both features were present together. While these effects generalised across human- and robot-looking agents, the study did not control for whether participants *believed* that the agent's behaviour was human- or machine-generated.

As users interact with increasingly human-like agents, their belief about an agent's humanness—in particular, the assumption that they have human-like mental states and capabilities (i.e., *intentional stance belief*)²⁴—becomes an important factor in shaping the human–machine interaction⁴³. Adopting an intentional stance, that is, assuming an entity possesses intentional qualities and agency^{44–46}, may be an effective way for non-expert users to interpret and predict the behaviour of artificial agents, given their complexity⁴⁷. For instance, users may perceive a virtual assistant as 'intending' to optimise their schedule instead of trying to understand the underlying algorithms and code. This focus on the assistant's 'intentions' rather than its technical complexity may simplify the interaction. When people adopt an intentional stance, they may perceive artificial agents as more trustworthy⁴⁸ or engaging⁴⁹ and benefit from increased communication quality and collaboration^{24,50}. Specifically, the adoption of an intentional stance may modulate how people respond to gaze cues as they tend to attribute greater relevance to gaze cues when they believe them to be indicative of the intentions of a sentient being^{19,40,51–54}. Gaze-cueing studies could show how believing screen-based virtual agents or physical robots to be human- instead of computer-controlled may modify behavioural strategies for social coordination, including increased degrees to which gaze is followed^{52,53,55}, more positive subjective experiences⁵⁵, and modulated neural processing of gaze shifts that signal the achievement or avoidance of joint attention^{53,56,57}. For instance, Abubshait and Wykowska⁵¹ showed that attention orienting to robot gaze-cues was positively related to how much intention was attributed to the robot. Thus, the perceived physical humanness of an agent and the belief that changes in gaze direction are intentional rather than unintentional may impact the social relevance of gaze and, thus, increase or decrease the likelihood of following it^{58,59}. While previous research has primarily compared the discrete categories human- and computer-controlled agents, it remains unclear whether these effects persist when comparing human- and *AI*-controlled agents. In recent years, significant advancements in *AI* may have elevated public expectations regarding the intelligence and capabilities of *AI*-controlled agents. Consequently, the perceived difference between humans and artificial agents in basic social interactions may be narrowing, with the expectation that these computer systems can now adapt to different contexts and, consequently, increase attributions of intelligence, sentience and consciousness⁶⁰. Research has shown that the tendency to attribute mental states decreases from humans to robots and then to computers⁶¹. Given that *AI* is narrowing the gap between human and artificial agents, it is important to understand where *AI*-controlled agents fit within this hierarchy.

The current study builds on the experimental design of Caruana and colleagues^{41,42} to investigate how agent gaze behaviour and user beliefs about the agent's agency shape the interpretation of eye gaze as communicative. Understanding the role of the intentional stance in ascribing communicative intent is essential for defining the conditions under which people are likely to be most receptive to the social behaviours of artificial agents. In our task, the agent performed short gaze sequences toward three objects. We manipulated in each trial whether the agent displayed eye contact and whether or not the agent repeatedly looked at the same object in the gaze sequence. Participants used their keyboard to decide on each trial whether to 'give' one of the objects to the agent. We anticipated replicating previous findings on the effects of gaze dynamics on ascriptions of communicative intent. Specifically, based on Hechler and colleagues⁴², we expected participants to most likely ascribe communicative intent when the agent made eye contact and repeated averted gaze displays (e.g., object 1, eye contact, object 1), followed by eye contact without gaze repetition (e.g., object 1, eye contact, object 2), followed by gaze repetition without eye contact (e.g., object 1, object 2, object 1), and, finally, no eye contact or gaze repetition (e.g., object 1, object 3, object 2). These predictions are consistent with previous findings showing that communicative intent is most strongly inferred when eye contact precedes an object-directed gaze shift, whereas private intentions are inferred primarily from gaze toward objects alone⁶². These findings highlight eye contact as a particularly powerful ostensive cue for communication and suggest that combining eye contact with additional gaze signals enhances ascriptions of communicativeness. Consistent with this, Senju and colleagues⁶³ showed that six-month-old infants follow gaze only when preceded by ostensive cues such as direct gaze or infant-directed speech, underscoring the critical role of ostension in triggering communicative interpretations. From a theoretical perspective, these observations are in line with the proposal that human communication is supported by a dedicated system of *natural pedagogy*^{64,65}. According to this framework, ostensive signals such as direct gaze serve to indicate that communication is intended, induce referential expectations, and bias observers to interpret subsequent signals as conveying generalisable information. Thus, ostensive gaze cues can

be understood not merely as attention-grabbing, but as part of an evolved communicative system that scaffolds the attribution of communicative intent.

Importantly, evidence from infant gaze-following studies also suggests that communicative attributions can arise from attention-grabbing actions that are not conventionally ostensive. Six-month-old infants have been demonstrated to follow gaze not only after direct eye contact but also after non-ostensive yet salient cues such as shivering or nodding^{66,67}. These findings challenge the assumption that ostensive signals are uniquely required for gaze following and instead support the view that social attention mechanisms broadly facilitate communicative interpretations. In addition, Senju and colleagues⁶⁸ provided converging evidence that nine-month-old infants are sensitive to the referential relation between gaze direction and object location, but only when the gaze shift is preceded by direct eye contact. Their findings underscore that eye contact acts as a communicative signal enabling infants to treat subsequent gaze as referential, and that repeated gaze shifts further strengthen such encoding. Together, these studies align with our prediction that repeated gaze shifts, even in the absence of eye contact, would still elicit intermediate levels of ascribed communicative intent, whereas eye contact combined with gaze shifts would maximise attributions of communicativeness.

We further expected faster reaction times for 'give' than 'no-give' responses in the more communicative conditions, and the inverse effect in the less communicative conditions. This pattern would indicate greater certainty of communication when eye contact and gaze repetition cues were displayed and greater certainty of no communication when both cues were absent. Previous work by Jording and colleagues showed that participants base their interpretation of gaze as communicative or private on the occurrence of eye contact and the location of the subsequent gaze shift⁶² and are slower and less coherent in inferring interactivity from gaze patterns when these are more ambiguous⁶⁹.

Critically, for the current study, we anticipated that the bottom-up effects of gaze dynamics on evaluations of communicative intent would be modulated by top-down beliefs about human agency. To test this hypothesis, we manipulated within-subjects whether participants believed that the eye movements of the artificial agent were modelled on human- or AI-generated data. We expected that perceptual signals of communicative intent (e.g., eye contact) would have greater impact when perceived as reflecting the intentional behaviours of a human agent rather than an AI system.

Results

In this semi-interactive task, participants observed a virtual human-like agent performing short gaze sequences toward one of three objects, embedded in a collaborative block-construction scenario. Participants judged whether the agent's gaze behaviour indicated a communicative request ('give') or a private visual search ('no-give'). Reaction times were recorded to index response certainty. On each trial, two gaze parameters were manipulated *within-subjects*: (1) **Eye Contact**: whether the agent made direct gaze toward the participant during the middle gaze shift and (2) **Repeated Gaze**: whether the agent looked at the same object before and after the middle gaze shift. This 2x2 manipulation produced four distinct gaze conditions: (1) **No Eye Contact + No Repeated Gaze**; (2) **No Eye Contact + Repeated Gaze**; (3) **Eye Contact + No Repeated Gaze**; and (4) **Eye Contact + Repeated Gaze**. We also implemented a belief manipulation that varied participants' assumptions about the agent's control source (**Belief**): in one block, they were told the gaze behaviour was based on human eye-movement data (**Human-Belief**); in the other, it was said to be AI-generated (**AI-Belief**). Block order was counterbalanced across participants. This design allowed us to assess main and interactive effects of *bottom-up* gaze cues (**Eye Contact**, **Repeated Gaze**) and *top-down* beliefs (**Belief**) on the ascription of communicative intent.

Behavioural measures

'Give' frequencies

We examined whether participants' **Belief** influenced the tendency to 'give' (i.e., ascribe communicative intent) across the gaze conditions (**Eye Contact** and **Repeated Gaze**). The final model after model selection accounted for within-subject and between-trial variability with random intercepts for subject, subject:**Eye Contact**, subject:**Repeated Gaze**, subject:**Belief**, subject:**Eye Contact:Repeated Gaze**, subject:**Eye Contact:Belief**, subject:**Repeated Gaze:Belief**, subject:**Eye Contact:Repeated Gaze:Belief**, and trial. Descriptive statistics for both behavioural measures are summarised by condition in Table 1. The random effect components of the model explained a significant amount of variance, as indicated by a marginal R^2 of .31 (which reflects the variance explained by the fixed effects alone) and a conditional R^2 of .83 (which includes the variance explained by both the fixed and random effects).

We replicated, as expected, the significant main effects of **Eye Contact** ($\beta = -1.99$, $SE = 0.17$, $p < .001$) and **Repeated Gaze** ($\beta = -1.31$, $SE = 0.14$, $p < .001$) as well as the significant interaction between **Eye Contact** and **Repeated Gaze** ($\beta = 0.45$, $SE = 0.11$, $p < .001$). Participants were more likely to 'give' when the agent displayed eye contact ($M = 87.36$, $SD = 25.01$) than when it displayed no eye contact ($M = 50.51$, $SD = 39$); and when it repeated averted gaze displays ($M = 79.22$, $SD = 32.62$) than when it displayed gaze shifts to unique locations ($M = 58.65$, $SD = 39.38$). Critically, the likelihood to 'give' was higher when the agent displayed eye contact together with a repeated averted gaze display. Pairwise comparisons revealed significant differences in 'give' frequencies between all four gaze conditions. People were significantly more likely to 'give' when the agent displayed eye contact and repeated averted gaze displays (**Eye Contact+Repeated Gaze**: $M = 97.77$, $SD = 5.45$) than when averted gaze displays were unique (**Eye Contact+No Repeated Gaze**: $M = 76.95$, $SD = 31.73$; Estimate = -3.52 , $SE = 0.39$, 95% CI $[-4.46, -2.58]$, $z = -8.93$, $p < .001$); when the agent displayed eye contact with unique averted gaze displays than no eye contact and repeated gaze (**No Eye Contact+Repeated Gaze**: $M = 60.68$, $SD = 37.58$; Estimate = 1.36 , $SE = 0.4$, 95% CI $[0.29, 2.43]$, $z = -8.93$, $p < .001$); and when the agent displayed no eye contact and repeated gaze than no eye contact with unique gaze displays (**No Eye Contact+No Repeated Gaze**: $M =$

Gaze condition	Belief	Give frequency %	Give RT	No-Give RT
Eye Contact, No Repeated Gaze	AI	76.6 (31.6)	0.83 (0.51)	0.95 (0.54)
	Human	77.31 (31.99)	0.85 (0.51)	0.88 (0.55)
Eye Contact, Repeated Gaze	AI	98.28 (3.9)	0.63 (0.35)	0.96 (0.58)
	Human	97.25 (6.62)	0.64 (0.35)	1.12 (0.64)
No Eye Contact, No Repeated Gaze	AI	41.18 (38.04)	0.84 (0.5)	0.75 (0.42)
	Human	39.52 (37.67)	0.87 (0.5)	0.73 (0.42)
No Eye Contact, Repeated Gaze	AI	59.92 (37.95)	0.81 (0.46)	0.74 (0.46)
	Human	61.44 (37.33)	0.83 (0.5)	0.73 (0.44)

Table 1. Descriptive statistics by Gaze Condition, Belief, and Response. 'Give' frequencies are summarised as the percentage of trials that participants responded by giving a block to the agent. Reaction times are summarised by response. Means and Standard deviations are reported in the format 'M (SD)'

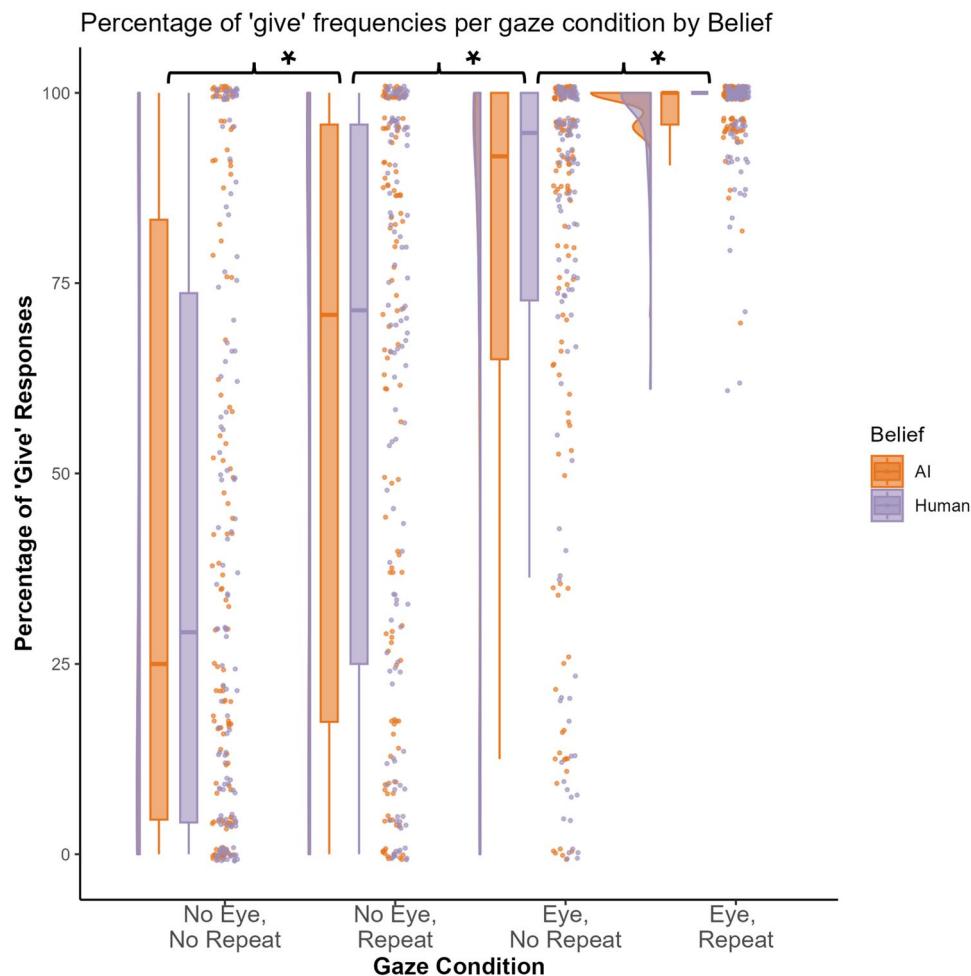


Fig. 1. 'Give' frequencies per gaze condition. AI Belief data is displayed in orange and Human Belief data in violet. The condition labels are abbreviated as follows: 'No Eye Contact, No Repeated Gaze' = 'No Eye, No Repeat', 'No Eye Contact, Repeated Gaze' = 'No Eye, Repeat', 'Eye Contact, No Repeated Gaze' = 'Eye, No Repeat', 'Eye Contact, Repeated Gaze' = 'Eye, Repeat'. Significant differences from pairwise comparisons between each Eye Contact+Repeated Gaze, averaged across Belief groups condition are indicated with an asterisk.

40.35, $SD = 37.79$; Estimate = -1.7 , $SE = 0.32$, 95% CI $[-2.53, -0.88]$, $z = -5.32$, $p < .001$). All other comparisons between conditions were also significant (all $p < .001$; see supplementary material on the OSF project page for a detailed report: <https://osf.io/wcm75>). The pattern of give frequencies across the four gaze conditions and between belief conditions are visualised in Fig. 1 (see the 03_Plots.Rmd output for a visualisation of the two blocks separately).

The primary aim of this study was to investigate whether participants' **Belief** about the intentionality of the agent influenced their tendency to 'give'. While we found no evidence for a significant main effect of **Belief** ($p = .300$), there was evidence for a three-way interaction between **Belief**, **Eye Contact**, and **Repeated Gaze** ($\beta = 0.12$, $SE = 0.04$, $p = .003$), indicating higher 'give' frequencies in the ambiguous conditions when participants believed to observe human- (**No Eye Contact+Repeated Gaze**: $M = 61.44$, $SD = 37.33$; **Eye Contact+No Repeated Gaze**: $M = 77.31$, $SD = 31.99$) rather than AI-modelled behaviour (**No Eye Contact+Repeated Gaze**: $M = 59.92$, $SD = 37.95$; **Eye Contact+No Repeated Gaze**: $M = 76.6$, $SD = 31.6$). However, follow-up pairwise comparisons between **Belief** conditions for each individual gaze condition showed no significant differences in 'give' frequencies (all $ps > .37$). To examine the **Belief x Eye Contact x Repeated Gaze** interaction further, we explored whether the absence of significant pairwise comparisons between **Belief** conditions was due to the averaging of response frequencies across block 1 and block 2 of the experiment; across which the belief order was counterbalanced. To this end, we examined whether there was any evidence for a different pattern of results when examining the block 1 and 2 datasets separately; essentially treating the **Belief** factor as a between-subjects factor, and removing the potential influence of order effects on the **Belief** manipulation. We found that the **Belief x Eye Contact x Repeated Gaze** interaction was neither significant for block 1 ($p = .123$) nor block 2 ($p = .243$).

Reaction times

We also investigated the influence of the above factors on RTs for a **Give** or **No-Give Response**. The frequency analysis reported above evaluated the influence on the *likelihood* of ascribing communicative intent, whereas the RT analysis evaluates *certainty/ambiguity* in the ascription of communicative intent. The final model after model selection included random intercepts for subject, subject:**Give Response**, subject:**Belief**, subject:**Eye Contact:Repeated Gaze**, subject:**Give Response:Eye Contact**, subject:**Give Response:Repeated Gaze**, subject:**Give Response:Belief**, subject:**Give Response:Eye Contact:Repeated Gaze:Belief**, and trial. Descriptive statistics for RTs are summarised by condition in Table 1 and visualised in Fig. 2. The random effect parameters in the model explained a significant amount of variance, as reflected in both the marginal R^2 value (.06) and conditional R^2 value (.45).

Once more, we replicated the findings reported by Hechler and colleagues⁴². Our analysis showed evidence for a **Response x Eye Contact** interaction ($\beta = -0.106$, $SE = 0.01$, $p < .001$; $\eta_p = 0.260$, 95% CI [0.200, 1.000]); and a **Response x Repeated Gaze** interaction ($\beta = -0.05$, $SE = 0.01$, $p < .001$; $\eta_p = 0.076$, 95% CI [0.041, 1.000]). We also found evidence for a **Response x Eye Contact x Repeated Gaze** interaction ($\beta = 0.026$, $SE = 0.01$, $p < .001$; $\eta_p = 0.030$, 95% CI [0.007, 1.000]). This three-way interaction suggested that, in the most communicative condition (i.e., **Eye Contact+Repeated Gaze**) participants were significantly faster to 'give' ($M = 633.33$, $SD = 199.36$) than 'not give' ($M = 1045.06$, $SD = 550.72$; Estimate = 0.48, $SE = 0.06$, 95% CI [0.36, 0.57], $t = 8.32$, $p < .001$); and, in the least communicative condition (i.e., **No Eye Contact+No Repeated Gaze**), participants were significantly faster to 'not give' ($M = 823.51$, $SD = 329.64$) than 'give' ($M = 928.49$, $SD = 364.28$; Estimate = -0.15, $SE = 0.03$, 95% CI [-0.23, -0.05], $t = -4.3$, $p < .001$). For the more ambiguous conditions, differences in RTs between 'give' and 'no-give' responses were smaller. When the agent made no eye contact but repeated the averted gaze display, RTs were not found to differ (**No Eye Contact+Repeated Gaze**: $p = .675$). However, when the agent displayed eye contact but did not repeat the gaze, pairwise comparisons between 'give' and 'no-give' responses revealed significantly shorter 'give' (**Give Response+Eye Contact+No Repeated Gaze**: $M = 897.86$, $SD = 345.09$) than 'not give' response RTs (**No Give Response+Eye Contact+No Repeated Gaze**: $M = 1153.16$, $SD = 523.73$; Estimate = 0.17, $SE = 0.04$, 95% CI [0.07, 0.25], $t = 4.82$, $p < .001$).

Finally, we found evidence for a **Response x Belief x Eye Contact x Repeated Gaze** interaction ($\beta = -0.015$, $SE = 0.01$, $p = .026$; $\eta_p = 0.002$, 95% CI [0.000, 1.000]). Pairwise comparisons of this four-way interaction only confirmed, for both groups, the same **Response x Eye Contact x Repeated Gaze** interaction reported above, since comparisons between **Belief** conditions within each of the gaze conditions were not significant (all $ps > .70$).

Agent perception

We explored whether people formed different expectations and perceptions of the agent stimulus depending on their intentional stance. Participants rated their expectations regarding the agent and the upcoming interaction before each **Belief** block and made similar retrospective ratings on the agent and the interaction after each block. The results of a cumulative link model analysis revealed that, before each **Belief** block, participants had significantly lower ratings of the AI- than the human-modelled agent with respect to their expectation that the agent would be (1) a good communicator ($\beta = 0.96$, $SE = 0.24$, $z = 3.99$, $p < .001$); enjoyable to interact with ($\beta = 0.5$, $SE = 0.23$, $z = 2.14$, $p = .033$); and (3) easy to interact with ($\beta = 0.79$, $SE = 0.24$, $z = 3.37$, $p < .001$). After each block, participants provided the same ratings about the agent under each **Belief** condition. Retrospective judgements did not significantly differ between **Belief** conditions (all $ps > .607$).

However, when participants compared the two agents on a single scale after both blocks, they rated the human as more preferable than the AI agent across several dimensions. Specifically, we conducted Wilcoxon Rank-Sum tests to determine whether participants' preferences for the human agent were significantly greater than the neutral midpoint of the scale. This indicated that participants significantly preferred the human over the AI with respect to its communication ($M = 3.18$, $SD = 1.03$), $V = 6928$, $p = .002$), pleasantness ($M = 3.16$, $SD = 0.78$), $V = 1548$, $p = .002$), and naturalness ($M = 3.42$, $SD = 0.88$), $V = 5629$, $p < .001$).

58 participants indicated, with an overall preference rating above 3, a preference for the agent in the **Human-Belief** condition, 38 participants indicated, with a rating below 3, a preference for the AI agent, and 38 participants indicated, with a rating of 3, no preference. While the variation in these preferences for either agent was not significant ($M = 3.14$, $SD = 1.25$), $V = 9378$, $p = .063$), free-text responses indicated that participants who preferred the human mostly did so because they perceived this agent as more "natural" ($n = 11$), whereas

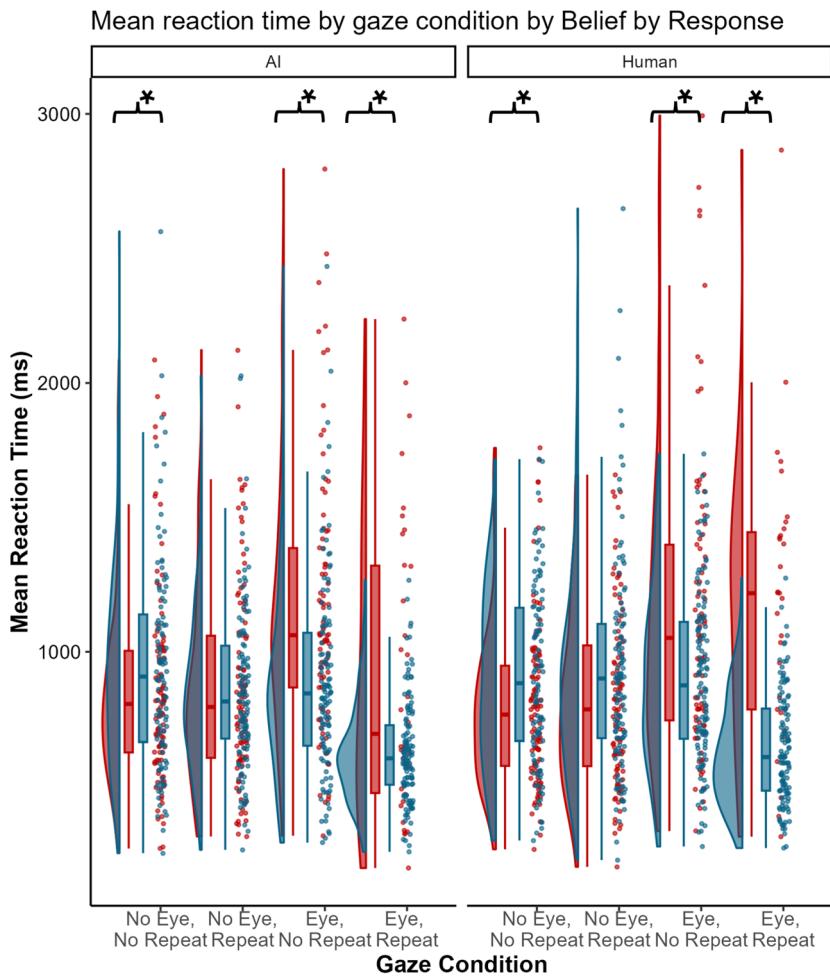


Fig. 2. Reaction times per gaze condition. AI Belief data is displayed on the left and Human Belief data on the right. ‘No-give’ response data is displayed in red and ‘give’ response data in blue. The condition labels are abbreviated as follows: ‘No Eye Contact, No Repeated Gaze’ = ‘No Eye, No Repeat’, ‘No Eye Contact, Repeated Gaze’ = ‘No Eye, Repeat’, ‘Eye Contact, No Repeated Gaze’ = ‘Eye, No Repeat’, ‘Eye Contact, Repeated Gaze’ = ‘Eye, Repeat’. Significant differences from pairwise comparisons between Response conditions for each Eye Contact+Repeated Gaze condition, separately for each Group are indicated with an asterisk.

participants who preferred the AI mostly did so because they perceived this agent as a clearer communicator ($n = 12$), as can be seen in the word clouds in Fig. 3. A shared prominent reason in both groups was that the respective agent was “easier” (Human: ($n = 14$); AI: $n = 8$).

Discussion

The current study examined how beliefs about an artificial agent’s intentionality shape the evaluation of gaze behaviour as signals of communicative intent—an essential process for recognising and responding to communication opportunities (e.g., joint attention). As expected, we replicated the pattern observed by Hechler and colleagues⁴², with individual and cumulative effects of **Eye Contact** and **Repeated Gaze** on the ascription of communicative intent. Such gaze dynamics need not uniquely signal communicativeness and could also reflect personal interest or uncertainty; however, in our task they biased participants toward ascribing communicative intent. Importantly, consistent with our previous work, we do not argue that repeated gaze alone is communicative. Rather, we suggest that it likely functions as a cue to relevance, which, when combined with the communicative potential of eye contact, promotes the impression of an intentional and relevant communicative gaze behaviour. Participants were more likely to ‘give’ when the agent displayed eye contact compared to when it did not, and when it displayed repeated averted gaze shifts to the same location than to unique locations. Critically, ascriptions of communicative intent were strongest—in terms of response frequency and speed—when eye contact and gaze repetition cues were presented together. This replication aligns with evidence from developmental psychology research showing that eye contact is a critical ostensive signal for eliciting gaze following and communicative interpretations in infancy^{63,68}, and with studies demonstrating that non-ostensive but attention-grabbing cues (e.g., shivering, nodding) can also trigger gaze following^{66,67}. Our RT data further supports this conclusion. We found that the absence of eye contact or repeated averted

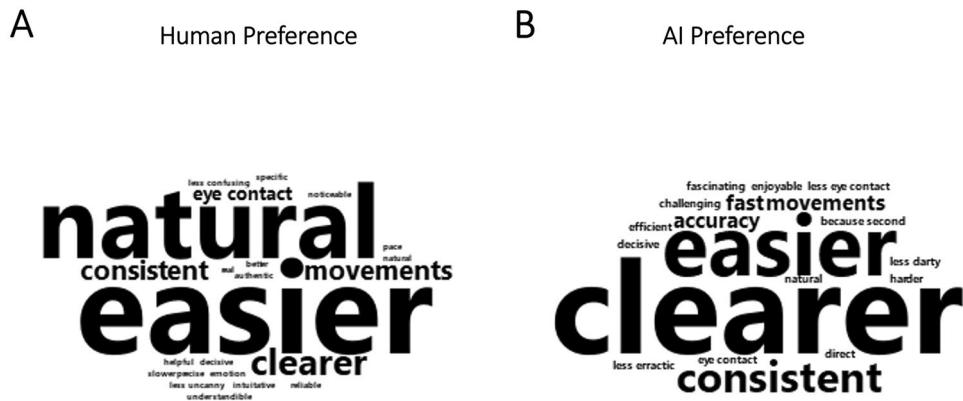


Fig. 3. Free-text responses to the question ‘Which partner did you prefer? Why?’ Responses of participants who preferred the human are shown in **(A)** and responses of participants who preferred the AI in **(B)**. Phrases that appeared more frequently are displayed larger.

gaze displays independently increased the certainty to 'not give' and their presence increased the certainty to 'give'. At the same time, conditions that involved partial signals of communicative intent (e.g., repeated averted gaze displays without eye contact) led to greater variability in 'give' responses and less differentiation in RTs to 'give' or 'not give', suggesting greater ambiguity in the evaluation of communicative intent. This interpretation also resonates with findings by Jording and colleagues⁶⁹, where judgements of interactivity and corresponding reaction times varied systematically with the ambiguity of a virtual agent's gaze signals. More ambiguous states (e.g., introspective gaze without any obvious attentional focus in the environment) elicited longer reaction times and less consistent judgements of interactivity, whereas clearer communicative signals (e.g., initiating joint attention) led to shorter reaction times and higher interactivity ratings. However, like Hechler and colleagues⁴², we also found faster 'give' than 'no-give' responses in the **Eye Contact + No Repeated Gaze** condition. In the original study, this unexpected pattern was mainly driven by heterogeneity induced by autistic participants, who constituted half of the sample⁴². However, given that the pattern was replicated in the current study, it may also reflect how eye contact and repeated averted gaze signal joint attention opportunities in different ways.

One possible explanation for different information provided by eye contact and gaze repetition, is that eye contact may directly signal communication *readiness* or *intention*, while repeated gaze may signal the *relevance* of one's focus of attention and provide context to the communication content where joint attention is being signalled. Such an interpretation is in line with the 'Relevance Theory' proposed by Sperber and Wilson⁷⁰, which posits that communication relies on the principle of relevance, where the signaller aims to be as relevant as possible, and the receiver seeks the most relevant meaning in order to minimise the cognitive costs of processing communicative signals. Within this framework, the concept of ostensive-inferential communication proposes that the signaller provides an ostensive gesture (e.g., eye gaze) that indicates an intention to convey a message (ostension). The receiver must then interpret this stimulus by inferring the communicator's intention based on the context and/or their own knowledge. In the current study, eye contact serves as a clear ostensive signal, indicating the communicator's intention to communicate, leading to increased certainty in the decision to 'give'. Repeated averted gaze shifts, on the other hand, provide contextual information that influences the interpretation of the communicator's intent. When these repeated gaze shifts occur without eye contact, this creates an ambiguous context that demands additional cognitive processing to infer the intention behind the signal. This would explain why we see more variability in RTs when deciding whether to 'give' when repeated gaze is present but eye contact is not. In summary, eye contact acts as a strong ostensive signal, signalling *communicativeness*, while repeated averted gaze shifts may provide contextual cues that can either clarify or obscure the intended message by signalling *relevance*. This distinction is important, as people often follow gaze cues even when the gazer is not intentionally guiding them⁷¹. Given that joint attention is sometimes achieved when the initiator does not intend to communicate their focus of attention (e.g., looking at the target in a game of 'eye spy' or looking at the clock when bored in conversation), it would also be valuable to examine evaluations of gaze as signalling intentional and unintentional communication and joint attention opportunities. This would require a modified task that does not explicitly demand evaluations of whether an agent is intending to communicate, but rather whether the agent is attending to something of interest or relevance. Interactive eye tracking studies of joint attention, similar to those conducted by Caruana and colleagues^{38,72} may also be adapted to examine how eye contact and gaze repetition influence spontaneous joint attention in unconstrained interactions, thus removing the need for explicit judgements altogether. This line of research would shed light on how artificial agents can strategically use gaze behaviour to facilitate intuitive social coordination with humans.

The primary focus of the current study was to determine how beliefs about the intentionality of the agent (AI vs. human) influenced the evaluation of its gaze behaviour. We found evidence for a significant **Belief x Eye Contact x Repeated Gaze** interaction, with higher 'give' frequencies in the ambiguous conditions when participants believed to observe human- rather than AI-modelled behaviour. However, we found no evidence for specific **Belief** effects within each gaze condition when examined using follow-up pairwise comparisons. This may be explained by similar *implicit* beliefs about the agent's intentionality being held across the two **Belief** blocks.

While participants stated after the experiment to have explicitly believed that the behaviour they observed was human- or AI-modelled respectively, they might have implicitly adopted an intentional stance towards both the human and the AI agent. This aligns with the notion that the intentional stance is the default for understanding artificial agents⁵⁰ and it might reflect how the gap between perceptions of human and AI agents is decreasing as (1) AI technologies become more advanced, and (2) people become more exposed to these AI technologies. For instance, a longitudinal study on public opinion and human–AI interaction conducted in the US reported that, while only a small minority of participants perceived current AIs as sentient, a significant majority believed AI could achieve sentience in the future or were uncertain about its potential⁷³. The authors concluded that perceptions of mind can enhance, diminish, or add complexity to well-established human-computer interaction dynamics. Together with the current findings, this suggests that, despite explicit beliefs about the nature of the agent, implicit assumptions about intentionality may be influencing participants' responses, highlighting the need for further research into the cognitive processes underlying human–AI interactions.

It is also possible that the observation of gaze dynamics—and in particular the establishment of eye contact—may have contributed to the implicit adoption of an intentional stance across both **Belief** blocks, even in the condition where people explicitly believed the agent was AI-modelled. Low-level perceptual gaze cues may have automatically engaged mentalising mechanisms, leading participants to infer intentions. Such an interpretation is consistent with theories and associated evidence that suggest that observing or experiencing eye contact automatically activates mentalising mechanisms in the brain, such as the Communicative Intention Detector⁷⁴, Fast-Track Modulator⁷⁵, and Watching Eyes account^{76–78}. This interpretation is also in line with the Intentional Stance Model put forward by Wykowska and colleagues⁵³, which proposes a bidirectional relationship between bottom-up sensory input and top-down cognitive processes. Gaze perception can recruit brain areas associated with social cognition (e.g., medial prefrontal cortex, temporoparietal junction), which can, in turn, influence the responsiveness of regions involved in gaze processing (e.g., superior temporal sulcus), and vice versa.

Finally, cognitive adjustments did not only influence the way participants interpreted gaze cues but also shaped subjective experiences of the interaction. Participants rated communication with the 'human' agent only *before* but not *after* the respective interaction as easier and more enjoyable than with the 'AI' agent. The lack of a difference in retrospective ratings may be due to the behaviour of the agent that was the same across blocks. Moreover, the **Belief** manipulation also influenced subjective experiences. Participants indicated that they significantly preferred the human partner with respect to their communication, pleasantness, and naturalness. This finding is in line with findings from Hechler and colleagues⁷⁹. The authors found that believing an agent in VR-based social interactions to be human- instead of AI-controlled positively affected the subjective perception of both the agent and the interaction, even though this did not influence behavioural gaze use strategies. This is also consistent with earlier studies, which found human agent beliefs to result in more positive subjective experiences than artificial agent beliefs in studies that implemented more discreet belief manipulations (i.e., human- vs. computer-controlled)^{55–57}. In the current study, there were not only **Belief** effects on subjective ratings but also differences in how participants explained their preferences for interacting with the "human–" and "AI- generated" agents. These differences, at least in communication-style preferences, may influence how **Belief** effects shape subjective experiences. Participants who, overall, preferred the human, explained that this was because the agent seemed more "natural". Those participants who preferred the AI mostly explained that this was because the agent seemed to be a clearer communicator. Given that these perceptions could not be based on perceptual inputs—as the agent's behaviour in both **Belief** conditions was the same—this could reflect the role that knowledge, expectations, and past experiences play on subjective experiences during human–AI interactions. Current AI systems vary widely in their capabilities, including computer vision, machine learning, and natural language processing⁸⁰. Likewise, people vary in their experience with and knowledge about these systems⁸¹, which may impact how people approach artificial agents⁸². Future work should further evaluate how knowledge about, experience with, and attitudes towards AI shape the response to gaze cues in human–AI interactions.

Limitations

The current study provides insights into how user beliefs about the intentionality of artificial agents shape the perception of gaze dynamics in *virtual* interactions. However, several limitations should be considered before generalising these findings to contexts involving physically embodied artificial agents (e.g., robots). First, while the virtual agent allowed for precise control over gaze cues and experimental conditions, it lacked multimodal cues that are typically present in real-world human–robot interactions, such as speech, facial expressions, and haptic feedback⁸³. Thus, our scenario may not fully capture the complexities of more naturalistic interactions. For instance, it is possible that beliefs about human agency will have greater influence on other modalities of communication (e.g., facial expression interpretation, body posture or speech). Second, physical robots may elicit more positive social responses than virtual agents^{48,84,85}. A review by Li and colleagues⁸⁶ supports this pattern, although some findings suggest that social, rather than physical, presence is more critical for interaction⁸⁷. As such, future work is needed to examine the extent to which the behavioural effects observed in the current study generalise to interactions with embodied artificial agents (i.e., physical robots).

In addition, we did not collect detailed demographic information such as participants' educational background or cultural context, nor their broader experience with AI or robotics. These factors may shape how people approach artificial agents^{81,88,89,89–92} and thus limit the generalisability of our findings.

Conclusion

Our study used a semi-interactive paradigm to investigate how beliefs about the human agency of an artificial agent impacts the evaluation of communicative intent from gaze behaviour. We directly tested whether the frequency of 'give' responses to gaze sequences was dependent on whether the agent's behaviour was believed

to be human- or AI-modelled. Our findings confirmed that eye contact plays a fundamental role in signalling communicative intent. Repeated gaze shifts also influenced evaluations of communicativeness. However, this may be a result of the contextual information repeated gaze cues provide regarding the relevance of gaze behaviour, rather than signalling communicative intent per se. The presence of both eye contact and repeated gaze maximised the likelihood of ascribing communicative intent. Importantly, we found numerical evidence that participants' belief about the gaze behaviour being human- rather than AI-generated may increase the ascription of communicativeness when gaze cues were ambiguous—because they comprised mixed gaze signals (e.g., repeated gaze without eye contact)—suggesting that top-down agency beliefs inform evaluations in social conditions of uncertainty. Consistent with earlier findings, we also found beliefs to have robust influences on subjective social experiences, again highlighting how expectations of artificial agents are likely to shape interactions with artificial agents. Our findings have implications for the design of AI-driven communication systems, emphasising the need to consider both the behavioural cues of the agent and the evolving expectations that human users have about them.

Methods

The current study builds upon a series of recent preregistered investigations that examined how the characteristics of perceived eye contact impact the interpretation of communicative intent. For a comprehensive theoretical rationale behind this research direction and an in-depth explanation of the original experimental design, please refer to the previous preregistration document available at: <https://osf.io/w68ut/>.

Participants

We recruited 170 participants in total. However, following participant exclusion, based on our pre-registered protocol, we entered a data set comprising 125 participants into our final analyses ($M_{age} = 40$; 61 identified as female; see Pre-Processing Data for a detailed summary of participant exclusion). The majority of people were born (31.62%) and resided in the United Kingdom at the time of completing the study (47.44%), with 43.59% identifying as white (see the supplementary material on the Open Science Framework project page (<https://osf.io/wcm75> for detailed information on country of birth and residence, ethnicity, and student and employment status). Data was collected using the online data collection platform Prolific. Participants had to be at least 18 years old, be fluent in English, reside in Western English-speaking countries (United Kingdom, United States, Ireland, Australia, New Zealand, Canada), report no language-related disorders, and have a minimum Prolific approval rating of 95%. We compensated participants with £9.50 per hour for an approximate duration of 31 minutes. All participants provided written informed consent prior to completing the study, following the procedures approved by the Flinders University Human Research Ethics Committee (6804). All experiments were performed in accordance with relevant guidelines and regulations, in particular, with the Declaration of Helsinki (2013 revision).

Experimental task and stimuli

We framed the task within a 'collaborative' context, instructing participants to 'help' an agent complete the construction of an off-screen block model. Participants were told that on each trial the agent was required to select a block to complete the model. In some trials, the block objects would be accessible to the agent, while in others the agent required the participant to 'give' them the required block. Participants observed the agent's behaviour and decided whether to 'give' a block (i.e., because they interpreted the agent's gaze behaviour as a communicative request) or not (i.e., because they interpreted the agent's gaze behaviour as a private visual search) via a keyboard response. Full task instructions and experimental task code can be found on the corresponding OSF project page (<https://osf.io/wcm75>).

Participants viewed a female agent directing the gaze towards one of three objects. We decided to use the same avatar throughout the study because previous work showed no evidence for avatar gender effects⁷⁹ and we wanted to minimise adding additional factors in our within-subjects design. The avatar was created using Adobe Fuse CC (Beta Version 2014.3.14; San Jose, CA, USA) and was designed to be moderately realistic, with a slightly stylised appearance, and to appear ethnically ambiguous. This perception was validated in a separate sample of 47 participants ($M_{age} = 37.28$, $SD = 10.33$; 31 identified as female)³⁹. To simulate dynamic gaze behaviour, participants were presented with rapidly sequenced static images of the avatar, creating the impression of apparent motion and smooth gaze shifts. The second (middle) gaze shift was presented with two different counterbalanced durations, a short one ranging from 400–800 ms, and a long one ranging from 1100–1500 ms (see Fig. 4 A). These different durations were included, and controlled for, to increase the variability and ecological validity of gaze behaviour, and to enable future exploratory analyses on the potential role of gaze duration on ascriptions of communicative intent (not a focus of this current study). The avatar maintained a neutral facial expression to avoid introducing emotional cues or possible uncanny valley effects⁹³.

In each trial, we manipulated within-subjects whether the avatar made eye contact with the participant or not in the second gaze shift (**Eye Contact**) and whether or not the avatar repeatedly looked at the same object before and after the second gaze shift (**Repeated Gaze**^{41,42}). These two variables resulted in the following four different types of gaze stimuli (see also Fig. 4B):

1. **No Eye Contact+No Repeated Gaze** Three gaze shifts with no eye contact (i.e. the agent looks in different locations).
2. **No Eye Contact+Repeated Gaze** Averted gaze in between two other gaze shifts both directed at the same location (i.e. as condition 1 but the agent gazes at the same location in the first and third gaze shift).
3. **Eye Contact+No Repeated Gaze** Eye contact in between gaze shifts at different locations (i.e. as condition 1 but the agent makes eye contact in the second gaze shift).

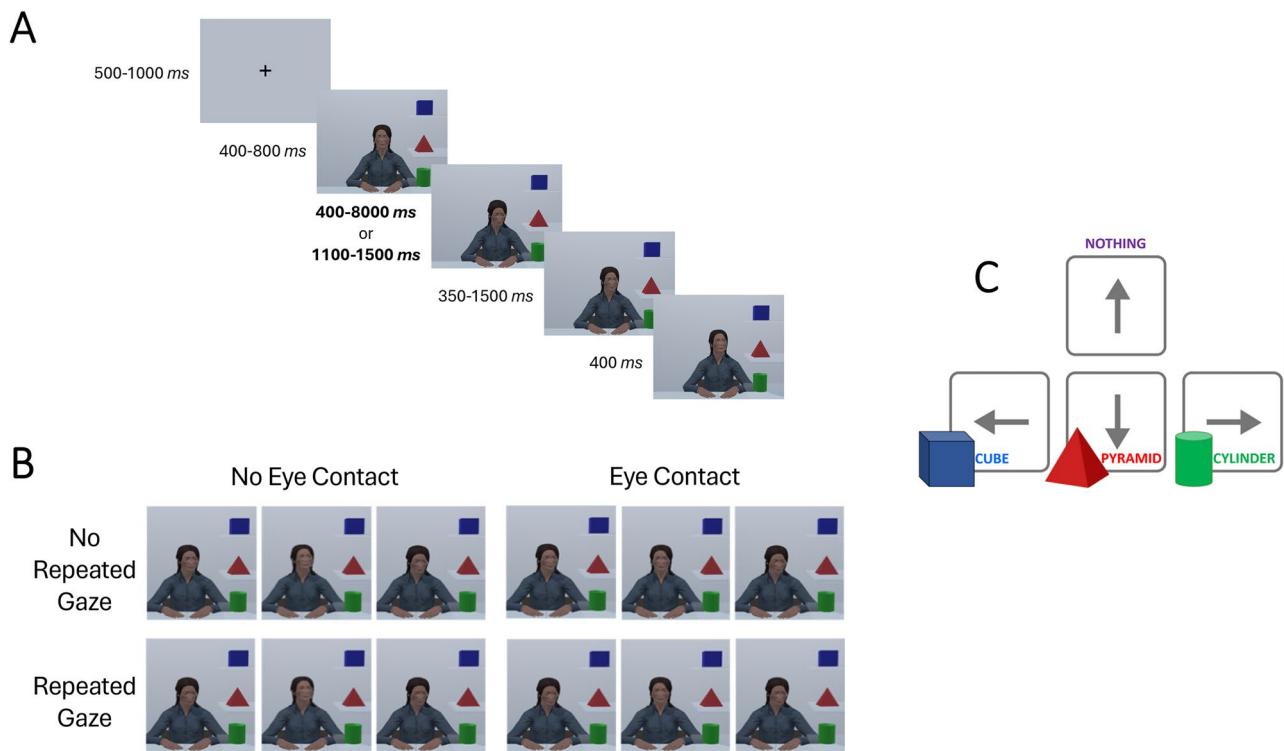


Fig. 4. Trial sequence examples. In one block, participants completed the task assuming the data was modelled on human behaviour. In a second block, participants were led to believe they were seeing an agent reflecting AI-modelled behaviour. The blocks appeared in counterbalanced order across participants. Gaze duration of the second gaze shift, which displayed either eye contact or averted gaze depending on the gaze condition, summarised in (B), was counterbalanced as shown in (A). In each trial, participants decided whether to ‘give’ the agent one of the three blocks, or nothing at all, using the arrow keys on a standard keyboard (C).

4. **Eye Contact+Repeated Gaze** Eye contact, a gaze shift, and eye contact again (i.e. as condition 3 but the agent gazes at the same location in the first and third gaze shifts).

At the conclusion of each trial, participants decided whether to ‘give’ one of the objects via a keyboard response or to ‘not give’. At this stage, we presented an image of the response-key mapping as a visual reminder (see Fig. 4C).

Procedure

The experiment was programmed in PsychoPy (v 2023.2.3) and deployed online via the Pavlovia platform. Participants accessed the task remotely on their own desktop or laptop computers. No specific restrictions or instructions were given regarding the testing environment (e.g., quiet surroundings). Prior to beginning the experimental task, participants were given the chance to practice the response-key mapping over the course of eight practice trials (two for each response key). During these practice trials, participants saw only the key-response mapping accompanied by a text prompt indicating the required response (e.g., “give nothing” or “give cylinder”). We provided feedback after each practice trial. The aim of this practice session was to minimise cognitive load during the task associated with response key mapping.

Participants completed the experimental task in two blocks. In one block, they were informed that the agent’s behaviour was modelled on real recordings of human eye movements. In the other block, participants were informed that the agent’s behaviour was generated by an AI system. Block order was counterbalanced across participants (**Human-Belief** vs. **AI-Belief**). Both of these blocks consisted of the same trial composition, with trials internally counterbalanced regarding the variation and sequence of gaze directions within each condition (24 trials for each experimental condition). Each block was divided into two sessions allowing participants to take two self-paced breaks. Participants completed 96 trials for each block (three sessions x 32 trials), and 192 trials in total (96 trials per block x two agents).

Before each of the two blocks, participants were informed whether their partner’s movements were “AI-modelled” or “human-modelled”. To check that participants attended to this belief manipulation, they were immediately prompted to indicate whether the agent in the upcoming block was human- or AI-modelled. Then, participants responded to questions about their expectations and saw another summary of the task instructions before the behavioural task began.

After completing a block, participants responded to questions on how they perceived their partner. The final block was followed by questions on task experience and partner perception. At the conclusion of the

Gender	Number of participants		Age		CATI total score		ATTARI total score	
	n	%	Mean	SD	Mean	SD	Mean	SD
Female	64	51.2	41	13	106.34	28.97	166.36	39.56
Male	61	48.8	38	12	102.58	29.36	171.91	42.08
Total	125	100.0	40	1	104.46	0.28	169.14	1.78

Table 2. Exposure groups.

entire experimental task, participants completed a self-report measure of autistic traits⁹⁴ and the Attitudes Towards Artificial Intelligence Scale (ATTARI-12)⁹⁵. Finally, participants were debriefed and indicated whether they accepted the instruction that they completed the task with agents displaying human- and AI-modelled movements.

Measures

Behavioural measures We collapsed responses across the four buttons into two **Response** categories (**Give**, **No-Give**). We calculated the proportion of trials in each condition with a **Give Response** to index the tendency of participants to interpret the agent's behaviour as communicative (i.e., making a request). We also defined reaction times (RTs) as the time interval between the end of the trial and the participant's button press. RTs were calculated for both **Give** and **No-Give Responses**, which we compared as an additional measure of response certainty/ambiguity within and across conditions. Specifically, shorter RTs indicated more certainty in the decision.

Agent perception We further explored whether top-down beliefs of the agent's agency influenced subjective experiences during the task and evaluations of the agent. Previous work found these experiences to be influenced by a similar **Belief** manipulation⁷⁹. To this end, participants completed throughout the experiment subjective rating and free-text impression questionnaires on partner perception and task experience created for this study (see supplementary material on the Open Science Framework project page (<https://osf.io/wcm75>)). Specifically, before each **Belief** block, participants rated on 5-point Likert scale ("Not at all" - "Extremely") their expectations regarding the agent's communication clarity and enjoyment and ease of the interaction. After each **Belief** block, participants provided the same ratings retrospectively. Upon completing the entire experimental task, participants rated their preference for the AI or human on 5-point Likert scale ("AI was the best" - "Human was the best") in terms of communication, pleasantness, and naturalness. Additionally, participants noted the most notable difference between the two agents in a free-text response and indicated their overall partner preference on a 5-point Likert scale, providing additional information in another free-text response. Finally, participants judged whether they accepted the instruction about the agents displaying human- and AI-modelled movements on a 10-point Likert scale and in a free-text response.

Autistic traits We also assessed participants' autistic traits with the Comprehensive Autistic Traits Inventory (CATI)⁹⁴ to characterise our sample. Autistic people often express difficulties perceiving or evaluating gaze information. While Hechler and colleagues found no evidence for impairments in autistic people in this task⁴², they reported, nevertheless, subtle differences in the evaluation of gaze information. The CATI is a reliable 42-item inventory that evaluates 'sub-threshold' autistic traits that may be present in individuals who do not meet the diagnostic criteria for autism. The questionnaire measures the trait dimensions *Social Interactions*, *Communication Difficulty*, *Social Camouflage*, *Repetitive Behaviours*, *Cognitive Rigidity*, and *Sensory Sensitivity*. Participants responded on a five-point Likert scale, with higher scores indicating a greater endorsement of autistic traits. Mean CATI scores are reported in Table 2.

Attitudes towards AI Finally, since the general set of attitudes towards artificial agents may influence the tendency to adopt an intentional stance, in that negative attitudes decrease the attribution of mental states⁹⁶, we assessed attitudes towards AI with the *Attitudes Towards Artificial Intelligence Scale* (ATTARI-12).[\[@stein_2024\]](#) The ATTARI is a recently developed measure designed to measure both positive and negative attitudes on a five-point Likert scale (1 = strongly disagree, 5 = strongly agree). The scale comprises general attitudes, considering cognitive, affective, and behavioural facets, with four items each, leading to twelve items in total. Mean ATTARI scores are reported in Table 2. The data was collected for characterising the sample and for future exploratory purposes. A Welch Two Sample *t*-test was conducted to compare the total scores of participants who interacted first with the "AI-modelled" agent and participants who interacted first with the "human-modelled" data. The results indicated that participants who interacted first with the "AI-modelled" agent ($M = 172.1$, $SD = 38.33$) scored significantly higher than participants who interacted first with the "human-modelled" agent ($M = 166.88$, $SD = 42.02$), $t(1040.26) = 2.12$, $p = .034$. The data is available on the OSF framework <https://osf.io/wcm75>.

Pre-processing data

We followed the pre-processing procedure specified in our preregistration (<https://osf.io/wcm75>). We initially collected, as pre-registered, 160 participants. At the end of the task, participants rated their acceptance of the instruction that they completed the task with agents displaying human- and AI-modelled movements on a scale of 1–10. We noticed that nine participants provided ratings less than '6', indicating that they were sceptical of the manipulation. Hence, we recruited ten new participants (final sample of 170 participants) and we excluded those nine participants from all analyses. Next, we excluded nine participants from all analyses since they answered

incorrectly to one of the two attention-check questions regarding whether the agent in the upcoming block was to be human- or AI-modelled.

While there were not objectively ‘correct’ or ‘incorrect’ responses in the experimental task, we assessed illogical responses that might indicate a random response style, including a ‘give’ response for an object that the agent never gazed towards. Note that this was not possible in the **No Eye Contact** condition in which the agent gazed at all three objects in these trials. We conducted analyses only on ‘logical’ responses. To begin, we removed 1493 trials (4.83%) with excessively short (i.e., < 150 ms) or long RTs (i.e., > 3000 ms) because these were likely pre-emptive or ‘guess’ responses. We then excluded seven participants whose rates of illogical answers exceeded 35%. Of the remaining participants, eight were excluded for having a rate of illogical answers that were more than 2.5 *SD* away from the sample mean. For the remaining participants, we checked for evidence of acquiescent response styles by examining the distribution of their ‘give’ responses across the three object options. Given the task’s counterbalancing of the agent’s gaze location, we expected ‘give’ responses to be equally distributed across the three ‘give’ response keys, with an *SD* of response frequency close to zero. We excluded five participants whose *SDs* were 2.5 *SDs* greater than the average *SD* observed across these three response options, as this suggested a marked preference to ‘give’ a particular object significantly more (or less) than the others. Following this, we removed all ‘illogical’ response trials from the remaining participants (2.38% of trials) and, then, four participants for whom we had to remove an excessive number of trials (number of retained trials < 2.5 *SDs* the sample mean). Lastly, three participants who were either too slow or too fast were removed (RTs < or > 2.5 *SDs* than the sample mean). The final dataset comprised 125 participants. The demographics of this final dataset can be found in Table 2.

Analyses

All analyses followed our analysis plan pre-registered on the Open Science Framework (<https://osf.io/wcm75>) and were conducted using a custom R Markdown script with R Statistical Software (v4.5.0; R Core Team 2025), which can be found alongside all data on our OSF project page.

Behavioural measures

For the behavioural data (‘give’ frequencies and RTs), we used the same comprehensive data analysis approach as Caruana and colleagues⁴¹, to examine main effects and interactions related to **Eye Contact**, **Repeated Gaze**, **Belief**. For RT analyses, we also included the **Response** type (**Give**, **No-Give**) as an additional fixed effect, since the use of RT as an index of response certainty requires RTs to be examined as a function of the response provided. We transformed RTs using the box-cox *powerTransform* function from the *car* R package (v3.1.3)⁹⁷ to address the often skewed RT distribution⁹⁸, which may induce bias in the estimation of model parameters⁹⁹.

Next, we applied linear mixed-effects (LME) and generalised linear mixed-effects (GLME) analyses on unaggregated data, using the maximum likelihood estimation method within the *Lme4* R package (v1.1.37)¹⁰⁰. This approach was chosen because we expected that some participants might not have consistently provided a ‘give’ or ‘no-give’ response within each condition, leading to bias when data is aggregated because aggregated RTs in one condition would be based on more observations than RTs in others. To ensure a maximal random factor structure, we followed the analysis pipeline recommended by Scandola and Tidoni¹⁰¹ incorporating model optimisation and implementation of complex random intercept (CRI) models. This pipeline entails the formulation of random effects structures, where intercepts are determined by several interacting factors, in cases where maximal models with random slopes for all effects cannot be reliably estimated using typical human datasets. Further details on the full and reduced models can be found in the supplementary material on the OSF project page (<https://osf.io/wcm75>). All *p*-values for the (G)LME models were estimated with the *afex* R package (v1.4.1)¹⁰², using a significance criterion of $\alpha < 0.05$. We conducted post-hoc pairwise comparisons of individual conditions averaged across all other conditions using the *emmeans* R package (v1.11.0). For example, for the **Eye Contact** x **Repeated Gaze** interaction, we compared the **Eye Contact** and **No Eye Contact** conditions for both **No Repeated Gaze** and **Repeated Gaze**, averaging across all other conditions. This approach was applied to all significant interactions. Additionally, a Holm correction was used where appropriate to adjust contrasts for multiple comparisons.

Agent perception

To further investigate whether the perception of the agent differed between **Belief** conditions, we compared subjective ratings of the agent and the interaction made before and after each block between the **Belief** conditions using cumulative link models using the *ordinal* R package (v2023.12.4.1)¹⁰³.

To test whether, after the interaction, participants were more likely to prefer the AI (i.e., preference scores closer to 1 on the scale) or human (i.e., a preference score closer to 5 on the scale) with respect to its communication, pleasantness, and naturalness, we performed separate Wilcoxon Rank-Sum tests, since Shapiro-Wilk tests for normality indicated that the data was non-normally distributed (all *p*s < .001).

Finally, to explore participants’ most common reasons for their agent preferences, we visually represented key themes and concepts from free-text responses. To do this, we converted text responses into separate word lists for those that explained AI and human preferences, and generated word clouds out of these lists using the *wordcloud2* R package (v0.2.2)¹⁰⁴. Preference ratings below ‘3’ were considered as indicating a preference for AI, ratings of ‘3’ were classified as indicating no preference, and ratings above ‘3’ were interpreted as a preference for the human.

Data availability

All materials, data, and analysis scripts are available on the Open Science Framework repository (<https://osf.io/wcm75>). We have pre-registered our study design and analysis plan, which can be accessed at the same link.

Received: 12 April 2025; Accepted: 1 October 2025

Published online: 17 October 2025

References

- Al Kuwaiti, A. et al. A review of the role of artificial intelligence in healthcare. *J. Personal. Med.* **13**(6), 951. <https://doi.org/10.3390/jpm13060951> (2023).
- Chen, L., Chen, P. & Lin, Z. Artificial intelligence in education: A review. *IEEE Access.* **8**, 75264–75278. <https://doi.org/10.1109/ACCESS.2020.2988510> (2020).
- Van Pinxteren, M. M. E., Pluymaekers, M. & Lemmink, J. G. A. M. Human-like communication in conversational agents: A literature review and research agenda. *J. Serv. Manag.* **31**(2), 203–225. <https://doi.org/10.1108/JOSM-06-2019-0175> (2020).
- Vlačić, B., Corbo, L., Costa e Silva, S. & Dabić, M. The evolving role of artificial intelligence in marketing: A review and research agenda. *J. Bus. Res.* **128**, 187–203. <https://doi.org/10.1016/j.jbusres.2021.01.055> (2021).
- Kemp, A. Competitive advantage through artificial intelligence: Toward a theory of situated AI. *Acad. Manag. Rev.* **49**(3), 618–635. <https://doi.org/10.5465/amr.2020.0205> (2024).
- Holzinger, A., Keiblanger, K., Holub, P., Zatloukal, K. & Müller, H. AI for life: Trends in artificial intelligence for biotechnology. *New Biotechnol.* **74**, 16–24. <https://doi.org/10.1016/j.nbt.2023.02.001> (2023).
- Yin, J., Ngiam, K. Y. & Teo, H. H. Role of artificial intelligence applications in real-life clinical practice: Systematic review. *J. Med. Internet Res.* **23**(4), e25759. <https://doi.org/10.2196/25759> (2021).
- Stone, P., Brooks, R., & Brynjolfsson, E., et al. Artificial intelligence and life in 2030: The one hundred year study on artificial intelligence. arXiv (Cornell University). Published online (2022). <https://doi.org/10.48550/arXiv.2211.06318>
- Shibata, T., Tashima, T., & Tanie, K. Emergence of emotional behavior through physical interaction between human and robot. In *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C)*, vol. 4, 2868–2873 (1999). <https://doi.org/10.1109/ROBOT.1999.774032>.
- Hortensius, R., Hekele, F. & Cross, E. S. The perception of emotion in artificial agents. *IEEE Trans. Cogn. Dev. Syst.* **10**(4), 852–864. <https://doi.org/10.1109/TCDS.2018.2826921> (2018).
- Ruhland, K. et al. A review of eye gaze in virtual agents, social robotics and HCI: Behaviour generation, user interaction and perception. *Comput. Graph. Forum.* **34**(6), 299–326. <https://doi.org/10.1111/cgf.12603> (2015).
- Carpenter, J. et al. Gender representation and humanoid robots designed for domestic use. *Int. J. Soc. Robot.* **1**(3), 261–265. <https://doi.org/10.1007/s12369-009-0016-4> (2009).
- Epley, N., Waytz, A. & Cacioppo, J. T. On seeing human: A three-factor theory of anthropomorphism. *Psychol. Rev.* **114**(4), 864–886. <https://doi.org/10.1037/0033-295X.114.4.864> (2007).
- Eyssel, F., Kuchenbrandt, D., & Bobinger, S. Effects of anticipated human–robot interaction and predictability of robot behavior on perceptions of anthropomorphism. In *Proceedings of the 6th International Conference on Human-Robot Interaction*. HRI '11. Association for Computing Machinery, 61–68 (2011). <https://doi.org/10.1145/1957656.1957673>.
- Krach, S. et al. Can machines think? Interaction and perspective taking with robots investigated via fMRI. *PLoS ONE* **3**(7), e2597. <https://doi.org/10.1371/journal.pone.0002597> (2008).
- Wiese, E., Buzzell, G. A., Abubshait, A. & Beatty, P. J. Seeing minds in others: Mind perception modulates low-level social-cognitive performance and relates to ventromedial prefrontal structures. *Cogn. Affect. Behav. Neurosci.* **18**(5), 837–856. <https://doi.org/10.3758/s13415-018-0608-2> (2018).
- Bossi, F. et al. The human brain reveals resting state activity patterns that are predictive of biases in attitudes toward robots. *Sci. Robot.* **5**(46), eabb6652. <https://doi.org/10.1126/scirobotics.abb6652> (2020).
- Cross, E. S. & Ramsey, R. Mind meets machine: Towards a cognitive science of human-machine interactions. *Trends Cogn. Sci.* **25**(3), 200–212. <https://doi.org/10.1016/j.tics.2020.11.009> (2021).
- Abubshait, A. & Wiese, E. You look human, but act like a machine: Agent appearance and behavior modulate different aspects of human–robot interaction. *Front. Psychol.* **8**, 1393. <https://doi.org/10.3389/fpsyg.2017.01393> (2017).
- Cornelius, S., & Leidner, D. E. Acceptance of anthropomorphic technology: A literature review. In *Proceedings of the Annual Hawaii International Conference on System Sciences. Hawaii International Conference on System Sciences* (2021). <https://doi.org/10.24251/HICSS.2021.774>.
- Duffy, B. R. Anthropomorphism and the social robot. *Robot. Auton. Syst.* **42**(3), 177–190. [https://doi.org/10.1016/S0921-8890\(02\)00374-3](https://doi.org/10.1016/S0921-8890(02)00374-3) (2003).
- Fink, J. Anthropomorphism and human likeness in the design of robots and human–robot interaction. In *Social Robotics* (eds Ge, S. S., Khatib, O., Cabibihan, J. J., Simmons, R. & Williams, M. A.) 199–208 (Springer, 2012). https://doi.org/10.1007/978-3-642-34103-8_20.
- Kim, J. & Im, I. Anthropomorphic response: Understanding interactions between humans and artificial intelligence agents. *Comput. Hum. Behav.* **139**, 107512. <https://doi.org/10.1016/j.chb.2022.107512> (2023).
- Perez-Osorio, J. & Wykowska, A. Adopting the intentional stance toward natural and artificial agents. *Philos. Psychol.* **33**(3), 369–395. <https://doi.org/10.1080/09515089.2019.1688778> (2020).
- Axelrod, L. & Hone, K. Uncharted passions: User displays of positive affect with an adaptive affective system. In *Affective Computing and Intelligent Interaction* (eds Tao, J., Tan, T. & Picard, R. W.) 890–897 (Springer, 2005). https://doi.org/10.1007/11573548_114.
- Breazeal, C., Kidd, C. D., Thomaz, A. L., Hoffman, G., & Berlin, M. Effects of nonverbal communication on efficiency and robustness in human–robot teamwork. In: *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems* 708–713 (IEEE, 2005). <https://doi.org/10.1109/IROS.2005.1545011>.
- Li, M., Guo, F., Wang, X., Chen, J. & Ham, J. Effects of robot gaze and voice human-likeness on users' subjective perception, visual attention, and cerebral activity in voice conversations. *Comput. Hum. Behav.* **141**, 107645. <https://doi.org/10.1016/j.chb.2022.107645> (2022).
- Onnasch, L. & Hildebrandt, C. L. Impact of anthropomorphic robot design on trust and attention in industrial human–robot interaction. *J. Hum.-Robot Interact.* **11**(1), 2. <https://doi.org/10.1145/3472224> (2021).
- Gobel, M. S., Kim, H. S. & Richardson, D. C. The dual function of social gaze. *Cognition* **136**, 359–364. <https://doi.org/10.1016/j.cognition.2014.11.040> (2015).
- Kleinke, C. Gaze and eye contact: A research review. *Psychol. Bull.* **100**(1), 78–100. <https://doi.org/10.1037/0033-2909.100.1.78> (1986).
- Mundy, P. & Newell, L. Attention, joint attention, and social cognition. *Curr. Dir. Psychol. Sci. J. Am. Psychol. Soc.* **16**(5), 269–274. <https://doi.org/10.1111/j.1467-8721.2007.00518.x> (2007).
- Bruner, J. S. From communication to language: A psychological perspective. *Cognition* **3**(3), 255–287. [https://doi.org/10.1016/0010-0277\(74\)90012-2](https://doi.org/10.1016/0010-0277(74)90012-2) (1974).
- Loveland, K. A. & Landry, S. H. Joint attention and language in autism and developmental language delay. *J. Autism Dev. Disord.* **16**(3), 335–349. <https://doi.org/10.1007/BF01531663> (1986).
- Mundy, P. & Crowsom, M. Joint attention and early social communication: Implications for research on intervention with autism: Preschool issues in autism. *J. Autism Dev. Disord.* **27**(6), 653–676. <https://doi.org/10.1023/a:1025802832021> (1997).
- Gregory, S. E. A. & Kessler, K. Investigating age differences in the influence of joint attention on working memory. *Psychol. Aging* **37**(6), 731–741. <https://doi.org/10.1037/pag0000694> (2022).

36. Williams, J. H. G., Waiter, G. D., Perra, O., Perrett, D. I. & Whiten, A. An fMRI study of joint attention experience. *Neuroimage* **25**(1), 133–140. <https://doi.org/10.1016/j.neuroimage.2004.10.047> (2005).
37. Caruana, N., McArthur, G., Woolgar, A. & Brock, J. Simulating social interactions for the experimental investigation of joint attention. *Neurosci. Biobehav. Rev.* **74**(Pt A), 115–125. <https://doi.org/10.1016/j.neubiorev.2016.12.022> (2017).
38. Caruana, N. et al. The effect of non-communicative eye movements on joint attention. *Q. J. Exp. Psychol.* (2006) **73**(12), 2389–2402. <https://doi.org/10.1177/1747021820945604> (2020).
39. Caruana, N., Inkley, C., Nalepka, P., Kaplan, D. M. & Richardson, M. J. Gaze facilitates responsivity during hand coordinated joint attention. *Sci. Rep.* <https://doi.org/10.1038/s41598-021-00476-3> (2021).
40. Willemse, C. & Wykowska, A. In natural interaction with embodied robots, we prefer it when they follow our gaze: A gaze-contingent mobile eyetracking study. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **374**(1771), 20180036. <https://doi.org/10.1098/rstb.2018.0036> (2019).
41. Caruana, N., Hechler, F. C., Cross, E. S. & Tidoni, E. The temporal context of eye contact influences perceptions of communicative intent. *R. Soc. Open Sci.* **12**(7), 250277. <https://doi.org/10.1098/rsos.250277> (2025).
42. Hechler, F. C. et al. Autistic and non-autistic people evaluate eye contact cues in context to identify communicative opportunities [preprint]. *Published online* https://doi.org/10.31234/osf.io/zcvp6_v1 (2025).
43. Schellen, E. & Wykowska, A. Intentional mindset toward robots: Open questions and methodological challenges. *Front. Robot. AI* <https://doi.org/10.3389/frobt.2018.00139> (2019).
44. Dennett, D. C. Intentional systems. *J. Philos.* **68**(4), 87–106. <https://doi.org/10.2307/2025382> (1971).
45. Dennett, D. C. True believers: The intentional strategy and why it works. In *Scientific Explanation: Papers Based on Herbert Spencer Lectures Given in the University of Oxford* (ed. Heath, A. F.) 150–167 (Clarendon Press, 1981).
46. Dennett, D. C. *The Intentional Stance* (The MIT Press, 1987).
47. Papagni, G. & Koeszegi, S. A pragmatic approach to the intentional stance semantic, empirical and ethical considerations for the design of artificial agents. *Minds Mach. (Dordr.)* **31**(4), 505–534. <https://doi.org/10.1007/s11023-021-09567-6> (2021).
48. Kiesler, S., Powers, A., Fussell, S. & Torrey, C. Anthropomorphic interactions with a robot and robot-like agent. *Soc. Cogn.* **26**, 169–181. <https://doi.org/10.1521/soco.2008.26.2.169> (2008).
49. Yamazaki, A., Yamazaki, K., Burdelski, M., Kuno, Y. & Fukushima, M. Coordination of verbal and non-verbal actions in human-robot interaction at museums and exhibitions. *J. Pragmat.* **42**(9), 2398–2414. <https://doi.org/10.1016/j.pragma.2009.12.023> (2010).
50. Wiese, E., Metta, G. & Wykowska, A. Robots as intentional agents: Using neuroscientific methods to make robots appear more social. *Front. Psychol.* <https://doi.org/10.3389/fpsyg.2017.01663> (2017).
51. Abubshait, A. & Wykowska, A. Repetitive robot behavior impacts perception of intentionality and gaze-related attentional orienting. *Front. Robot. AI.* <https://doi.org/10.3389/frobt.2020.565825> (2020).
52. Wiese, E., Wykowska, A., Zwicker, J. & Müller, H. I see what you mean: How attentional selection is shaped by ascribing intentions to others. *PLoS ONE* **7**(9), 1–7. <https://doi.org/10.1371/journal.pone.0045391> (2012).
53. Wykowska, A. & Wiese, E. Beliefs about the minds of others influence how we process sensory information. *PLoS ONE* **9**(4), 1–11. <https://doi.org/10.1371/journal.pone.0094339> (2014).
54. Özdem, C. et al. Believing androids - fMRI activation in the right temporo-parietal junction is modulated by ascribing intentions to non-human agents. *Soc. Neurosci.* **12**(5), 582–593. <https://doi.org/10.1080/17470919.2016.1207702> (2017).
55. Caruana, N., Spiro, D. & Brock, J. Human agency beliefs influence behaviour during virtual social interactions. *PeerJ (San Francisco, CA)* **5**, e3819. <https://doi.org/10.7717/peerj.3819> (2017).
56. Caruana, N., de Lissa, P. & McArthur, G. Beliefs about human agency influence the neural processing of gaze during joint attention. *Soc. Neurosci.* **12**(2), 194–206. <https://doi.org/10.1080/17470919.2016.1160953> (2017).
57. Caruana, N. & McArthur, G. The mind minds minds: The effect of intentional stance on the neural encoding of joint attention. *Cogn. Affect. Behav. Neurosci.* **19**(6), 1479–1491. <https://doi.org/10.3758/s13415-019-00734-y> (2019).
58. Scandola, M., Cross, E. S., Caruana, N. & Tidoni, E. Body form modulates the prediction of human and artificial behaviour from gaze observation. *Int. J. Soc. Robot.* **15**(8), 1365–1385. <https://doi.org/10.1007/s12369-022-00962-2> (2023).
59. Tidoni, E. et al. Human but not robotic gaze facilitates action prediction. *iScience* **25**(6), 104462. <https://doi.org/10.1016/j.isci.2022.104462> (2022).
60. Guingrich, R. E. & Graziano, M. S. A. Ascribing consciousness to artificial intelligence: Human–AI interaction and its carry-over effects on human–human interaction. *Front. Psychol.* <https://doi.org/10.3389/fpsyg.2024.1322781> (2024).
61. Thellman, S., de Graaf, M. & Ziemke, T. Mental state attribution to robots: A systematic review of conceptions, methods, and findings. *J. Hum.-Robot Interact.* <https://doi.org/10.1145/3526112> (2022).
62. Jording, M., Engemann, D., Eckert, H., Bente, G. & Vogeley, K. Distinguishing social from private intentions through the passive observation of gaze cues. *Front. Hum. Neurosci.* **13**, 442. <https://doi.org/10.3389/fnhum.2019.00442> (2019).
63. Senju, A. & Csibra, G. Gaze following in human infants depends on communicative signals. *Curr. Biol.* **18**(9), 668–671. <https://doi.org/10.1016/j.cub.2008.03.059> (2008).
64. Csibra, G. & György, G. Natural pedagogy. *Trends Cogn. Sci.* **13**, 148–153. <https://doi.org/10.1016/j.tics.2009.01.005> (2009).
65. Csibra, G. & Gergely, G. Natural pedagogy as evolutionary adaptation. *Philos. Trans. R. Soc. B Biol. Sci.* **366**(1567), 1149–1157. <https://doi.org/10.1098/rstb.2010.0319> (2011).
66. Szufnarowska, J. et al. Is ostension any more than attention? *Sci. Rep.* **4**, 5304. <https://doi.org/10.1038/srep05304> (2014).
67. Gredebäck, G., Astor, K. & Fawcett, C. Gaze following is not dependent on ostensive cues: A critical test of natural pedagogy. *Child Dev.* **89**(6), 2091–2098. <https://doi.org/10.1111/cdev.13026> (2018).
68. Senju, A., Csibra, G. & Johnson, M. H. Understanding the referential nature of looking: Infants' preference for object-directed gaze. *Cognition* **108**(2), 303–319. <https://doi.org/10.1016/j.cognition.2008.02.009> (2008).
69. Jording, M., Hartz, A., Bente, G., Schulte-Rüther, M. & Vogeley, K. Inferring interactivity from gaze patterns during triadic person-object-agent interactions. *Front. Psychol.* **10**, 1913. <https://doi.org/10.3389/fpsyg.2019.01913> (2019).
70. Sperber, D. & Wilson, D. *Relevance: Communication and Cognition* (Harvard University Press, 1986).
71. Ricciardelli, P., Carcagno, S., Vallar, G. & Bricolo, E. Is gaze following purely reflexive or goal-directed instead? Revisiting the automaticity of orienting attention by gaze cues. *Exp. Brain Res. Exp. Hirnforsch. Exp. Cereb.* **224**, 93–106. <https://doi.org/10.1007/s00221-012-3291-5> (2012).
72. Alhasan, A. & Caruana, N. Evidence for the adaptive parsing of non-communicative eye movements during joint attention interactions. *PeerJ (San Francisco, CA)* **11**, e16363. <https://doi.org/10.7717/peerj.16363> (2023).
73. Anthis, J. R., Pauketat, J. V. T., Ladak, A., & Manoli, A. What do people think about sentient AI? [arXiv preprint]. *Published online* (2024).
74. Kampe, K. K. W., Frith, C. D. & Frith, U. “Hey John”: Signals conveying communicative intention toward the self activate brain regions associated with “mentalizing,” regardless of modality. *J. Neurosci.* **23**(12), 5258–5263. <https://doi.org/10.1523/JNEUROSCI.23-12-05258.2003> (2003).
75. Senju, A. & Johnson, M. The eye contact effect: Mechanisms and development. *Trends Cogn. Sci.* **13**(3), 127–134. <https://doi.org/10.1016/j.tics.2008.11.009> (2009).
76. Conty, L., George, N. & Hietanen, J. K. Watching eyes effects: When others meet the self. *Conscious. Cogn.* **45**, 184–197. <https://doi.org/10.1016/j.concog.2016.08.016> (2016).

77. Hietanen, J. O. & Hietanen, J. K. Genuine eye contact elicits self-referential processing. *Conscious. Cogn.* **51**, 100–115. <https://doi.org/10.1016/j.concog.2017.01.019> (2017).
78. Caniguer, R. & de Hamilton, A. F. C. The role of eye gaze during natural social interactions in typical and autistic people. *Front. Psychol.* **10**, 560–560. <https://doi.org/10.3389/fpsyg.2019.00560> (2019).
79. Hechler, F. C. et al. The role of mentalising and agent appearance on gaze use during joint attention [preprint]. *Published online* https://doi.org/10.31234/osf.io/ser72_v2 (2025).
80. Stanford Emerging Technology Review (Stanford University, 2025). <https://setr.stanford.edu/technology-areas-2025>
81. *Attitudes Towards the Impact of Digitisation and Automation on Daily Life* (European Commission, 2017). <https://doi.org/10.2759/835661>.
82. Stibe, A. & Dinh, T. H. N. Exploring human artificial intelligence using the knowledge behavior gap model. In *Mobile Web and Intelligent Information Systems* (eds Younas, M. et al.) 189–203 (Springer, 2024).
83. Su, H. et al. Recent advancements in multimodal human–robot interaction. *Front. Neurorobot.* <https://doi.org/10.3389/fnbot.2023.1084000> (2023).
84. Bartneck, C. Interacting with an embodied emotional character. In *Proceedings of the 2003 International Conference on Designing Pleasurable Products and Interfaces. DPPI '03. Association for Computing Machinery* 55–60 (2003). <https://doi.org/10.1145/782896.782911>.
85. Wainer, J., Feil-Seifer, D. J., Shell, D. A., & Mataric, M. J. Embodiment and human–robot interaction: A task-based perspective. In *RO-MAN 2007: The 16th IEEE International Symposium on Robot and Human Interactive Communication* 872–877 (2007). <https://doi.org/10.1109/ROMAN.2007.4415207>.
86. Li, J. The benefit of being physically present: A survey of experimental works comparing copresent robots, telepresent robots and virtual agents. *Int. J. Hum. Comput. Stud.* **77**, 23–37. <https://doi.org/10.1016/j.ijhcs.2015.01.001> (2015).
87. Thellman, S., Sillevarg, A., Gulz, A., & Ziemke, T. Physical vs. virtual agent embodiment and effects on social interaction. In *Intelligent Virtual Agents* (eds Traum, D., Swartout, W., Khooshabeh, P., Kopp, S., Scherer, S. & Leuski, A.) 412–415 (Springer, 2016). https://doi.org/10.1007/978-3-319-47665-0_44.
88. Jecker, N. S. & Nakazawa, E. Bridging east-west differences in ethics guidance for AI and robotics. *AI* **3**(3), 764–777. <https://doi.org/10.3390/ai3030045> (2022).
89. Mantello, P., Ho, T., Nguyen, M. H. & Vuong, Q. H. Bosses without a heart: Socio-demographic and cross-cultural determinants of attitude toward emotional AI in the workplace. *AI Soc.* **36**, AOP. <https://doi.org/10.1007/s00146-021-01290-1> (2021).
90. Belanche, D., Casaló, L. V. & Flavián, C. Artificial intelligence in FinTech: Understanding robo-advisors adoption among customers. *Ind. Manag. Data Syst.* **119**(7), 1411–1430. <https://doi.org/10.1108/IMDS-08-2018-0368> (2019).
91. Schepman, A. & Rodway, P. The general attitudes towards artificial intelligence scale (GAAIS): Confirmatory validation and associations with personality, corporate distrust, and general trust. *Int. J. Hum.-Comput. Interact.* **39**(13), 2724–2741. <https://doi.org/10.1080/10447318.2022.2085400> (2023).
92. Zhang, B., & Dafoe, A. Artificial intelligence: American attitudes and trends. *SSRN Electronic Journal*. Published online January (2019). <https://doi.org/10.2139/ssrn.3312874>.
93. Sollfrank, T. et al. The effects of dynamic and static emotional facial expressions of humans and their avatars on the EEG: An ERP and ERD/ERS study. *Front. Neurosci.* **15**, 651044. <https://doi.org/10.3389/fnins.2021.651044> (2021).
94. English, M. et al. The comprehensive autistic trait inventory (CATI): Development and validation of a new measure of autistic traits in the general population. *Mol. Autism.* <https://doi.org/10.1186/s13229-021-00445-7> (2021).
95. Stein, J. P., Messingschlager, T., Gnambs, T., Hutmacher, F. & Appel, M. Attitudes towards AI: Measurement and associations with personality. *Sci. Rep.* **14**, 2909. <https://doi.org/10.1038/s41598-024-53335-2> (2024).
96. Spatola, N., Marchesi, S. & Wykowska, A. The intentional stance test-2: How to measure the tendency to adopt intentional stance towards robots. *Front. Robot. AI* **8**, 666586. <https://doi.org/10.3389/frobt.2021.666586> (2021).
97. Fox, J., Weisberg, S. & Fox, J. *An R Companion to Applied Regression* 2nd edn. (SAGE Publications, 2011).
98. Bono Cabré, R., Blanca Mena, M. J., Arnau Gras, J. & Gómez Benito, J. Non-normal distributions commonly used in health, education, and social sciences: A systematic review. *Front. Psychol.* <https://doi.org/10.3389/fpsyg.2017.01602> (2017).
99. Balota, D. A., Aschenbrenner, A. J. & Yap, M. J. Additive effects of word frequency and stimulus quality: The influence of trial history and data transformations. *J. Exp. Psychol. Learn. Mem. Cogn.* **39**(5), 1563–1571. <https://doi.org/10.1037/a0032186> (2013).
100. Bates, D., Mächler, M., Bolker, B. & Walker, S. Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* **67**(1), 1–48. <https://doi.org/10.18637/jss.v067.i01> (2015).
101. Scandola, M. & Tidoni, E. Reliability and feasibility of linear mixed models in fully crossed experimental designs. *Adv. Methods Pract. Psychol.* **7**(1), 25152459231214456. <https://doi.org/10.1177/25152459231214454> (2024).
102. Singmann, H., Bolker, B., Westfall, J., Aust, F. & Ben-Shachar, M. S. Afex: Analysis of Factorial Experiments (2024). <https://CRAN.R-project.org/package=afex>.
103. Christensen, R. H. B. *Ordinal—Regression Models for Ordinal Data* (2023). <https://doi.org/10.32614/CRAN.package.ordinal>.
104. Zhao, Y. wordcloud2: Word cloud for r. Published online (2019). <https://cran.r-project.org/web/packages/wordcloud2/index.html>.

Author contributions

FCH: Conceptualisation, Methodology, Validation, Formal Analysis, Investigation, Resources, Data Curation, Writing—Original Draft, Writing—Review & Editing, Visualisation, Supervision, Project Administration. ET: Conceptualisation, Methodology, Validation, Formal Analysis, Resources, Data Curation, Writing—Review & Editing, Visualisation, Supervision, Project Administration. ESC: Conceptualisation, Methodology, Writing—Review & Editing. NC: Conceptualisation, Methodology, Validation, Formal Analysis, Resources, Data Curation, Writing—Original Draft, Writing—Review & Editing, Visualisation, Supervision, Project Administration.

Funding

Open Access funding enabled and organized by Projekt DEAL. FCH received a Ph.D. scholarship from Hans-Boeckler-Stiftung. Experimental development was supported by funding from the Experimental Psychology Society funded to TD and NC. ET and NC contributions were also supported in part by the BIAL foundation (137/24) EC's contributions were supported in part by the BIAL foundation (326/2024).

Declarations

Competing interests

The authors declare that they have no competing interests.

Ethics approval and consent to participate

The study was approved by the Flinders University Human Research Ethics Committee (6804). All participants provided written informed consent.

Additional information

Correspondence and requests for materials should be addressed to F.C.H.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025