



# Examining Social Cognition with Embodied Robots: Does Prior Experience with a Robot Impact Feedback-associated Learning in a Gambling Task?

ABDULAZIZ ABUBSHAIT 

CRAIG G. MCDONALD 

EVA WIESE 

*\*Author affiliations can be found in the back matter of this article*

SPECIAL COLLECTION:  
EXPLORING  
COGNITION USING  
NEW TECHNOLOGICAL  
ADVANCES

RESEARCH ARTICLE

][ubiquity press

## ABSTRACT

Social agents rely on the ability to use feedback to learn and modify their behavior. The extent to which this happens in social contexts depends on motivational, cognitive and/or affective parameters. For instance, feedback-associated learning occurs at different rates when the outcome of an action (e.g., winning or losing in a gambling task) affects oneself (“Self”) versus another human (“Other”). Here, we examine whether similar context effects on feedback-associated learning can also be observed when the “other” is a social robot (here: Cozmo). We additionally examine whether a “hybrid” version of the gambling paradigm, where participants are free to engage in a dynamic interaction with a robot, then move to a controlled screen-based experiment can be used to examine social cognition in human-robot interaction. This hybrid method is an alternative to current designs where researchers examine the effect of the interaction on social cognition during the interaction with the robot. For that purpose, three groups of participants (n total = 60) interacted with Cozmo over different time periods (no interaction vs. a single 20 minute interaction in the lab vs. daily 20 minute interactions over five consecutive days at home) before performing the gambling task in the lab. The results indicate that prior interactions impact the degree to which participants benefit from feedback during the gambling task, with overall worse learning immediately after short-term interactions with the robot and better learning in the “Self” versus “Other” condition after repeated interactions with the robot. These results indicate that “hybrid” paradigms are a suitable option to investigate social cognition in human-robot interaction when a fully dynamic implementation (i.e., interaction and measurement dynamic) is not feasible.

CORRESPONDING AUTHOR:  
**Abdulaziz Abubshait**

George Mason University, US;  
Italian Institute of Technology,  
IT

[Abdulaziz.Abubshait@iit.it](mailto:Abdulaziz.Abubshait@iit.it)

KEYWORDS:

Cognitive Control; Learning;  
Reward processing; Social  
cognition

TO CITE THIS ARTICLE:

Abubshait, A., McDonald  
C. G., & Wiese, E. (2021).  
Examining Social Cognition  
with Embodied Robots: Does  
Prior Experience with a Robot  
Impact Feedback-associated  
Learning in a Gambling Task?  
*Journal of Cognition*, 4(1):  
28, pp. 1–16. DOI: [https://doi.  
org/10.5334/joc.167](https://doi.org/10.5334/joc.167)

As social creatures, a considerable part of our lives revolves around interactions with other human beings. However, due to an increased availability of artificially intelligent agents in modern societies, our future social interactions will likely expand to nonhuman agents like avatars or robots (Wiese et al., 2017). In fact, robots have already been implemented in elderly care as social assistants to increase emotional comfort (e.g., Birks et al., 2016; Tapus et al., 2007), in therapeutic settings with children with autism spectrum disorder to train social-cognitive skills (Bekele et al., 2014; Warren et al., 2015), as well as in rehabilitation settings to improve sensorimotor skills (Basteris et al., 2014). Nonetheless, despite considerable progress in equipping artificial agents with social capabilities, they are still limited in their ability to interact with humans in a natural way (Wiese et al., 2017), and the public remains skeptical concerning the introduction of robot assistants to everyday life (Bartneck & Reichenbach, 2005). Specifically, the use of robots as companions for children has been discussed controversially due to concerns about privacy, fear of reduced interest in human-human interactions due to increased attachment to the robot companion, and anticipation of negative impacts on learning outcomes and development (Sharkey, 2016). While these concerns should be taken seriously, there is a lack of empirical studies that systematically examine the impact of robots on social cognition, development and wellbeing; this is particularly true with regard to long-term social interactions in everyday environments. Since robots will employ social roles in our society and share our environments with us in the future, it is essential to understand how to design them so that they can foster rewarding long-term social interactions by activating relevant social schemes, behaviors and emotions without negatively impacting human-human interactions.

To examine whether and how humans develop social attachments to robot companions, experimental paradigms need to (i) adequately reflect the dynamics of long-term social interactions with embodied robots in everyday environments (i.e., external validity) but at the same time (ii) allow for controlled and reproducible measures of social cognition (i.e., internal validity). Most importantly, behavioral and neurophysiological studies have shown that social-cognitive mechanisms unfold differently in dynamic or “online” interactions (in contrast to the simulation of social interactions on the screen or “offline” paradigms), and activate brain areas involved in social cognition differently (Schilbach et al., 2013). For instance, while offline paradigms have shown that mechanisms of social attention (i.e., the extent to which observers shift their attention to locations that are gazed-at (see Frischen et al., 2007; for a review) differ between human-human and human-robot interaction (Admoni et al., 2011), online paradigms examining the same mechanisms in face-to-face interactions using embodied robot platforms show that these mechanisms resemble human-human interaction more closely than originally thought (Kompatsiari et al., 2018). It was also shown that depending on whether “offline” or “online” paradigms are employed, different brain regions are implicated in social attention: offline fMRI studies identify brain regions in the right hemifield as important neural correlates of social attention (Capozzi & Ristic, 2020; for a review), whereas online studies that use face-to-face paradigms implicate structures in the left hemifield (e.g., Cole et al., 2016). These studies show that in order to get a realistic idea of how social-cognitive mechanisms are engaged in human-robot interaction, paradigms should include an “online” component, such that at least one part of the experiment should allow participants to engage in dynamic interactions with an embodied robot.

Longer-term social interactions with others require monitoring the behavior of others and adjusting our behaviors to theirs to ensure successful exchanges of knowledge, affiliation and support (Insel & Fernald, 2004). In order to be able to adapt behaviors to ever-changing environments and learn from previous experiences, we rely on feedback of others to tell the difference between behavioral responses that are appropriate and those that are not. Receiving positive feedback (e.g., smile) positively reinforces a given behavior and increases the likelihood that it will be shown again in the future; negative feedback (e.g., frown) negatively reinforces a given behavior and decreases the likelihood that it will be shown again in the future (Krigolson et al., 2009). There is also a direct link between feedback processing and learning, such that the response to feedback is stronger (or weaker) the less (or more) advanced the learning progress is: participants use feedback to map their expectations regarding an outcome to the

actual outcome, with positive feedback indicating a good match (and positively reinforcing an expectation) and negative feedback indicating a bad match (and negatively reinforcing an expectation) between expectation and outcome (i.e., reinforcement learning; Holroyd & Coles, 2002; Schultz, 2017). While learners rely on feedback at the beginning of the learning process, they depend increasingly less on feedback in the later stages of a learning process but place more weight on their own responses to inform upcoming behaviors (Holroyd & Coles, 2002; Krigolson et al., 2009). This suggests that how we learn in social interactions is a consequence of how we process (i.e., feedback monitoring) and are reinforced by others' feedback (i.e., reinforcement learning; Hobson & Inzlicht, 2016). While feedback processing is relatively well understood in human-human interactions, it has not been examined in human-robot interactions yet (to our knowledge). This is of particular importance, however, given that robots are already used in educational and therapeutic settings (and will be used even more in these fields in the future; Rahwan et al., 2019) where a reduced response to feedback could have measurable negative effects on learning outcomes. Feedback monitoring and reward processing are also linked to prosocial behavior (de Bruijn et al., 2017), and as such utterly important for supporting positive long-term human-robot interactions.

One paradigm that has traditionally been used to examine feedback processing (Gehring & Willouby, 2002) and has been adapted to investigate social reinforcement learning is the gambling task (e.g., Hassall et al., 2016; Kimura & Katayama, 2016; Krigolson et al., 2013; Lemoine & Roland-Lévy, 2017; Leng & Zhou, 2010; Rigoni et al., 2010). In one variant of this task, participants are asked to gamble by picking (on a trial-by-trial basis) one of two differently colored squares (one of which is associated with a higher chance of winning than the other one) shown next to each other on a screen over a consecutive sequence of trials. After each selection, participants receive feedback whether they won ("Win", i.e., positive feedback) or lost ("Lose", i.e., negative feedback) this trial. Electrophysiological studies using this task have shown that event-related potentials (ERPs) implicated in reward processing, such as the *Feedback-Related Negativity* (FRN) and/or *Reward Positivity* (RewP), are more pronounced when gambling for oneself ("Self") versus another person ("Other"; Hassall et al., 2016; Krigolson et al., 2013). The extent to which participants learn to select the option that is associated with a higher chance of winning can be predicted by the extent to which participants process feedback, such that a stronger reliance on feedback cues is associated with worse learning outcomes (Lohse et al., 2020).

Importantly, these mechanisms are sensitive to the context in which feedback is provided. For instance, feedback processing is altered by whether a person's performance has a direct impact on a partner's performance (de Bruijn et al., 2011; Koban et al., 2012; Koban et al., 2013), whether the interaction is cooperative or competitive (Czeszumski et al., 2019; de Bruijn et al., 2011; Radke et al., 2011; Van Meel & Van Heijningen, 2010) or whether a social interaction partner is present during the delivery of feedback (Simon et al., 2014). Most importantly, the relationship between the gambler and the recipient of the gambling outcome (i.e., the "Other" in the gambling task) modulates feedback processing, such that when strangers are the recipients of winning outcomes, feedback processing is attenuated compared to when participants themselves receive the outcome (e.g., Hassall, Silver, Turk, & Krigolson, 2016; Krigolson, Hassall, Balcom & Turk, 2013). The presence of a social ingroup member (as opposed to a social outgroup member) also attenuates Self-Other differences, as the motivation to win is higher for a social agent that is believed to be similar to oneself (versus dissimilar; Hobson & Inzlicht, 2016). In line with this notion, there is no observable Self-Other difference in feedback processing when the recipient of a positive outcome is a friend versus a stranger (Leng & Zhou, 2010) indicating that feedback processing is sensitive to the social relationship between the gambler and the recipient of the outcome.

These findings suggest that feedback processing may differ in human-human- vs. human-robot interaction, partially because nonhuman social entities are often perceived as outgroup members (Hackel et al., 2014); they also suggest that similar to human-human interactions, feedback processing in human-robot interaction may change over time if participants start to get attached to robots and perceive them as friends. Support for this assumption comes from a recent neurophysiological study using the gambling task to examine whether familiarization

with the robot Cozmo<sup>1</sup> modulates feedback processing after a one-time interaction: one group of participants engaged in a 20-minute interaction with Cozmo (playing interactive games with the robot via the associated mobile app) and the other group played the Simon Says game (an interactive game that does not involve Cozmo) before gambling for themselves (“Self”) or Cozmo (“Other”). The data shows larger RewP amplitudes and slower learning of the contingencies between color and chance of winning after one-time familiarization with Cozmo (Abubshait et al., 2021). Although this effect was not specific to the “Other” condition, but affected the “Self” condition as well, the results indicate that familiarization with a social robot does indeed modulate feedback-based learning in short-term interactions. What causes this effect and whether similar observations can be made for longer interaction durations warrants further investigation.

The goal of this paper is to present a proof-of-concept that hybrid paradigms that separate a dynamic, relatively unscripted interaction component (here: interaction with Cozmo via its associated app) from a highly controlled, scripted data collection component are suitable to investigate mechanisms of social cognition in human-robot interaction. Specifically, we explore how different degrees of familiarity with a robot due to prior interactions of different durations impact behavioral correlates of reward processing in form of feedback-associated learning in a gambling task (see Abubshait et al., 2021). For that purpose, we compare feedback-associated learning (i.e., speed with which participants learn to pick the square associated with the higher chance of winning) after a **one-time interaction** in the lab (about 20 minutes) to feedback-associated learning after **repeated interactions** with Cozmo in participants’ homes (over five consecutive days); both conditions are compared to a **no interaction** condition, in which participants do not familiarize themselves with Cozmo before performing the gambling task but instead play an interactive game that does not involve Cozmo (i.e., *Simon says*). In all three conditions, the participants then perform a computer-based version of the gambling task with Cozmo sitting underneath the screen in the participants’ field of view, making occasional eye blinks to signal that it is “on”.

This setup creates high external validity by allowing unrestricted social interactions with Cozmo to occur AND high internal validity due to the controlled setting in which reward processing is examined. As feedback-associated learning is impacted by a range of factors including social (e.g., social closeness between gambler and recipient; Leng & Zhou, 2010), motivational (e.g., better learning when intrinsically motivated; Wilhelm et al., 2019) and physiological (e.g., arousal: worse learning in engaging vs. calm environments; Lohse et al., 2020; Wilhelm et al., 2019) factors, it is challenging to formulate specific directed hypotheses regarding the effect of familiarization duration on reward processing. We generally expect that increased intrinsic motivation to perform well for Cozmo due to increased familiarity with the robot after longer initial familiarization, should positively affect feedback processing and lead to more effective learning of the association between “color” and “chance of winning” during gambling. Increased levels of arousal, in contrast, for instance due to an increased level of engagement, should negatively impact feedback processing and learning.

## METHODS & MATERIALS

### PARTICIPANTS

60 participants were recruited from George Mason University’s undergraduate participant pool (Mean Age = 22.34, Range = 18–55, SD = 5.9, 38 females) in exchange for course credit or monetary compensation. Participants have neither interacted with nor owned a Cozmo before participating in the study. Participants were assigned to the no-interaction condition (NI;  $n = 16$ ), the one-time interaction condition (OTI;  $n = 17$ ) or the repeated-interaction condition (RI;  $n = 18$ ). All subjects were right-handed, had normal or corrected-to-normal vision.<sup>2</sup> The data of

---

1 Cozmo is the ideal platform to examine social-cognitive processes in human-robot interaction due to its affordability, ease of use, socio-emotional expressiveness, availability of source code and accessibility of interaction data (see Chaudhury et al., 2020; *preprint*; for a tutorial), and as such is increasingly used in psychology and computer science research (Charrier et al., 2019; Ciardo et al., 2020; Cross et al., 2019; Hsieh et al., 2020; Lefkeli et al., 2020; McNeill & Kennington, 2019; Pelikan et al., 2020).

2 Please note that in addition to the behavioral data, we collected EEG data in all three conditions. The EEG data from the no (NI) and the one-time (OTI) interaction condition are reported in Abubshait et al. (2021). The EEG specific procedures are described there. The EEG data of the repeated interaction (RI) condition has not been published yet.

several participants were excluded from the final analyses due to technical difficulties ( $n = 3$ ) or participants not following the task protocol ( $n = 2$ ) or not completing the study ( $n = 2$ ); two additional participant data sets were excluded due to corrupted data files. Data handling and collection was in accordance to George Mason University's ethics board. Raw data can be found on the OSF page of the experiment: [osf.io/6dfky/](https://osf.io/6dfky/).

## APPARATUS

In the social interaction conditions (OTI, RI), participants interacted with the social robot Cozmo (Digital Dreamlabs, USA, version 3.2.0) using a Samsung Galaxy tablet (**Figure 1**). Cozmo is equipped with three cubes that are used to play interactive games with the robot (e.g., tapping a cube as fast as possible); all interactions are pre-programmed using the Cozmo app. In the no-interaction condition (NI), participants played the Simon Says game (Hasbro, Inc, USA), an engaging task that requires them to respond to different color patterns (**Figure 1**). The gambling task was programmed and presented using Matlab (Mathworks, Natick, MA) and Psychtoolbox (Brainard, 1997). Analyses were conducted using R (version 3.6), including the packages tidyverse (Wickham et al., 2019) and lme4 (Bates et al., 2014).



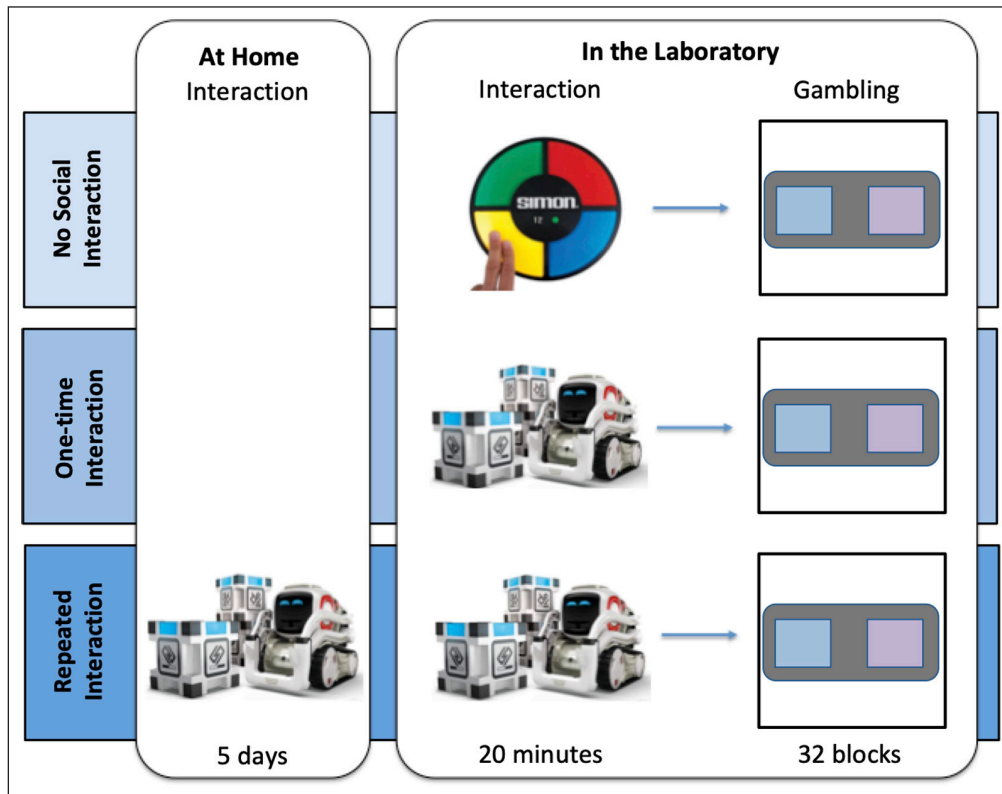
**Figure 1** The Cozmo robot and the Simon Says game:

Cozmo (used for one-time and repeated interaction conditions) is a tank-like robot with a screen-like face and blue squares as eyes. It can express a wide range of emotions using its eyes and – in addition – sounds. More information on the robot Cozmo can be found here: [www.digitaldreamlabs.com/pages/cozmo](http://www.digitaldreamlabs.com/pages/cozmo). The Simon Says game (used for the no social interaction condition) is an electronic game that is equipped with pressable buttons that light up. When playing Simon Says, players see a series of buttons light up that are associated with a tone in a specific sequence that the player needs to mimic. After the player successfully mimics the sequence, the sequence would increase in length. The Simon Says game was chosen due to its resemblance to the Quick-tap game.

## PROCEDURE

After providing consent, all participants were given a cover story that the experiment was a collaborative effort between the psychology and engineering departments at GMU. Specifically, they were told that “the engineering department was trying to decide which of their robots to upgrade and that the psychology department is helping them make that decision by conducting user-centered testing”. Afterwards, participants were introduced to Cozmo and familiarized with its basic functionalities. All participants were instructed that the latter part of the study will include a part where they will play a gambling task where they would be gambling for a chance to win a gift card for themselves (“Self”) for half of the blocks, and gambling for new hardware and software upgrades for Cozmo (e.g., batteries, “Cozmo”) for the other half of the blocks.

Participants were assigned to one of three interaction conditions: (i) **no interaction** where participants did not interact with Cozmo prior to the gambling task but played the Simon Says game once for 20 minutes instead; (ii) **one-time interaction** where participants interacted once for 20 minutes with Cozmo prior to the gambling task and (iii) **repeated interaction** where participants interacted with Cozmo over the course of five days for a minimum of 20 minutes a day in their homes. In the *no social interaction* and *one-time social interaction*, participants performed the interaction task (see below for more details) and immediately afterwards completed the gambling task. In the *repeated social interactions* condition, participants took Cozmo home and interacted with it for at least 20 minutes every day (see below for more details). On day number 5, they came back to the lab and performed the gambling task (after a refresher of the instructions). At the end of the experiment, participants were debriefed and thanked for their participation. A timeline of events for the three interaction conditions can be found in **Figure 2**.



**Figure 2 Overview of the three interaction conditions:** Participants in the no interaction condition arrived to the lab and played the Simon Says game before completing the gambling task. Participants in the one-time interaction arrived to the lab and interacted with Cozmo once before completing the gambling task. Participants in the repeated interaction condition came in at the beginning of the week and were given a Cozmo robot to take home. They were instructed to interact with Cozmo at least 20 minutes per day. No other restrictions were made on their interaction. On the fifth day, they came back to the lab and then completed the gambling task.

## INTERACTION CONDITIONS


In the **no-interaction** condition, participants played the *Simon Says* game for about 20 minutes, which required them to imitate a sequence of tones and lights presented on an electronic device; see **Figure 1**. Every time the participant successfully imitated a sequence, the game would add an extra tone/light to the sequence, which made the game more challenging and gave participants an idea of how well they were performing throughout the game.

In the **one-time interaction** condition, participants played two games with Cozmo for about 20 minutes through the Cozmo app: *Keep Away* and *Quick Tap*. On a given trial of *Keep Away*, participants held one of the cubes in front of Cozmo (within the robot's reaching distance). In the meanwhile, Cozmo attempted to tap the top of the cube: if participants managed to pull the cube away before Cozmo tapped the top, they earned a point; if Cozmo was faster and tapped the cube before it was removed, Cozmo earned a point. In *Quick Tap*, participants played a color matching game with Cozmo involving two cubes: at the beginning of the game, one cube was placed in front of the participant and the other was placed in front of Cozmo (i.e., Cozmo and the participant faced each other with the cubes in between). On a given trial, the cubes would light up with a specific color (either same or different colors). If both cubes showed the same color (e.g., both cubes light up with the color blue), both parties are asked to tap the top of the cube in front of them as fast as possible (i.e., Go-trial). However, if both cubes lit up in red or showed different colors, both parties were asked to refrain from tapping (i.e., Nogo-trial). Participants won the round if (i) they tapped faster than Cozmo on a Go-trial or (ii) correctly withheld their response on a Nogo-trial (only if Cozmo incorrectly taps its cube at the same time). Cozmo won the round if (i) it tapped faster than participants on a Go-trial or if participants do not tap on a Go-trial, (ii) if participants incorrectly tap on a Nogo-trial (only if Cozmo correctly withheld its response at the same time). Neither the participant nor Cozmo received a point if (i) both Cozmo and the participant failed to tap on a Go-trial and/or (ii) both Cozmo and the participant correctly withheld their response on a Nogo-trial. The order in which participants played *Keep Away* and *Quick Tap* was counterbalanced across participants. All of the games were prompted and initiated by the researcher by controlling the app via the tablet.

Participants in the **repeated interaction** condition interacted with Cozmo over the course of five days in their homes. They were free to choose (i) how they wanted to interact with Cozmo (all features and games in the app were possible) and (ii) for how long (with a minimum of 20 minutes every day). The five-day period started with participants coming to the laboratory to receive “their” Cozmo and to receive a short tutorial, which was supposed to familiarize them with the robot and the app. After providing consent, the researcher showed them how to use the app and the tablet to prompt Cozmo’s actions and games. After completing the tutorial, participants were asked to explore the app and play with Cozmo for 20 minutes in the lab (to make sure that all participants were able to use the platform on their own). Afterwards, they were instructed about the “at home” part of the study: they were told that they could choose how to engage with Cozmo and for how long each day, but that they should (i) interact with Cozmo for at least 20 minutes every day, and (ii) document their interactions (type and length) via an online questionnaire to ensure that they interacted with Cozmo daily for at least 20 minutes; no other constraints were provided, see **Table 1** and **Figure 3**. for the average duration per day and percentage of games played. Finally, participants were asked to come back to the lab on day 5 to return the robot and complete the gambling part of the study.

DAY	AVERAGE DURATION	SD
1	39.3	14.0
2	35.2	7.5
3	36.8	9.9
4	34.8	7.6
5	34.5	4.9

**Table 1** Average reported duration of interaction per day.  
 Note: Duration reported is in minutes.

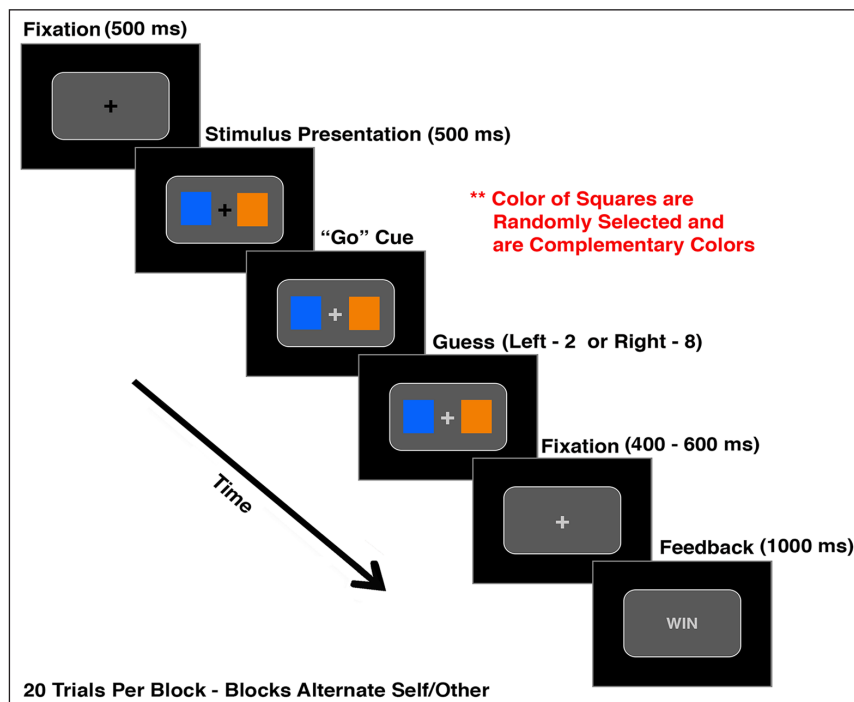


**Figure 3 Proportion of games played:** The graph illustrates participants’ reported proportion of games across all participants. The majority of the games played included “Quick tap”, “Memory match” and “Keep away”.

## GAMBLING TASK

After the interaction phase, all participants performed a gambling task that required them to either gamble for themselves (“Self”) or Cozmo (“Cozmo”) by determining which of two differently colored squares (e.g., blue and orange; presented on the horizontal midline of a screen, left and right of a centrally presented fixation cross) produced a winning outcome with a higher frequency. Participants completed 32 blocks (i.e., 16 for Self and 16 for Cozmo) total with a break at the mid-point of the experiment; each block consisted of 20 trials. Whether participants gambled for Self or Cozmo was announced at the beginning of each block by presenting either “Self” or “Cozmo” on the screen; “Self” and “Cozmo” blocks were randomized throughout the experiment. Although the colors of the squares changed after each block (together with the associated likelihood of producing a winning outcome), one color was always associated with a 60% chance of winning and the other color was associated with a 10% chance of winning.

The trial sequence is shown in **Figure 4**. Each trial started with the presentation of a central, black fixation cross for 500 ms. The two differently colored squares would then appear left and right of the fixation cross. Although the color scheme would be randomly selected for each block, the colors of the squares would always be complementary colors (e.g., orange and blue); the same two colors were used for all 20 trials within a given block. After the colored squares had been presented for 500 ms, the fixation cross would change its color from black to light grey to indicate to participants that they could now start to gamble (i.e., pick the square they believed was associated with the higher chance of winning). The squares remained on the screen until participants responded with either the “2” key (with their left index finger) to select the square on the left or the “8” key (with their right index finger) to select the square on the right. The probability of each colored square being presented on the left or right side of the screen was equi-probable. After participants submitted their response, feedback would be presented for 1000 ms to inform them about the gambling outcome (i.e., “WIN” or “LOSE”). If participants selected a square prior to the change of the color of the fixation cross, the trial was not included in the analyses. The inter-trial interval (ITI) was jittered between 400–600 ms.



**Figure 4 Trial sequence of the gambling task:** At the beginning of each trial, a black fixation cross was presented centrally for 500 ms, followed by the presentation of two differently colored squares - left and right of the fixation cross. After another 500 ms, the fixation cross changed its color from black to white, which indicated to participants that they could now make their choice by either pressing “2” for the left square or “8” for the right square. After participants made their selection, the fixation cross was presented again for a jittered time interval of 400 to 600 ms. Afterwards, participants received feedback if they lost (“LOSE”) or won (“WIN”) the trial. “Self” versus “Cozmo” was blocked in this experiment and was only shown once at the beginning of a block (not shown).

## ANALYSIS

The impact of different interaction modes on feedback processing during the gambling task was analyzed by examining participants’ learning performance over time (i.e., within a block) using a log-log growth curve mixed-model. The growth model examined differences in learning rates as a function of *Condition* (No- vs. One-time vs. Repeated interaction) and *Recipient* (Self vs. Cozmo); both factors were dummy coded. The growth model predicted whether participants chose the color associated with the higher probability of winning or the color associated with the lower probability of winning (i.e., a dichotomous outcome). Interactions between the dummy coded variables and the growth variable examined differences in learning behavior over time between the conditions. To construct the random-effects part of the log-log growth curve, we used a nested model comparison approach to compare four different models to a model that only included the intercept as a predictor, which was the reference of comparison. The models included a log-log growth model that varied the intercept for each participant, a log-log growth model that varied the intercept for each participant and each trial, a log-log growth model that varied the intercept for each participant, each trial and each block and finally, a log-log growth model that varied the intercept for each participant, each trial, each block and a random intercept for the actual score of the gambling task. The nested model comparison allows us to test if the models fit significantly different from the intercept-only model using a Chi-squared test. We then examine the fit indices of the models to tell which of



the significantly different models fits the data best. The fixed effects equation can be found in Equation 1 below.

$$\ln(\gamma) = \ln(\beta_0) + \ln(\beta_1 X_1) + \beta_2 X_2 + \beta_3 X_3 + \ln(\beta_4 X_1) X_2 + \ln(\beta_5 X_1) X_3 + \ln(\beta_6 X_1) X_2 X_3 \quad (1)$$

## RESULTS

The nested model comparison showed that the model that varied the intercept for each participant (2 (11) = 1204.52,  $p < .001$ , AIC = 35332, BIC = 35441, -LL = -17653), the model that varied the intercept for each participant and each trial (2 (1) = 4.98,  $p = .02$ , AIC = 35329, BIC = 35447, -LL = -17650), the model that varied the intercept for each participant, trial and block (2 (1) = 136.29,  $p < .001$ , AIC = 35195, BIC = 35321, -LL = -17582) and the model that varied the intercept for the each participant, trial, block and participants' overall score (2 (1) = 3722.04,  $p < .001$ , AIC = 31475, BIC = 31609, -LL = -15721) all fit significantly in comparison to the reference model (i.e., the intercept-only model). The fact that the model that varied the intercept for each participant, trial, block and participants' overall score had the lowest fit indices for all of AIC, BIC and the log likelihood suggests that it is the model that fits the data best. As such, we report the findings from this model.


Results of the growth model are summarized in **Table 2**; the most important main and interaction effects are highlighted here: (1) *Condition* (One-time vs. Repeated) showed a significant main effect of learning overall. This suggests that participants who were exposed to repeated interactions with Cozmo learned more overall across all blocks and trials by choosing the correct color more often compared than those in the one-time interaction condition:  $b = -.44$ ,  $SE = .16$ ,  $z(33039) = -2.65$ ,  $p < .01$ . This effect did not reach statistical significance for the One-time vs. No interaction contrast:  $b = -.34$ ,  $SE = .17$ ,  $z(33039) = -1.95$ ,  $p = .051$ . Although these mean differences are evident between the repeated and the one-time interaction conditions, they do not inform about learning rates.

VARIABLE	B	SE	Z VALUE	P
Intercept	-.28	.37	-.75	.45
Growth Curve	.45	.04	10.14	<.001
Condition:	-.44	.16	-2.65	<.01
<i>Repeated interaction vs. One-time interaction contrast</i>				
Condition:	-.34	.17	-1.95	.051
<i>No interaction vs. One-time interaction contrast</i>				
Recipient:	.08	.12	.62	.53
<i>Self vs. Cozmo contrast</i>				
Growth × Condition:	.28	.05	5.06	<.001
<i>Repeated interaction vs. One-time interaction contrast</i>				
Growth × Condition:	.28	.06	4.81	<.001
<i>No interaction vs. One-time interaction contrast</i>				
Growth × Recipient:	-.03	.05	-.58	.56
<i>Self vs. Cozmo contrast</i>				
Condition × Recipient:	.23	.17	1.32	.18
<i>Repeated interaction vs. One-time interaction for Self vs. Cozmo contrast</i>				
Condition × Recipient:	-.08	.19	-.43	.66
<i>No-Interaction vs. One-time interaction for Self vs. Cozmo contrast</i>				
Growth × Condition × Recipient:	-.17	.07	-2.18	.02
<i>Repeated interaction vs. One-time interaction for Self vs. Cozmo contrast</i>				
Growth × Condition × Recipient:	.01	.08	.2	.83
<i>No-Interaction vs. One-time interaction for Self vs. Cozmo contrast</i>				

**Table 2** Results of the log-log growth curve model.

*Note:* Significance testing was based on 33039 Degrees of Freedom. The “b” denotes the variable estimate and the “SE” denotes the Standard Error of the estimate.

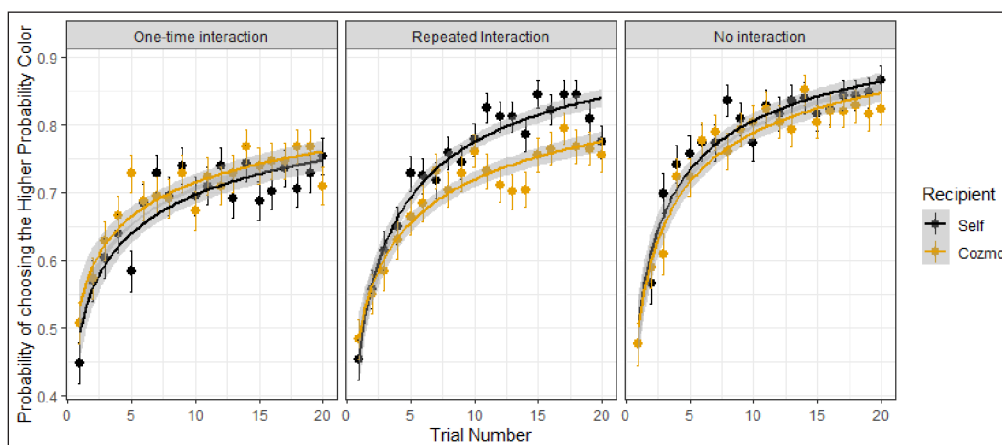
(2) *Condition* (No vs. One-time vs. Repeated) had a significant impact on learning rates (i.e., changes in likelihood of choosing the square with the higher chance of winning over the course of one block), as indicated by a significant growth function  $\times$  *Condition* interaction (for the repeated interaction vs. one-time interaction contrast:  $b = .28, SE = .05, z(33039) = 5.06, p < .001$ ; no interaction vs. one-time interaction contrast:  $b = .28, SE = .06, z(33039) = 4.81, p < .001$ ). Specifically, learning the contingency between color and chance of winning was fastest in the no interaction condition, followed by the repeated interaction condition, followed by the one-time interaction condition (**Figure 5**).



**Figure 5 Learning rates as a function of condition:**

Growth curve models allow us to examine differences in learning the higher probability outcome between conditions. The figure illustrates the significant *Growth X Condition* interactions, illustrating that learning rates were significantly different between the No-interaction and the One-time Interaction condition and the One-time interaction vs. the repeated interaction condition. The differences were such that those in the No interaction learned the fastest, followed by those who had repeated interactions and then those who had a one-time interaction with Cozmo. The points in the graph illustrate the percentage of picking the higher probability outcome.

(3) *Recipient* (Self vs. Cozmo) did not significantly impact learning rates across conditions, as indicated by an insignificant growth function  $\times$  *Recipient* interaction ( $b = -.03, SE = .05, z(33039) = -.58, p = .56$ ). However, learning rates were faster for Self vs. Cozmo in the repeated interaction condition compared to the one-time interaction condition, as indicated by a significant growth function  $\times$  *Recipient x Condition* interaction (for the repeated vs. one-time contrast:  $b = -.17, SE = .07, z(33039) = -2.18, p = .83$ ). There was no Self vs. Cozmo difference in learning rates in the no vs. one-time interaction conditions, as is indicated by a non-significant growth function  $\times$  *Recipient x Condition* interaction (for the no interaction vs. one-time interaction contrast:  $b = .01, SE = .08, z(33039) = .2, p = .83$ );<sup>3</sup> see **Figure 6**.



**Figure 6 Learning rates as a function of condition and recipient:**

The growth curve model showed significant learning differences in the *Growth x Recipient x Condition* interaction such learning rates were faster for Self vs. Cozmo in the repeated interaction condition compared to the one-time interaction condition. However, no differences were shown between the one-time interaction and the no interaction condition; points represent percentages of picking the higher probability outcome.

## DISCUSSION

The goal of the present study was to examine the impact of prior social interactions with a robot on feedback processing embedded in a paradigm with a dynamic (“online”) interaction component. In the social interaction conditions, participants were asked to familiarize

<sup>3</sup> We ran an additional analysis to examine the moderating effect of the reported time of interacting on learning rates for self vs. cozmo for those in the repeated-interaction condition, however, the model did not converge and as such, the estimates are deemed unstable. Therefore, we opted not to report the model.

themselves with the robot, Cozmo, by either interacting with it for 20 minutes (one-time interaction condition) or for five consecutive days (repeated interaction condition) before performing the gambling task. In the control condition (no social interaction), participants did not familiarize themselves with Cozmo before performing the gambling task but played the game *Simon Says* instead. In all conditions, participants gambled for themselves (i.e., reward goes to participant) or Cozmo (i.e., reward goes to Cozmo).

The findings reveal that the presented hybrid paradigm separating the dynamic (but unstructured) interaction component from the controlled (but structured) measurement component is suitable to detect differences in the effect of different familiarization durations on feedback-based learning without unnecessarily restricting the interaction component to a pre-scripted sequence. The results show that learning contingencies between the square color and chance of winning was fastest in the no interaction condition, followed by the repeated interaction condition, and the one-time interaction condition. Furthermore, a significant difference was found in feedback-based learning between “Self” and “Cozmo” in the repeated interaction condition showing that the paradigm was sensitive to the experimental manipulation of both independent variables (i.e., recipient and condition). This is important to note given that dynamic paradigms (i.e., paradigms, where interaction and measurement take place during a live interaction) often place constraints on the interaction (e.g., large number of repetitions required, interruptions due to measurements being taken) and are difficult to implement (e.g., sending EEG markers in real-time). This is particularly the case when objective behavioral (e.g., reaction times) or neurophysiological (e.g., event-related potentials) measures need to be obtained since recording such measures often requires pre-defined sequences or at least a certain number of critical events (e.g. a human/robot showing a certain behavior), which is often challenging to realize in a fully dynamic interaction and can have negative consequences in terms of comparability and data analysis (e.g., different number of critical events for each participant). The paradigm presented here shows that a “hybrid” solution, that is: the combination of an unscripted interaction component and a scripted measurement component, is worth considering as it does not place significant constraints on the social interaction, but allows behavioral and neurophysiological correlates of social cognition to be obtained with high internal validity. A few early studies pioneering this approach have shown its feasibility for behavioral (current study), neurophysiological (e.g., Abubshait et al., 2021; “hybrid” with an ERP component) and neural (e.g., Cross et al., 2019; “hybrid” with fMRI component) outcome measures.

Two findings require further discussion: (1) learning rates were lower in the one-time interaction than in the repeated interaction condition but highest in the no social interaction condition, and (2) significant differences in feedback-associated learning between “Self” and “Other” were observed only in the repeated interaction condition, but not the one-time or the non-social interaction conditions. These findings suggest potential effects of prior familiarization specific to the robot (i.e., specific effects on the “Self” vs. “Other” condition) might not manifest after a one-time familiarization period, but only when people are more familiarized with robots. While the current results do not allow us to explain all this variance in the data, certain hypotheses (which need to be tested in future experiments) that are based on prior research seem reasonable. For instance, it is possible that the differences in learning between the one-time and repeated social interaction condition is due to a “novelty effect” associated with the robot (Lemaignan et al., 2014) that increases participants’ engagement in the one-time interaction condition but has already washed out after repeated interactions. This interpretation would be in line with previous studies that have suggested that enriched environments modify response to feedback and negatively impact learning due to increased arousal and positive affect (Lohse et al., 2020). It would also be in line with other studies showing that artificial social agents like avatars or robots often cause a distraction in learning environments, which can hinder people’s ability to learn (Kennedy et al., 2015; Momen et al., 2016; Yadollahi et al., 2018). These interpretations are further supported by the fact that feedback-associated learning in the one-time interaction condition was attenuated for “Self” and “Cozmo” instead of being specific to the “Cozmo” condition (as would be expected for a familiarity effect). Future studies will have to be conducted to examine these hypotheses further.

Two challenges for future versions of this paradigm will be to (1) determine which modifications are necessary to replicate the Self-Other difference in feedback processing reported in

the original version of the paradigm, and (2) examine the effect of different familiarization protocols (e.g., different durations and/or interaction modes) on feedback processing. Both challenges are related to the fact that the Self-Other difference in feedback processing in the gambling task was not observed in two out of three conditions in the current experiment. Prior work suggests that better responses to feedback for Self vs. Other is explained by an increased intrinsic motivation to win when positive rewards are assigned to oneself versus an unknown player (Hassall et al., 2016; Wilhelm et al., 2019), and been shown to vanish when the “Other” is a known friend versus an unknown stranger (indicating that the “relationship” between the gambler and the recipient of a positive outcome might play a significant role in feedback processing; Leng & Zhou, 2010). Thus, one would have expected a significant Self-Other difference in the control condition (no familiarization), but not in the familiarization conditions.

While the current experiment does not allow any data-driven interpretations, we suspect that the reported findings might be impacted by the participants’ affective state since positive emotions have been shown to impact outcome monitoring (Bakic et al., 2014). Specifically, it is possible that the fact that the participants played engaging games (Simon Says and playing with Cozmo) in the control and one-time interaction (but not the repeated interaction) conditions immediately before performing the gambling task may have elevated their arousal and in turn impacted feedback processing. Additionally, worse learning outcomes specifically for the “Other” condition after repeated interaction would be in line with recent studies showing that repeated exposure to social robots negatively affects social cognitive processes (Abubshait & Wykowska, 2020). We suggest that reduced feedback monitoring for “Other” could be attributable to devaluation of the reward following repeated interactions with the robot. The current study also does not allow us to completely control for differences in reward valuation of winning a gift card for one’s self vs. upgrades for the robot. However, since prior work has failed to find differences in RewP amplitudes for “Self” versus “Other” in the one-time interaction condition (Abubshait et al., 2021), this suggests that the reward values for these two conditions were comparable after a short interaction. This interpretation is also supported by other work that has shown that RewP amplitudes are sensitive to the magnitude of the reward in the gambling task (Hassall et al., 2016). While it is possible that the reward values for “Self” and “Other” did not differ in the one-time interaction condition, it might be that changes in affect over the course of repeated interactions with the robot specifically impacted the degree to which rewards were valued for “Self” vs. “Other”.<sup>4</sup> Whether participants’ attitudes towards Cozmo were indeed negatively impacted in the repeated interaction condition and whether this potential decrease in positive emotions is strong enough to impact feedback processing warrants further investigation.

We would like to point out that the findings need to be interpreted with care due to a few limitations. First, although we did observe significant effects of all experimental manipulations on the obtained outcome measures, a sample size of 60 participants can be considered small for a three-factor mixed design; this makes it necessary to ensure that the reported findings can be replicated with larger sample sizes. Second, due to the lack of systematic investigations of social cognition in long-term human-robot interactions, it is difficult to empirically determine what the appropriate duration of a familiarization period is supposed to be in order to label an interaction “long-term”. While a duration of five days, as was chosen in the current study, is longer than is reported in most human-robot interaction studies, it is unclear whether this time frame is long enough to induce sustained changes in perceptions of and behaviors towards robots (and to call the interaction “long-term”). Given how important it is to investigate mechanisms of social cognition in human-robot interaction in naturalistic settings over representative periods of time (Dautenhahn, 2007; Giusti & Marti, 2006), it is essential for the field of social robotics to establish systematic standards (e.g., type of tasks and robots, interaction protocols, duration, etc.) of examination (e.g. see Chaudhury et al., 2020; *preprint*; for a tutorial on how to use the robot Cozmo for HRI research).

The reported findings have implications for the fields of social cognition and human-robot interaction. First, the results show that hybrid paradigms that separate the interaction and

---

<sup>4</sup> Please note that although it cannot be ruled out that differences in reward relevance (i.e., being rewarded with the chance of winning a gift card vs. winning upgrades for Cozmo) might play a role in the current paradigm, one would expect this to equally impact all three familiarity conditions and can therefore not explain Self-Other differences in learning in the repeated interaction condition only.

measurement components are sensitive to examine the impact of social interactions on social-cognitive parameters and might be considered as a pragmatic compromise if fully dynamic paradigms are not feasible. Second, the results indicate that the impact of repeated interactions with an embodied robot on social cognition is multidimensional and changes over time, which requires studies in human-robot interaction to incorporate a repeated measures component. Third, although it is to be expected that social robots will be able to display more sophisticated social behaviors in the future, it seems that current social robot platforms might not be sophisticated enough to engage adult participants positively over longer periods of time. This could lead to negative attitudes towards the robot once the novelty effect has vanished. Despite these challenges, being able to use affordable, easy to program and emotionally engaging robots to examine social cognition in human-robot interaction is an exciting new avenue and will open new possibilities for empirical research in the future.

## DATA ACCESSIBILITY STATEMENTS

Behavioral data are made available and can be found on [osf.io/6dfky/](https://osf.io/6dfky/).

## ETHICS AND CONSENT

This study was approved by and carried out in accordance with George Mason University's Internal Review Board (IRB).

## FUNDING INFORMATION


Part of this study was funded by the Human Factors and Applied Cognition Fellowship from the Human Factors and Applied Cognition program of the Psychology Department at George Mason University. The award was granted to Abdulaziz Abubshait in 2019.


## COMPETING INTERESTS

The authors have no competing interests to declare.

## AUTHOR AFFILIATIONS

**Abdulaziz Abubshait**  [orcid.org/0000-0001-5063-6355](https://orcid.org/0000-0001-5063-6355)  
George Mason University, US;  
Italian Institute of Technology, IT

**Craig G. McDonald**  [orcid.org/0000-0002-2901-9157](https://orcid.org/0000-0002-2901-9157)  
George Mason University, US

**Eva Wiese**  [orcid.org/0000-0002-0322-1250](https://orcid.org/0000-0002-0322-1250)  
Berlin Institute of Technology, US

## REFERENCES

- Abubshait, A., Beatty, P. J., McDonald, C. G., Hassall, C. D., Krigolson, O. E., & Wiese, E.** (2021). A win-win situation: Does familiarity with a social robot modulate feedback monitoring and learning? *Cognitive, Affective, & Behavioral Neuroscience*. DOI: <https://doi.org/10.3758/s13415-021-00895-9>
- Abubshait, A., & Wykowska, A.** (2020). Repetitive Robot Behavior Impacts Perception of Intentionality and Gaze-Related Attentional Orienting. *Frontiers in Robotics and AI*, 7. DOI: <https://doi.org/10.3389/frobt.2020.565825>
- Admoni, H., Bank, C., Tan, J., Toneva, M., & Scassellati, B.** (2011). Robot gaze does not reflexively cue human attention. *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*, 1983–1988.
- Bakic, J., Jepma, M., De Raedt, R., & Pourtois, G.** (2014). Effects of positive mood on probabilistic learning: Behavioral and electrophysiological correlates. *Biological Psychology*, 103, 223–232. DOI: <https://doi.org/10.1016/j.biopsycho.2014.09.012>
- Bartneck, C., & Reichenbach, J.** (2005). Subtle emotional expressions of synthetic characters. *International Journal of Human-Computer Studies*, 62(2), 179–192. DOI: <https://doi.org/10.1016/j.ijhcs.2004.11.006>
- Basteris, A., Nijenhuis, S. M., Stienen, A. H., Buurke, J. H., Prange, G. B., & Amirabdollahian, F.** (2014). Training modalities in robot-mediated upper limb rehabilitation in stroke: A framework for classification based on a systematic review. *Journal of NeuroEngineering and Rehabilitation*, 11(1),

111. DOI: <https://doi.org/10.1186/1743-0003-11-111>
- Bates, D., Mächler, M., Bolker, B., & Walker, S.** (2014). Fitting Linear Mixed-Effects Models using lme4. *Journal of Statistical Software*, 67(1). DOI: <https://doi.org/10.18637/jss.v067.i01>
- Bekele, E., Crittendon, J. A., Swanson, A., Sarkar, N., & Warren, Z. E.** (2014). Pilot clinical application of an adaptive robotic system for young children with autism. *Autism*, 18(5), 598–608. DOI: <https://doi.org/10.1177/1362361313479454>
- Birks, M., Bodak, M., Barlas, J., Harwood, J., & Pether, M.** (2016). Robotic Seals as Therapeutic Tools in an Aged Care Facility: A Qualitative Study. *Journal of Aging Research*, 2016, 1–7. DOI: <https://doi.org/10.1155/2016/8569602>
- Brainard, D. H.** (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4), 433–436. DOI: <https://doi.org/10.1163/156856897X00357>
- Capozzi, F., & Ristic, J.** (2020). Attention AND mentalizing? Reframing a debate on social orienting of attention. *Visual Cognition*, 1–9. DOI: <https://doi.org/10.1080/13506285.2020.1725206>
- Charrier, L., Rieger, A., Galdeano, A., Cordier, A., Lefort, M., & Hassas, S.** (2019). The RoPE Scale: A Measure of How Empathic a Robot is Perceived. *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 656–657. DOI: <https://doi.org/10.1109/HRI.2019.8673082>
- Ciardo, F., Beyer, F., De Tommaso, D., & Wykowska, A.** (2020). Attribution of intentional agency towards robots reduces one's own sense of agency. *Cognition*, 194, 104109. DOI: <https://doi.org/10.1016/j.cognition.2019.104109>
- Cole, G. G., Skarratt, P. A., & Kuhn, G.** (2016). Real Person Interaction in Visual Attention Research. *European Psychologist*, 21(2), 141–149. DOI: <https://doi.org/10.1027/1016-9040/a000243>
- Cross, E. S., Riddoch, K. A., Pratts, J., Titone, S., Chaudhury, B., & Hortensius, R.** (2019). A neurocognitive investigation of the impact of socializing with a robot on empathy for pain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 374(1771), 20180034. DOI: <https://doi.org/10.1098/rstb.2018.0034>
- Czeszumski, A., Ehinger, B. V., Wahn, B., & König, P.** (2019). The Social Situation Affects How We Process Feedback About Our Actions. *Frontiers in Psychology*, 10. DOI: <https://doi.org/10.3389/fpsyg.2019.00361>
- de Bruijn, E. R. A., Miedl, S. F., & Bekkering, H.** (2011). How a co-actor's task affects monitoring of own errors: Evidence from a social event-related potential study. *Experimental Brain Research*, 211(3), 397. DOI: <https://doi.org/10.1007/s00221-011-2615-1>
- de Bruijn, E. R. A., Ruissen, M. I., & Radke, S.** (2017). Electrophysiological correlates of oxytocin-induced enhancement of social performance monitoring. *Social Cognitive and Affective Neuroscience*, 12(10), 1668–1677. DOI: <https://doi.org/10.1093/scan/nsx094>
- Frischen, A., Bayliss, A. P., & Tipper, S. P.** (2007). Gaze cueing of attention: Visual attention, social cognition, and individual differences. *Psychological Bulletin*, 133(4), 694–724. DOI: <https://doi.org/10.1037/0033-2909.133.4.694>
- Hackel, L. M., Looser, C. E., & Van Bavel, J. J.** (2014). Group membership alters the threshold for mind perception: The role of social identity, collective identification, and intergroup threat. *Journal of Experimental Social Psychology*, 52, 15–23. DOI: <https://doi.org/10.1016/j.jesp.2013.12.001>
- Hassall, C. D., Silver, A., Turk, D. J., & Krigolson, O. E.** (2016). We are more selfish than we think: The endowment effect and reward processing within the human medial-frontal cortex. *Quarterly Journal of Experimental Psychology*, 69(9), 1676–1686. DOI: <https://doi.org/10.1080/17470218.2015.1091849>
- Hobson, N. M., & Inzlicht, M.** (2016). The mere presence of an outgroup member disrupts the brain's feedback-monitoring system. *Social Cognitive and Affective Neuroscience*, 11(11), 1698–1706. DOI: <https://doi.org/10.1093/scan/nsw082>
- Holroyd, C. B., & Coles, M. G. H.** (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, 109(4), 679–709. DOI: <https://doi.org/10.1037/0033-295X.109.4.679>
- Hsieh, T.-Y., Chaudhury, B., & Cross, E. S.** (2020). Human-Robot Cooperation in Prisoner Dilemma Games: People Behave More Reciprocally than Prosocially Toward Robots. *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 257–259. DOI: <https://doi.org/10.1145/3371382.3378309>
- Insel, T. R., & Fernald, R. D.** (2004). HOW THE BRAIN PROCESSES SOCIAL INFORMATION: Searching for the Social Brain. *Annual Review of Neuroscience*, 27(1), 697–722. DOI: <https://doi.org/10.1146/annurev.neuro.27.070203.144148>
- Kennedy, J., Baxter, P., & Belpaeme, T.** (2015). The Robot Who Tried Too Hard: Social Behaviour of a Robot Tutor Can Negatively Affect Child Learning. *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 67–74. DOI: <https://doi.org/10.1145/2696454.2696457>
- Kimura, K., & Katayama, J.** (2016). Cooperative context is a determinant of the social influence on outcome evaluation: An electrophysiological study. *International Journal of Psychophysiology*, 100, 28–35. DOI: <https://doi.org/10.1016/j.ijpsycho.2015.12.005>
- Koban, L., Corradi-Dell'Acqua, C., & Vuilleumier, P.** (2013). Integration of Error Agency and Representation of Others' Pain in the Anterior Insula. *Journal of Cognitive Neuroscience*, 25(2), 258–

272. DOI: [https://doi.org/10.1162/jocn\\_a\\_00324](https://doi.org/10.1162/jocn_a_00324)
- Koban, L., Pourtois, G., Bediou, B., & Vuilleumier, P.** (2012). Effects of social context and predictive relevance on action outcome monitoring. *Cognitive, Affective, & Behavioral Neuroscience*, 12(3), 460–478. DOI: <https://doi.org/10.3758/s13415-012-0091-0>
- Krigolson, Olave E., Hassall, C. D., Balcom, L., & Turk, D.** (2013). Perceived ownership impacts reward evaluation within medial-frontal cortex. *Cognitive, Affective, & Behavioral Neuroscience*, 13(2), 262–269. DOI: <https://doi.org/10.3758/s13415-012-0144-4>
- Kompatsiari, K., Ciardo, F., Tikhonoff, V., Metta, G., & Wykowska, A.** (2018). On the role of eye contact in gaze cueing. *Scientific Reports*, 8(1), 17842. DOI: <https://doi.org/10.1038/s41598-018-36136-2>
- Krigolson, Olav E., Pierce, L. J., Holroyd, C. B., & Tanaka, J. W.** (2009). Learning to Become an Expert: Reinforcement Learning and the Acquisition of Perceptual Expertise. *Journal of Cognitive Neuroscience*, 21(9), 1833–1840. DOI: <https://doi.org/10.1162/jocn.2009.21128>
- Lefkeli, D., Ozbay, Y., Gürhan-Canlı, Z., & Eskenazi, T.** (2020). Competing with or Against Cozmo, the Robot: Influence of Interaction Context and Outcome on Mind Perception. *International Journal of Social Robotics*. DOI: <https://doi.org/10.1007/s12369-020-00668-3>
- Lemaignan, S., Fink, J., & Dillenbourg, P.** (2014). The Dynamics of Anthropomorphism in Robotics. 2014 9th ACM/IEEE International Conference on Human-Robot Interaction (HRI), 226–227. DOI: <https://doi.org/10.1145/2559636.2559814>
- Lemoine, J. E., & Roland-Lévy, C.** (2017). The effect of the presence of an audience on risk-taking while gambling: The social shield. *Social Influence*, 12(2–3), 101–114. DOI: <https://doi.org/10.1080/15534510.2017.1373697>
- Leng, Y., & Zhou, X.** (2010). Modulation of the brain activity in outcome evaluation by interpersonal relationship: An ERP study. *Neuropsychologia*, 48(2), 448–455. DOI: <https://doi.org/10.1016/j.neuropsychologia.2009.10.002>
- Lohse, K. R., Miller, M. W., Daou, M., Valerius, W., & Jones, M.** (2020). Dissociating the contributions of reward-prediction errors to trial-level adaptation and long-term learning | Elsevier Enhanced Reader. *Biological Psychology*, 149(107775). DOI: <https://doi.org/10.1016/j.biopsycho.2019.107775>
- McNeill, D., & Kennington, C.** (2019, June 22). Predicting Human Interpretations of Affect and Valence in a Social Robot. *Robotics: Science and Systems XV*. Robotics: Science and Systems 2019. DOI: <https://doi.org/10.15607/RSS.2019.XV.041>
- Momen, A., Sebrechts, M. M., & Allaham, M. M.** (2016). Virtual Agents as a Support for Feedback-Based Learning. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 60(1), 1780–1784. DOI: <https://doi.org/10.1177/1541931213601407>
- Pelikan, H. R. M., Broth, M., & Keevallik, L.** (2020). “Are You Sad, Cozmo?”: How Humans Make Sense of a Home Robot’s Emotion Displays. *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 461–470. DOI: <https://doi.org/10.1145/3319502.3374814>
- Radke, S., de Lange, F. P., Ullsperger, M., & de Bruijn, E. R. A.** (2011). Mistakes that affect others: An fMRI study on processing of own errors in a social context. *Experimental Brain Research*, 211(3), 405–413. DOI: <https://doi.org/10.1007/s00221-011-2677-0>
- Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J.-F., Breazeal, C., Crandall, J. W., Christakis, N. A., Couzin, I. D., Jackson, M. O., Jennings, N. R., Kamar, E., Kloumann, I. M., Larochelle, H., Lazer, D., McElreath, R., Mislove, A., Parkes, D. C., Pentland, A. ‘Sandy,’ ... Wellman, M.** (2019). Machine behaviour. *Nature*, 568(7753), 477–486. DOI: <https://doi.org/10.1038/s41586-019-1138-y>
- Rigoni, D., Polezzi, D., Rumiati, R., Guarino, R., & Sartori, G.** (2010). When people matter more than money: An ERPs study. *Brain Research Bulletin*, 81(4), 445–452. DOI: <https://doi.org/10.1016/j.brainresbull.2009.12.003>
- Schilbach, L., Timmermans, B., Reddy, V., Costall, A., Bente, G., Schlicht, T., & Vogeley, K.** (2013). Toward a second-person neuroscience. *Behavioral and Brain Sciences*, 36(4), 393–414. DOI: <https://doi.org/10.1017/S0140525X12000660>
- Schultz, W.** (2017). Reward prediction error. *Current Biology*, 27(10), R369–R371. DOI: <https://doi.org/10.1016/j.cub.2017.02.064>
- Sharkey, A. J. C.** (2016). Should we welcome robot teachers? *Ethics and Information Technology*, 18(4), 283–297. DOI: <https://doi.org/10.1007/s10676-016-9387-z>
- Simon, D., Becker, M. P. I., Mothes-Lasch, M., Miltner, W. H. R., & Straube, T.** (2014). Effects of social context on feedback-related activity in the human ventral striatum. *NeuroImage*, 99, 1–6. DOI: <https://doi.org/10.1016/j.neuroimage.2014.05.071>
- Tapus, A., Mataric, M. J., & Scassellati, B.** (2007). Socially assistive robotics [Grand Challenges of Robotics]. *IEEE Robotics Automation Magazine*, 14(1), 35–42. DOI: <https://doi.org/10.1109/MRA.2007.339605>
- Van Meel, C. S., & Van Heijningen, C. A. A.** (2010). The effect of interpersonal competition on monitoring internal and external error feedback. *Psychophysiology*, 47(2), 213–222. DOI: <https://doi.org/10.1111/>

- Warren, Z. E., Zheng, Z., Swanson, A. R., Bekele, E., Zhang, L., Crittendon, J. A., Weitlauf, A. F., & Sarkar, N.** (2015). Can Robotic Interaction Improve Joint Attention Skills? *Journal of Autism and Developmental Disorders*, 45(11), 3726–3734. DOI: <https://doi.org/10.1007/s10803-013-1918-4>
- Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L., François, R., Golemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T., Miller, E., Bache, S., Müller, K., Ooms, J., Robinson, D., Seidel, D., Spinu, V., ... Yutani, H.** (2019). Welcome to the Tidyverse. *Journal of Open Source Software*, 4(43), 1686. DOI: <https://doi.org/10.21105/joss.01686>
- Wiese, E., Metta, G., & Wykowska, A.** (2017). Robots as intentional agents: Using neuroscientific methods to make robots appear more social. *Frontiers in Psychology*, 8. DOI: <https://doi.org/10.3389/fpsyg.2017.01663>
- Wilhelm, R. A., Miller, M. W., & Gable, P. A.** (2019). Neural and Attentional Correlates of Intrinsic Motivation Resulting from Social Performance Expectancy. *Neuroscience*, 416, 137–146. DOI: <https://doi.org/10.1016/j.neuroscience.2019.07.039>
- Yadollahi, E., Johal, W., Paiva, A., & Dillenbourg, P.** (2018). When deictic gestures in a robot can harm child-robot collaboration. *Proceedings of the 17th ACM Conference on Interaction Design and Children*, 195–206. DOI: <https://doi.org/10.1145/3202185.3202743>

## TO CITE THIS ARTICLE:

Abubshait, A., McDonald C. G., & Wiese, E. (2021). Examining Social Cognition with Embodied Robots: Does Prior Experience with a Robot Impact Feedback-associated Learning in a Gambling Task? *Journal of Cognition*, 4(1): 28, pp. 1–16. DOI: <https://doi.org/10.5334/joc.167>

Submitted: 01 November 2020

Accepted: 30 April 2021

Published: 31 May 2021

## COPYRIGHT:

© 2021 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.

*Journal of Cognition* is a peer-reviewed open access journal published by Ubiquity Press.