Investigating the effect of heart rate synchrony on prosocial behaviour towards a social robot

Katie A. Riddoch¹, Emily S. Cross^{1,2}

¹Institute of Neuroscience and Psychology, University of Glasgow, Glasgow, Scotland

²Department of Cognitive Science, Macquarie University, Sydney, Australia

Abstract

Robots are being designed to alleviate the burden of social isolation and loneliness, particularly for older adults for whom these issues are more widespread. While good intentions underpin these developments, the reality is that many of these robots are abandoned within a short period of time. To encourage the longer-term use and utility of such robots, researchers are exploring ways to increase robot likability and facilitate attachment. Results from experimental psychology suggest that interpersonal synchrony (the overlap of movement/sensation between two agents) increases the extent to which people like one another. To investigate the possibility that synchrony could facilitate people's liking towards a robot, we undertook a between-subjects experiment in which participants interacted with a robot programmed to illuminate at the same rate, or 20% slower, than their heart rate. To quantify the impact of cardio-visual synchrony on prosocial attitudes and behaviours toward this robot, participants completed self-report questionnaires, a gaze-cueing task, and were asked to strike the robot with a mallet. Contrary to hypotheses, results revealed no differences in self-reported liking of the robot, gaze cueing effects, or the extent to which participants hesitated to hit the robot between the synchronous and asynchronous groups. The qualitative data collected in semi-structured interviews provided rich insights, though, and call into question the use of the broad "Likability" measurement, and the appropriateness of the 'hesitance to hit' paradigm as a measure of attachment to a robotic system.

Key Words Social robot · Human–robot interaction · Synchrony · Prosocial Attitudes · Heart Rate · Cardiovisual

1 Introduction

1.1 Companion Robots

Rapidly increasing elderly populations around the world, paired with an insufficient number of carers, has led to developments in "gerontechnologies" - devices designed to improve the health and wellbeing of older adults (Broekens, Heerink & Rosendal, 2009; Fozard, Rietsema, Bouma & Graafmans, 2000). Examples include location monitors, medication reminders, and fall detection systems (Bharucha et al., 2009). It is proposed that gerontechnologies will allow people to retain their independence and remain in their own homes as they age, as opposed to entering into a care home facility (Benefield & Holtzclaw, 2014; Piau et al., 2014). Although so-called "aging in place" is a desire reported by many people (Rantz et al, 2005), living alone is also associated with social isolation (lack of contact with others) and loneliness (distress that

results from the discrepancy between one's desired and actual social relationships; Pinquart & Sörenson, 2003). Longitudinal studies on older adults have demonstrated that feelings of loneliness are linked to more symptoms of depression (Cacioppo, 2006), reduced physical activity (Hawkey, Thisted & Cacioppo, 2009), and impaired cognition (O'Lunaigh et al., 2012). A study of 1604 elderly individuals also found that those who reported being lonely on average experienced greater difficulty with walking, stair climbing, and completing activities of everyday living (e.g. bathing, dressing, etc.) compared to those who did not report being lonely (Perissinotto, Cenzer & Covinsky (2012).

In an attempt to reduce loneliness and the myriad of associated health problems,

companies are in the process of developing ‘companion robots’ - machines designed to be engaging, comforting, and respond to the user in an intuitive manner (Broekens, Heerink & Rosendal, 2009; Young, Hawkins, Sharlin & Igarashi, 2009). As well as having capabilities of functionally assistive robots (e.g. carrying food and fetching medication), a further aim of companion robots is to design these machines in such a way that they can connect with a user socially. The need for such robots to be developed has been further reinforced by social distancing measures introduced as part of the recent COVID-19 pandemic, with prominent roboticists highlighting the role of social robots as ideal agents for providing care and companionship when contact with other people brings increased infection risk (Yang et al., 2020). One such robot that is already being introduced to social spheres, and which a number of organisations hope to deploy in such situations, is the “Pepper” robot (Softbank Robotics; Fig 1). This robot has the appearance of a friendly humanoid, and can recognise and speak to individuals, and is already being trialled in hospital and service industry contexts (Foster et al, 2016; Niemelä, Heikkilä & Lammi, 2017; Tanioka, 2019).

Fig 1. The Pepper robotic system

Studies using questionnaires have identified that healthy older adults are generally open to the idea of socially assistive robots

(Louie, McColl & Nejat, 2014), but Gurley and Norcio (2009) identified that social robots introduced into the homes of older adults were frequently abandoned within one year. While this finding is concerning, other researchers have described how most commercially available robots struggle to maintain the interest of their users even beyond 10 hours (Tanaka, Cicourel & Movellan, 2007).

For robots to have a long-term positive impact, it is vital to develop a clearer understanding of the features or behaviours that might facilitate bonding and attachment to robotic systems. As robots are often ascribed intentions and treated like social entities by people of all ages in a variety of laboratory and naturalistic experiments (Hortensius & Cross, 2018; Hortensius, Hekele & Cross, 2018; Wykowska, Chaminade & Cheng, 2016), research from psychology and the cognitive sciences has the potential to play a significant role in characterising which factors and attributes of robots will lead to the long term acceptance and enjoyment by human users.

1.2 Interpersonal Synchrony

One of the factors that researchers are examining in terms of building stronger bonds between humans and robots is interpersonal synchrony (e.g., Henschel & Cross, 2020; Lehmann et al., 2015, Mörtl, Lorenz, & Hirche, 2014). Interpersonal synchrony refers to the overlap of movement/sensation in time or form (for example, when we tap in sync, compared to out of sync, with another person; Hove & Risen, 2009). Experimental studies have demonstrated that our perception of, and behaviour towards, another person is altered when we engage in movement synchrony (Hove & Risen, 2009; Mogan, Fischer & Bulbulia, 2017; Rennung & Göritz, 2016). Specifically, the meta-analyses of Mogan, Fischer and Bulbulia (2017) and Rennung and Göritz (2016) identified that experimentally-induced movement synchrony had significant positive effects on prosocial behaviours (e.g. increased donating in trust and public goods games) and perceived social bonding between people (e.g. increased subjective liking and rapport).

Experts disagree on the mechanism underlying the positive effects of such synchrony (Mogan, Fischer & Bulbulia, 2017), however theories include that they are the result of: 1) a perceived blurring of the “self” and

“other” (Hove, 2009), 2) the allocation of attention to “group-centred cognition” (MacCrae et al, 2008), and 3) shared emotionality, termed “collective effervescence” (Durkheim, 1912/1995). Studies regarding the effect of interpersonal synchrony on object perception offer support for the theory of Hove (2009), as they demonstrate that participants exhibit increased self-identification with objects and images which illuminate in a synchronous manner relative to their heart rate (Sel, Azevedo & Tsakiris, 2017; Suzuki et al, 2013).

Suzuki et al (2013) conducted a version of the classic rubber hand illusion in which participants sit with their arms outstretched in front of them on a table, with one hand concealed by a sheet. In place of their concealed hand is a rubber hand which flashes in a manner synchronous, or asynchronous, relative to the individual’s heart rate. The study demonstrated that when the hand illuminated in synchrony with their heart rate, participants felt a stronger overlap between the “self” and the rubber hand. Sel, Azevedo and Tsakiris (2017) also investigated the effects of cardio-visual synchrony on self-other overlap, but instead using images of faces. In their study, participants were shown images of their own face morphed with the faces of others, that were illuminating in a manner synchronous or asynchronous relative to their heart rate. They found a similar result to Suzuki and colleagues - that participants had increased self-identification for the faces which pulsed in a synchronous manner to their heartbeat, compared to an asynchronous manner (Sel, Azevedo and Tsakiris, 2017). In addition to supporting the theory of Hove (2009), these studies indicate that interpersonal synchrony can not only alter our perception of human beings we encounter, but also our perception of inanimate objects and body parts. Given these findings, it is no surprise that the positive effects of synchrony are also demonstrated in human-robot interaction (HRI) research.

In an attempt to increase the liking and acceptance of robotic systems, researchers have experimentally-induced movement synchrony between individuals and robots (e.g., Henschel & Cross, 2020; Lehmann et al., 2015, Mörtl, Lorenz, & Hirche, 2014). To do so, the movement of the participant is monitored, and this information is relayed to the robotic system in real time. The robot can then react and adapt

to the changing movement of the individual. Mörtl, Lorenz and Hirche (2014) found that 11 out of 12 participants perceived such reactivity in a robotic arm as “pleasant”, and stated that it was “a nice feeling” and it made the robot appear “lively”. Upon implementing such reactivity into a Care-O-bot3, Lehmann et al (2015) also reported positive results of a synchrony manipulation on participant perceptions. Specifically, participants in this study perceived the system as more likable and intelligent when it engaged in a synchronous, compared to asynchronous body orientation relative to their own.

However, not all studies that have manipulated interpersonal synchrony between humans and robots report positive impacts on participants’ prosociality toward robots. In an experiment using the Pepper robot, Henschel and Cross (2020) found that when participants completed a synchronous compared to asynchronous drawing task with a robot, this had no effect on perceived likability of the robot, nor how many questions the participants asked the social robot (another type of prosocial behaviour termed “social motivation”). The lack of significant findings raises questions about the robustness of synchrony manipulations, and the authors also suggest that subtlety of the manipulation and the task-oriented nature of the task might further dilute impacts of synchrony manipulations (Henschel & Cross, 2020).

On balance, the studies reviewed here suggest that synchrony can have positive effects on our perception of other people, inanimate objects, and certain robotic systems. They also demonstrate that the effects of synchrony are not limited to the perception of synchronous physical movement, and that observing something illuminating in synchrony with our heart rate (termed “cardio-visual synchrony”) can also alter our perception of that object (Sel, Azevedo & Tsakiris, 2017; Suzuki et al, 2013). Given these findings, in the current study we explore whether cardio-visual synchrony might generalise to the way which we perceive and relate to a social robot.

2 The Experiment Overview

To investigate the effect of cardio-visual synchrony on the perception of a robotic agent, here we monitored the heart rate of the

participant, using a wrist-based heart monitor, and relay this information to the robotic system in real time. As a result, the robot's shoulder lights will illuminate in a manner synchronous (at the same rate) or asynchronous (20% slower) relative to the participants heart rate. If effective, this intervention would be relatively easy to implement in robotic systems used in social contexts with human users, as it involves minimal programming, uses a commercially available heart rate monitor, and can operate independently of the other behaviours of the robot. That is, the movement, speech, and other existing functions of the robot are unaffected by the presence of the flashing lights.

To determine whether cardio-visual synchrony leads participants to perceive a Pepper robot as more likable and behave in a more prosocial manner towards it, we used a between-subjects design and a number of qualitative and behavioural measures.

2.1 Likability Measure

In 2009, Bartneck and colleagues released the "Godspeed Questionnaire Series" (GQS) - a validated questionnaire with scales that could probe an individual's perception of a robot as 1) anthropomorphic, 2) animate, 3) likable, 4) intelligent, and 5) safe (Bartneck, Kulić, Croft & Zoghbi, 2009). Each scale was composed of 5 semantic differential scales (each containing two "anchors"). For example, the likability scale contains 5 semantic differential scales (each containing two "anchors") - 1) Dislike - Like, 2) Unfriendly - Friendly, 3) Unkind - Kind, 4) Unpleasant - Pleasant, and 5) Awful - Nice. The participant selects the number which best represents their opinion of the robot (experiment-dependent).

In the discussion of the original paper, the authors invited psychologists to develop and further validate the scales, yet over ten years later, the GQS remains the same, and has been used in at least 69 studies spanning numerous countries (Wiese and Bartneck, 2015). The overlapping nature of the "Anthropomorphism" and "Animacy" scales has been criticised (Carpinella et al., 2017) however the "Likability" scale boasts high internal validity, with Cronbach's alpha values ranging from 0.842 to 0.923 in experiments using the scale (Bartneck, Kanda, Ishiguro, & Hagita, 2007, 2009). The high internal validity of the scale, paired with the benefits of the semantic differential format (Friborg,

Martinussen & Rosenvinge, 2006), motivated our decision to use the Likability scale of the GSC to probe liking towards the robot in this experiment.

2.2 Hesitance to Hit Paradigm

To probe the extent to which robots are perceived as social agents as opposed to objects, researchers have asked participants to 'harm' robotic agents. For example, previous experimenters have asked people to administer increasing levels of electric shocks to a robot (Bartneck et al., 2005), turn a robot off and "wipe its memory" (Bartneck et al., 2007), and hit a robotic animal with a hammer (Darling, Nandy & Breazeal, 2015). The experimenter takes measures such as how many times a participant will actually hit a robot, the number of pieces it was broken into, and the amount of time between being given the instruction and compliance (termed "hesitation"). The rationale used is that maximal hesitation and minimal damage reflects empathy for the robot (Darling, Nandy & Breazeal, 2015), and perhaps even perception that the robot they are interacting with is a social agent with awareness and the capacity to feel pain (Bartneck & Hu, 2018).

Darling, Nandy and Breazeal (2015) demonstrated that individuals with high trait empathy hesitate for longer, suggesting an empathic component to the reaction time. The suggestion that hesitation reflects perceived sentience is not supported by evidence however, and Bartneck and Hu (2008) admit that this idea was included in an attempt to "stimulate discussion". As a result, the question remains - what are hesitation to harm/hit paradigms actually measuring?

In this experiment we used the hesitation to hit paradigm to yield an objective measure of harm compliance, in addition to a semi-structured post-experiment interviews, to probe why participants hesitated. By combining the reaction time data with qualitative insights, we believe we can better understand what the paradigm is measuring, and potentially gain insight into differences in empathy, liking, perceived sentience, and attachment to the robot, as influenced by the cardiovisual manipulation.

2.3 Gaze Cueing Task

A final objective measure we use to probe social perception and possibly feelings of

prosociality in the current study is the gaze cueing task. The direction of an individual's eye gaze offers insight into their internal mental state (e.g. their desire to interact with a specific object), as well as salient information in the surrounding environment (e.g. the presence of an approaching predator) (Martini, Buzzell & Wiese, 2015). As a result, it is no surprise that neurotypical humans shift their attention in the same direction as averted eyes (Admoni et al., 2011; Driver et al., 1999; Friesen & Kingstone, 1998).

To better understand the mechanisms underlying visual attention and gaze-following behaviours, the Posner (1980) attention cueing paradigm was modified, and the "gaze cueing" paradigm was developed (Driver et al., 1999; Friesen & Kingstone, 1998). In the now well-established gaze cueing protocol, the participant sees a centrally-presented pair of eyes gazing to the left or right side of the computer screen. A target (e.g. a letter or shape) then appears on the left or right side of the screen. "Congruent" or "valid" trials refer to those in which the target appears where the eyes are gazing (e.g. eyes gazing to the right, followed by target on the right). In contrast, when the target appears in the opposite direction to the eyes, these are termed "incongruent" or "invalid" trials. Studies have demonstrated that on average, participants are faster to respond to the target during congruent compared to incongruent trials, due to the misdirection of attention during incongruent trials (Friesen & Kingstone, 1998).

Such "attention orienting" was believed to be reflexive, resulting from our visual attention being captured and driven by the movement of the eyes ("bottom-up" processing; Driver et al., 1999; Bayliss, Pellegrino & Tipper, 2004), however new evidence suggests that our knowledge and beliefs can modulate the effect ("top-down" modulation; Teufel et al., 2010; Teufel et al., 2009; Morgan, Freeth & Smith, 2018; Wiese et al., 2012). In a series of gaze cueing experiments, Teufel and colleagues (2009, 2010) modified the paradigm to include the addition of goggles on the face stimuli. Results indicated that when the observer was believed to be wearing opaque goggles (unable to see), compared to transparent goggles (able to see), reflexive gaze following was reduced. This finding led to the authors' conclusion that our

perception of another person as knowledgeable, intentional, and ultimately as having a "mind" (as in the "seeing" condition) is crucial for reflexive gaze following.

Wiese and colleagues also found reduced gaze cueing effects when robotic faces were believed to be unintentional, compared to not, suggesting that the effect extends beyond our perception of human gaze. To ensure these findings were not the result of study design, and that the effect generalises to other types of tasks, Morgan, Freeth and Smith (2018) conducted a variation of the experiment using a different design, instruction, and task. The original results were replicated, supporting the suggestion that mind attribution and top-down processes can modulate reflexive gaze cueing.

The ability of the gaze cueing paradigm to apparently distinguish between the perception of stimuli as having a mind or not lends itself well to this current study. By implementing a gaze cueing paradigm in our experiment (using images of human faces, robot faces, and arrows) we can, arguably, gain further insight into the extent to which participants perceive the robot they interact with as a social agent compared to a mindless object. By conducting the task before and after interacting with the robot, we can also explore the malleability of such perceptions.

2.4 Measures Summary

To summarise, in this experiment we use three main measures. The Likability scale of the Godspeed Series Questionnaire is used to determine differences in perceived likability between the asynchronous and synchronous groups. After the interaction, participants are asked to hit the robot and reaction time differences between the groups is compared. To determine what the hesitance reflects, we administer a semi-structured interview at the end of the study asking participants why they hesitated and what they were thinking and feeling during the task. Finally, to gain insight regarding the perception of the robot as a social agent, opposed to a mindless object, the gaze cueing paradigm is administered before and after interacting with the robot.

2.5 Research Questions

The current study was designed to address three primary preregistered predictions.

Given the literature suggesting that interpersonal synchrony increases likability (See 1.2 Interpersonal Synchrony), we predict that the Synchronous group, compared to the Asynchronous group, will perceive the robot as more likable. We hypothesise that this difference will be reflected in the ratings on the validated ‘Liking’ scale of the Godspeed questionnaire - with the Synchronous group scoring higher than the Asynchronous group.

We propose that the increased liking will also be reflected in how long participants hesitate after being asked to hit the robot. Specifically, we predict that the Synchronous group, compared to the Asynchronous group, will hesitate for longer after being asked to hit the robot with a mallet. We also predict that more individuals in the Synchronous group, compared to the Asynchronous group, will refuse to hit the robot.

To investigate whether perceptions of the robot as a social agent differ between the two groups, we will also explore the Gaze Cueing data collected. On the basis of experiments suggesting that mind-perception modulates gaze cueing effects (Teufel et al., 2010; Teufel et al., 2009; Morgan, Freeth & Smith, 2018; Wiese et al., 2012), we anticipate that participants will be slower to respond to human faces, compared to the arrows, replicating the basic gaze cueing effect. We also expect that participants in the synchronous group will exhibit slower reaction times to the robot stimuli, compared to the asynchronous group. This is on the basis of research suggesting that a synchronous agent is perceived as more “like me” and as a social being opposed to an object as a result (Sel, Azevedo & Tsakiris, 2017; Suzuki et al, 2013)

3 Methodology

3.1 Preregistration & Ethics

Prior to data collection, all manipulations, measures, and the sample size justification and main hypotheses were pre-registered on the Open Science Framework (OSF):https://osf.io/d7c8t/?view_only=6198943ab9784339ac796c69fc4460c1. Consistent with recent proposals (Simmons, Nelson, & Simonsohn, 2011, 2012), we report all manipulations and all measures in the study. In addition, following open science initiatives

(Munafò et al., 2017), the data, stimuli, and analysis code associated with this study are freely available on the Open Science Framework. By making the data available, we enable others to pursue tests of alternative hypotheses, as well as more exploratory analyses. All study procedures were approved by the College of Science and Engineering Ethics Committee (University of Glasgow, Scotland) – approval number 300180265.

3.2 Sample Size Justification

Due to the time and resources associated with recruiting a sample including older individuals, the decision was made to use Bayesian Sequential Hypothesis Testing as outlined by Best, Barsalou and Papies (2018). To determine our minimum and maximum sample size, we undertook two power analyses in G*Power3.1. Both were undertaken on the basis of an independent T-Test, comparing scores on the Likeability scale of the Godspeed questionnaire (Bartneck et al., 2009).

To find a large effect (Cohen’s $d = 0.8$), G*Power3.1 indicated that using a power of 0.8 and an alpha level of 0.05 (5%), we would need a minimum of 42 participants (21/group). To find a medium effect (Cohen’s $d = 0.5$) G*Power3.1 indicated that using a power of 0.8 and an alpha level of 0.05 (5%), we would need to test 102 participants (51/group). Our maximum sample size was based on a medium effect size, as a small effect size might not be of as much interest commercially (at least initially). That is, if we found a small effect, we would argue that the synchrony setup would not be an interesting prospect for those designing and producing robots.

As outlined in the preregistration we initially tested 42 participants (the minimum sample size), then calculated the updated Bayes Factor (BF) after every 4 participants. As the BF was below 6 (considered “strong” evidence that the alternate hypothesis is true, opposed to null) we continued recruiting and testing until $n=89$. Testing was halted at 89 participants opposed to 102 (the maximum sample size we pre-registered), due to time constraints and difficulties recruiting individuals over the age of 60.

3.3 Participants

Eighty-nine individuals took part in the experiment, however the data from 12 individuals were excluded as they encountered problems which affected their experience with Pepper (error lights within Pepper, loss of Bluetooth/WIFI connection, and hearing problems). As a result, the final sample included 77 individuals aged 18-83 (mean age = 43.36, SD = 21.38), with 31 individuals over the age of 60 (“older adults”). Participants were individuals residing in Glasgow (Scotland, UK) and were initially recruited by word of mouth (in person, via email, and through social media advertisements) followed by snowball sampling. All individuals had normal or corrected-to-normal vision and hearing, and no previous experience interacting with the robot used in the study. Participants were randomly assigned to the Synchronous (n=40) or Asynchronous (n = 37) group prior to arrival and were compensated £10 for their participation.

Independent samples t-tests revealed no significant differences between the groups regarding age, nor on questionnaires assessing negative attitudes towards robots (Nomura et al., 2008), anthropomorphic tendencies (Waytz, Cacioppo & Epley, 2010), and general empathetic concern (Batchelder, Brosnan & Ashwin, 2017).

Fig 2. Plots to illustrate the two groups’ distributions regarding negative attitudes towards robots (top), anthropomorphic tendencies (middle), and general empathetic concern (bottom).

Both Synchronous and Asynchronous groups contained more women than men (57.50% and 62.16%, respectively). In the Synchronous group, one individual identified as “Agender”, and in the Asynchronous group, one person identified as “Non-Binary”.

3.4 Experiment Design

Participants assigned to the Synchronous group interacted with a humanoid robot (see ‘Apparatus’ for more details) whose shoulder lights were programmed to illuminate at the same rate as their heartbeat. This was achieved by sampling the participant’s heart rate using a wearable heart rate monitor and relaying this information to the robot in real time. In the asynchronous group, the heart rate data were sampled in the same way, however the lights of the robot were programmed to

flash at a rate 20% slower than the individual's heart rate (thus producing an asynchronous cardio-visual experience).

Regardless of group allocation, participants completed the same tasks with the robot, and the same measures. To account for between-group differences, baseline measures were taken from both groups before they interacted with the robot.

3.5 Apparatus

3.5.1 Robotic Platform

The robot used in the experiment was the Pepper Robotic system - a commercially available humanoid robot from SoftBank Robotics (Tokyo, Japan). Pepper is 120cm tall and features 2 in-built cameras, as well as microphones and tactile sensors, which allow it to detect objects and movement in the environment.

Pepper also has expressive movement and speech capabilities that can run autonomously, but for the purpose of experimental control, we used a Wizard of Oz set up wherein we controlled Pepper's behaviour. Specifically, for one part of the experiment we triggered a sequence of speech and movements via the 'Choreograph' software (See 'Procedure' for further details). In another section we created a panel of key phrases using html. Upon clicking a speech button, a corresponding line of Python code is triggered and Pepper speaks and moves accordingly. To maintain experimental control, for each interaction the experimenter systematically clicked from the first phrase "Hi there", through a series of closed questions and responses, to the final phrase "Thank you". Closed questions (E.g. do you prefer tea or coffee?) and option-specific responses were used to create the illusion that the robot was responding to the specific words of the participant.

In addition to controlling the speech and movement of the robot, it is also possible to co-opt the lights in the shoulder panels. Specifically, rather than flashing to indicate the 'mode' or 'state' of Pepper, we programmed the lights to illuminate either synchronously or asynchronously with participants' heart rate, depending on group assignment.

Fig 3. Illuminating shoulder lights on the Pepper robot

3.5.2 Heart Rate Monitor

To capture heart rate information, we used a Polar OH1 optical heart rate sensor. The Polar OH1 heart rate sensor was chosen due to its reliability, sport-focussed design (allowing for freedom of movement without diminished accuracy), Bluetooth Low Energy (BLE) capabilities (necessary for relaying the heart rate information to external devices), and ability to be worn on the wrist (as opposed to uncomfortable or obstructive devices such as finger clips or chest straps).

The data collected by the heart rate monitor were relayed via a laptop to the robot, allowing the shoulder lights of the robot to illuminate at the same rate as the participant's heart rate ('Synchronous'), or 20% slower ('Asynchronous'). We contemplated using random or stable rates as the control, however decided against this as we wished to ensure that the effect was the result of the synchrony manipulation - not other factors such as effects of dynamic vs stable pulsing.

3.6 Dependent Measures

3.6.1 Questionnaires

To probe 'liking', participants completed the validated 'Liking' scale of the Godspeed Questionnaire (Bartneck, Kulic, Croft & Zoghbi, 2009) before and after interacting with the robot. See '2.1 Likability Questionnaire' for background. To explore if

and how participants' perception of the robot is affected, we also administered questionnaires probing the extent to which participants perceived the robot as a social agent; specifically, the Inclusion of Self in Other task (Aron, Aron & Smollan, 1992) and the Robotic Social Attributes Scale (Carpinella, Wyman, Perez & Stroessner, 2017). To check for between-group differences in attitudes towards robots, participants also completed questionnaires to probe their history with robots via the Exposure to Cinematic Depictions of Robots (Riek, 2011), and Negative Attitudes towards Robots (Nomura et al., 2008) questionnaires. To account for between-group differences with regards to anthropomorphic tendencies and general empathy, the Individual Differences in Anthropomorphism Questionnaire (Waytz, Cacioppo & Epley, 2010), and the Empathy Components Questionnaire (Batchelder, Brosnan & Ashwin, 2017) were also administered. Participants were given the option of completing the questionnaires on paper, or via the online questionnaire platform form{r} (<https://formr.org/>).

3.6.2 Gaze Cueing

To investigate the extent to which participants perceive the Pepper robot as an intentional social agent, as opposed to an object, we administered a computer-based gaze cueing paradigm. See Introduction for further insight into this method. In the task participants saw images of an arrow, or the faces of Pepper, a different robot, or a human (dimensions: approximately 600mm x 800mm). The robot images are original - taken in the Social Brain in Action Laboratory. The face of the human was chosen from the Karolinska Emotional faces database (Lundqvist, Flykt & Ohman, 1998), on the basis of its resemblance to the robot in both form and contrast. All images were changed to black and white, to control for potential influences of colour on attention. To retain ecological validity, the decision was made not to control for contrast and composition between the agent types.

Classically, gaze cueing paradigms are conducted with arrows or eyes, however the limited eye movements of the robots forced us to adopt a whole-head shift in direction. Studies have demonstrated that gaze cueing effects are still present when stimuli depict the whole head

of the agent (Frischen, Bayliss & Tipper, 2007; Langton & Bruce, 2000), further validating this decision. The task was created using the PsychoPy3 experiment builder and was presented to participants on a 21.5-inch iMac desktop computer.

In the gaze cueing task, participants see a fixation cross, followed by a front-facing image. See 'Neutral' images on Fig. 4. They then see said agent orientating in the left or right direction. See 'Direction Cueing' on the same figure. The participant then sees a target (in this case, an asterisk) appear congruent, incongruent, or neutral relative to the face. The sequence of images was presented in quick succession, leading to the illusion of apparent motion of the head/face. Refer to Introduction for further details, see Fig. 4 for illustration of trial types and timings. The participants are instructed to focus on the fixation cross, then respond "as quickly and accurately as possible" upon seeing the 'target' appear. The participant responds to the left or right targets by pressing the corresponding arrow key on the keyboard (left, or right, respectively).

Fig. 4 Neutral and Direction Cueing stimuli (A) and stimuli presentation timings and examples (B)

Responses were removed if slower than 150ms ('preemptive') or above 2500ms ('unusually slow'). This is in line with previous work scrutinising attention cueing differences in young and older adults (Gayzur et al., 2014) and pieces focussed on the 'effective analysis of reaction time data' (Whelan, 2008). These

criteria led to the exclusion of approximately 0.26% of the dataset (35 data points).

3.6.3 Hesitance to Hit

To understand the extent to which participants empathise with the Pepper robot, and perceive it as a social agent (compared to an object), we administered an adapted version of the ‘hesitance to hit’ paradigm. The participant takes a seat at the table across from the robot, and the experimenter stands next to the robot. After asking the participant to don a pair of safety goggles, the experimenter then proceeds to read the script to the participant:

“Right, there is something I haven’t told you about this experiment. This Pepper is one of ten specially designed robots that I was given as part of a large research grant. By ‘specially designed’ I mean that they’re totally shatterproof – so if you hit one, the robot will break in a safe way that’s easy to repair. The reason Pepper is designed this way is because our lab is interested in what happens when someone has to hit a robot – for example, if a robot was to malfunction and you had to hit and disable one. Does that make sense? Great.”

Experimenter passes participant the hammer

“So, for this part, your task is to give the robot one hard hit on the head. So, when you’re ready, come round the table and I’ll get out of the way.”

After standing up to hit the robot (indicating the intention to hit the robot), the participant is told to pass the hammer to the experimenter, and is informed that they will not actually be hitting the robot. The participant is prompted to remove their safety goggles and is invited back to the control area for a task debriefing. If the participant verbally protests against hitting the robot, they are told “It’s just part of the experiment”. Upon protesting three times, the task is ended as indicated previously, and the participant is deemed to have ‘refused’ to hit the robot.

To better understand what the hesitance reflects, and validate the technique as a measure of social attribution, we interviewed participants regarding why they hesitated as part of the debriefing procedure.

3.7 Procedure

Upon entering the room (see Fig 5.) participants took a seat in the control space, in front of the iMac (depicted in blue). After participants provided written informed consent, they undertook the computer-based gaze-cueing paradigm, then saw a video of Pepper (Tech Insider, 2018) and were asked to complete questionnaires probing how much they liked the robot (See Dependent Measures for details). During this time the experimenter turned on the robot and sat in the testing space. After this, the participant donned the heart rate monitor and the robot was triggered to illuminate in a synchronous or asynchronous manner (group-dependent). The participant was then taken to the testing space and was seated in front of the robot.

Fig 5. Room setup illustrating a room divider separating the testing space (left) and control space (right).

Participants drew and made notes about the robot (5 minutes), then observed as the robot performed the ‘Tai Chi’, ‘Vacuuming’, and ‘Disco Dancing’ movements (2 minutes). These tasks were undertaken to subtly encourage attention towards the synchronous/asynchronous shoulder lights. Participants were then informed that Pepper would ask them about their food preferences, and they would create a shopping list together. The experimenter used the excuse “I’ll get out of the way so that Pepper doesn’t try to talk to me as well” and returned to the control space to control Pepper (as described in 3.5.1 ‘Robotic Platform’). The participant was asked a series of questions about their food preferences, and what they would like on a food shopping list. A small video recording device (Logitech webcam) mounted on the room separator was triggered, which allowed the experimenter to monitor and record this interaction.

After repeating the liking questionnaires and Gaze Cueing task,

participants were then invited to take a seat in front of the robot. The experimenter discreetly turned on the video camera again, and joined the participant in the testing space. After asking participants to don a pair of safety goggles, the hesitance to hit instruction was given (see 3.6.3 Hesitance to Hit). Upon completion of the task (either by agreeing or refusing to hit the robot) the participant is invited back into the testing space. The participant is asked a series of open questions (E.g. “After I asked you to hit the Pepper, what was going through your mind?”) to probe their thoughts and feelings during the hesitance to hit task. For the full list of questions see Supplementary Materials A.

To conclude the experiment, participants completed demographic and personality questionnaires (See Dependent Measures for details), then were debriefed and compensated for their time.

4 Results

This study was designed to evaluate three main hypotheses relating to the impact of cardio-visual synchrony between a human and robot. Specifically, we predicted that:

1. after interacting with the robot, the Synchronous group will rate the robot higher on the ‘Liking’ scale of the Godspeed questionnaire, compared to the Asynchronous group.
2. the Synchronous group, compared to the Asynchronous group, will hesitate for longer after being asked to hit the robot with a mallet.
3. more individuals in the Synchronous group, compared to the Asynchronous group, will refuse to hit the robot.

In the following, each hypothesis is addressed in turn.

4.1 Hypothesis 1: Liking

To determine whether the robot was rated as more likable, depending on whether it illuminated in a manner that was synchronous or asynchronous with a participant’s heart rate, scores on the ‘Likability’ scale of the Godspeed questionnaire were considered. As each of the 5 items in the Likability scale is rated from 1-5,

the minimum score that could be given is 5 and the maximum is 25. 25 indicates high scores on all dimensions – Dislike-Like, Unkind-Kind, Unfriendly-Friendly, Unpleasant-Pleasant, Awkward-Nice.

Before the interaction, descriptive statistics indicated little difference between the scores of the Synchronous ($M=20.61$, $SD=3.89$) and Asynchronous group ($M=20.39$, $SD=3.65$). After the interaction, both groups rated the robot slightly higher on the likability scale, however little difference between the Synchronous ($M=21.76$, $SD=4.00$) and Asynchronous ($M=21.14$, $SD=3.76$) groups was found (See Figure 6 for visualisation).

Fig 6. Ratings on the Likability Scale. Each point represents the score of an individual participant. The crosshair represents the mean of the participants in that condition, at the specified time point.

The results of a 2 x 2 ANOVA indicated a Bayes Factor of 0.215 when considering the difference between the synchrony groups. The analyses also indicated a Bayes Factor of 0.532 for the effect of Time (Before vs After interacting with the robot). These Bayes factor values are substantially lower than 6 – the value proposed to indicate “strong” evidence that the alternate hypothesis is true. This indicates that neither condition, nor time, changed the extent to which participants liked the robot in our particular manipulation.

4.2 Hypothesis 2: Hesitance to Hit

Compared to the Asynchronous group, we hypothesised the Synchronous group would hesitate for longer between being asked to hit the robot with a mallet, and agreeing to do so.

To account for individuals who refused to hit the robot, we adopted the method used by Darling, Nandy and Breazeal (2015) - adding 1 second onto the maximum measured hesitation. This method was originally justified as follows: “if a subject did not strike the robot, we considered this to be greater than the maximum measured hesitation” (Darling, Nandy & Breazeal, 2015). This led to three values of 101 in the dataset. For clarity, these values were not plotted in Figure 2, however they were included in the descriptive statistics and statistical tests. The ‘refusals’ will be discussed more in the later Results section ‘Hypothesis 3: Refusal to Hit’.

Descriptive statistics indicated that the Asynchronous group ($M = 18.50s$, $SD = 28.92$) hesitated for longer, on average, than the Synchronous ($M = 14.75$ seconds, $SD = 22.14$) group. See Figure 7 for visualisation.

Fig 7. Length of hesitation between being asked to hit the robot, and agreeing to do so. Each circle represents an individual participant and the crosshairs indicate the mean for each group. Note: three participants refused to hit the robot and have been excluded from this plot as a result. The next section of results will discuss these individuals in greater detail.

To compare statistical significance of these results we had intended to perform an independent samples T-Test, however the assumptions of normality and equal variances were violated (Shapiro-Wilkov and Levene’s $p < .05$). This led to the use of a two-tailed Mann-Whitney U test instead. This test indicated no significant differences between hesitation time of the two, $U = 698.0$, $p = 0.864$.

4.3 Hypothesis 3: Refusal to Hit

The final hypothesis states that more individuals in the Synchronous group, compared to the Asynchronous group, will refuse to hit the robot. A reminder: a ‘refusal’ results from a participant protesting (e.g. “I don’t want to...”, “Do I have to...?”) three times. To compare the number of refusals made by participants, between the groups, refer to Figure 8.

Fig 8. Chart to illustrate the frequency of individuals exhibiting the distinct number of protests (0, 1, 2, or 3 protests). Note: After protesting 3 times the experiment was terminated, and the participant was deemed to have ‘refused’ to hit the robot.

The results of a Mann Whitney U test indicated no significant difference between the number of protests made by the Synchronous group ($M = 0.52$, $SD = 0.68$), compared to the Asynchronous group ($M = 0.60$, $SD = 1.14$), $t(39) = 262.50$, $p = 0.917$. A Mann Whitney U test, opposed to an Independent T Test was used because a Levene’s test indicated inequality of variances between the groups ($p < .05$).

4.4 Exploratory Analyses

As pre-registered, we explored whether gaze cueing effects differed between the two groups. To do so we performed a $2 \times 2 \times 4 \times 3$ Mixed ANOVA: Group (Synchronous vs Asynchronous), Time (Pre-Interaction vs Post-Interaction), Agent (Human vs Pepper Robot vs Other Robot vs Arrows), Congruency

(Congruent vs Incongruent vs Neutral). The assumption of Sphericity was violated, therefore the Greenhouse-Geisser correction was used.

Fig 10. Plots to illustrate the gaze cueing reaction time data, split by agent type. Violin plots indicate the variance of the data within each condition.

In contrast to our hypothesis, the results also indicated that these effects were not significantly different between the two groups [agent x group: $F(2.470, 167.989) = 2.382$, $p = .083$, $\eta^2_p = .034$; congruency x group: $F(1.836, 124.852) = 3.092$, $p = .053$, $\eta^2_p = .043$].

The results of the ANOVA did indicate a main effect of agent ($F(2.470, 167.989) = 8.355$, $p < .001$, $\eta^2_p = .109$). Descriptive statistics further reveal that on average, participants were slower to respond to the arrow condition ($M = 0.4043$, $SD = 0.1179$) compared to the Human ($M = 0.3937$, $SD = 0.1837$), Other Robot ($M = 0.3963$, $SD = 0.1575$), and Pepper ($M = 0.3888$, $SD = 0.1087$) trials.

The results also indicate a main effect of congruency ($F(1.836, 124.852) = 52.969$, $p < .001$, $\eta^2_p = .438$). Descriptive statistics indicated that participants were, on average,

fastest to respond to congruent trials ($M = 0.3795$, $SD = 0.1678$) compared to incongruent ($M = 0.3957$, $SD = 0.1271$) and neutral ($M = 0.4121$, $SD = 0.1357$) trials.

No significant interaction emerged between agent x congruency after applying the Greenhouse Geisser correction, $F(3.305, 224.734) = 2.161$, $p = .087$, $\eta^2_p = 0.020$. Additionally, time had no significant effect on the effects of agent ($F(2.324, 158.006) = .926$, $p = .429$, $\eta^2_p = .013$) or congruency ($F(1.723, 117.144) = 1.170$, $p = .309$, $\eta^2_p = .017$).

5 Discussion

Our aim in the present study was to determine whether synchrony between a participant's heart rate and the illumination of lights on a robot's torso (cardio-visual synchrony) would lead participants to exhibit increased prosocial attitudes and behaviours towards the robot. Our results demonstrate that overall, heart rate synchrony did not impact quantitative or qualitative measures of positive attitudes toward the Pepper robot, including self-report questionnaire responses, or hesitancy to hit the robot when instructed. Furthermore, we did not find evidence that heart rate synchrony influenced a cognitive measure of social attribution – specifically, gaze cueing effects.

5.1 Liking

The Godspeed questionnaire results indicate that participants in the Synchronous group did not rate Pepper as more “likable”, “friendly”, “kind”, “pleasant”, or “nice”, compared to the Asynchronous group. As a result, no differences between the groups on the “Likability” scale of the Godspeed questionnaire emerged. As can be seen in Figure 4, it would appear that both groups report high liking ratings for Pepper both before and after the manipulation. As such, liking ratings that cluster near the ceiling of the scale do not have much room for improvement. An important limitation to note regarding the Likability scale of the Godspeed questionnaire is that it cannot currently account for nuances between social liking and liking in other forms. For example, you might not perceive your kettle as friendly and kind, however a new add-

on might make it more likable as a result of improved and usability. Consequently, it is not possible to rule out whether other factors of likeability not probed by this scale changed based on our manipulation. A challenge for future work, therefore, will be to explore ways of probing liking with more nuance and range. The focus on “likability” results from our interest in perceptions of Pepper as a social agent, however it would also be insightful to probe factors affecting the usage and uptake of technologies as Pepper is inherently a machine. Davis, Bagozzi, and Warshaw (1989) identified that “perceived usefulness” and “perceived ease” were positively correlated with current usage and future use of computers, so measures that take account of these factors would appear to be useful.

5.2 Hesitance to Hit

Moving to the objective measures we collected regarding prosocial behaviours toward the Pepper robot, our reaction time data indicated that participants’ hesitance to hit the robot did not vary significantly between the synchronous and asynchronous heart rate manipulation groups. It is of note, however, that the number of ‘absolute refusals’ (protesting against hitting the robot three times, resulting in the termination of the task) did differ between the groups. Contrary to predictions, in the synchronous group, no individual absolutely refused to hit the robot, whereas in the asynchronous group, three individuals absolutely refused to do so. Although it is impossible to conclude anything meaningful about our manipulation leading to this result, given such small numbers, this finding could nonetheless be followed up by future work.

If this finding were to be replicated, one possibility explaining why individuals in the asynchronous group refused to hit Pepper comes from literature regarding the effects of engineered music on relaxation (Leslie et al., 2019). It has been demonstrated that slow music relative to an individual’s breathing rate can induce slowed breathing and a calm state (Leslie et al., 2019). In a similar way, the observation of the slow pulsing lights relative to the person’s heart rate (as in the asynchronous condition) could have given rise to a similar relaxed state. This is supported by a statement from one participant in the asynchronous condition who commented that

the pulsing lights were “calming”. This calm state might have then led to an aversion to hit the robot, along the lines of empirical work that demonstrates induced relaxation can reduce aggression and increase prosocial behaviours (Whitaker & Bushman, 2011).

Many questions remain though, including why only three out of 37 people in the asynchronous group refused to hit Pepper, whereas the others agreed. Individual differences in attention to the lights, and susceptibility to the relaxing properties of the slow pulsing, offer suggestions as to why this might be the case. To gain greater insight regarding this theory though, it is necessary to undertake further experimentation targeting the use of illumination for relaxation, and using both subjective and physiological measurements of arousal.

Another relevant aspect of our experiment design that could have influenced participants’ hesitancy to hit was the lack of experimenter blinding. While we do not find evidence that experimenter bias resulted in more individuals in the synchrony group refusing to hit the robot, it could be the case that experimenter bias for some reason influenced participants’ behaviour in the opposite direction, leading to more hit refusals in the asynchronous group. We would encourage future studies wishing to use the hesitancy to hit paradigm to use double blinding procedures so that any impact of experimenter bias, no matter the direction, can be ruled out.

It is important to note that although refusals and hesitation occurred, all but three participants (n=74) agreed when asked to hit Pepper on the head with the hammer. These results are comparative to those in a study by Rosalia and colleagues (2005), in which all participants complied when asked to administer the high level of electric shocks to a robot. This study was modelled on the (in)famous work by Milgram (1965), wherein participants were asked to administer electric shocks to a human confederate. In the original work with humans performed by Milgram, he reported that the majority of participants (24 out of 40) did not comply - refusing to administer the highest level of shock to the other person. It could be argued that compliance was lower in the Milgram study due to the verbal protests of the confederate, however Rosalia and colleagues

(2005) also included vocal protests (“Please, please stop” etc.) in their robotic replication. While the Rosalia et al. (2005) study was very different from the current study in a number of important ways (differences in the setup, robot used, the nature of the harm, etc...) it nonetheless corroborates the suggestion that when instructed, people will more readily comply when asked to harm a robot, compared to a human. Such findings clearly underscore that we perceive living and artificial agents differently, and hold important ethical implications for human-robot interactions (Fosch-Villaronga & Albo-Canals, 2019; Ozcana et al., 2016).

5.3 Qualitative Insights

Despite finding no between-group difference in the hesitance to hit quantitative measures, the qualitative data gained from debriefing procedures was particularly insightful. Upon asking participants what they were thinking and feeling after being asked to hit the robot, participants indicated numerous reasons for hesitating. These included (but were not limited to) fondness for Pepper, worries about the cost of damaging the robot, concern for their own safety, disbelief about the instruction, feeling uncomfortable about being asked to commit an act of ‘violence’, social desirability effects, and feeling pressure from an authority figure. Upon being asked questions such as “what were you thinking after I asked you to hit the robot?” and “were you experiencing any emotions?” some people reported excitement about the prospect of hitting the robot, and described Pepper simply as “the robot”, “it”, an “object”, or a “machine”. In contrast, other individuals expressed that they felt very uncomfortable about hitting the robot, were upset about the thought of doing so, and referred to Pepper as “he”, “she”, a “social presence” and “like a child”. These contrasting perspectives were observed in both the synchronous and asynchronous groups.

The extent of the differences in anthropomorphism and feelings of upset are intriguing, especially when considering that all participants had a similar controlled experience with Pepper. This diverse range of feelings and responses to the request to hit a robot leads us to suggest that it is not advisable to use reaction time alone as a measure of liking for this paradigm, and, returning to considerations

presented in Section 5.1, that researchers should explore other methods of behaviourally probing liking and attachment.

In the literature it is frequently stated that humans have a natural tendency to anthropomorphise machines (Nie et al., 2012; Reeves & Nass, 1996), yet our results highlight that some people (those describing Pepper as “it”, “an object” etc) appear to exhibit no such tendency. Due to the self-reported nature of the apparent anthropomorphism, it could be the case that such individuals did perceive Pepper as more than a machine, but did not want to admit so (due to embarrassment or some other reason). Alternatively, it may be that our assumption that all humans anthropomorphise machines is inaccurate, and that some people lack the capacity for the attribution of human-like qualities to a robotic system. To probe apparent individual differences in anthropomorphism, and question our current understanding of the nature of anthropomorphism, it would be beneficial to quantify the qualitative data collected in this study (e.g. through content analysis) and compare these values to scores on the various questionnaires.

How we perceive and interact with a robot is proposed to be influenced by our tendency to anthropomorphism in general (Waytz, Cacioppo & Epley, 2010), our ability to empathise with others (Darling, Nandy & Breazeal, 2015), the negative attitudes we hold about towards robots in general (Nomura et al., 2008; Syrdal, 2009), and our exposure to certain films depicting robots (Riek, Adams & Robinson, 2011). Given prior work in this domain that has led to the development of a number of relevant questionnaires, future work might wish to explore in a more systematic way whether participant responses to one or several of these questionnaires hold predictive value for participants’ susceptibility to prosociality measures (c.f. Henschel & Cross, 2020). Anthropomorphism is proposed to be of huge importance for bonding between humans and animals (Mithen, 1996), so by conducting follow-up work probing anthropomorphic perceptions and tendencies more sensitively, it might be possible to identify populations which are more or less likely to bond with, and benefit from robotic agents, over the long term. Longitudinal studies in this domain would also

address whether quick and intense anthropomorphism (as seen in this study) is sustained long-term, and whether this translates to longer term uptake and bonding with social robots. Such longitudinal studies would also allow us to determine whether people who initially did not anthropomorphise the robots develop this tendency over time.

5.4 Gaze Cueing

A final task we performed to assess the extent to which participants viewed Pepper as a social agent, and whether the heart rate synchrony manipulation had any impact on such perceptions, was a gaze cueing task. The results from this task indicated no differences in gaze cueing effects between the two groups, or between pre- and post-manipulation testing points. While this finding is ultimately not surprising, given the null effects reported for other measures in this study, it is nonetheless valuable to consider what the gaze cueing findings suggest about social perception.

Previous literature indicates slowed responses for agents perceived as “mindful”, compared to those considered to be lacking knowledge or insight (See Introduction for details). Given these previous findings, we expected clear increases in response times when participants responded to the human face compared to the arrow stimuli. However, our findings indicated an almost opposite pattern. Participants were slower to respond to the arrow, compared to both the human and robot faces. While this finding is certainly at odds with what has frequently been reported by prior literature (Teufel et al., 2010; Teufel et al., 2009; Morgan, Freeth & Smith, 2018; Wiese et al., 2012), we suggest two reasons why people were on average slower to respond to the arrow than the human.

One explanation why gaze cueing effects were stronger for the arrow, relative to the faces, is because the arrow is an extremely strong and salient direction cue, and the task in our study was to respond solely to the direction of the target. In previous studies participants have been asked to respond to the nature of the target (e.g. what letter is presented) in addition to the direction (Driver et al., 1999; Greenwood, Parasuraman & Haxby, 1993), offering one explanation why previous results differ to our own. It is possible that with the addition of a target-discrimination element, participants

would be less distracted by the direction of the arrow, and we would see a pattern of results more reminiscent of previous findings (reduced gaze cueing effects for the arrow relative to eye gaze). Another reason for the unexpected finding could be because, contrary to many previous gaze cueing tasks, we used a full head shift, opposed to a movement of just the eyes.

Langton and Bruce (1999) demonstrated that head orientation and gaze direction both yield robust gaze cueing effects, leading us to initially deem the current stimuli as appropriate. However, upon reflecting on the result, we found evidence in the literature that when the face is averted by 30 degrees (similar to our stimuli), gaze effects were reduced (Hietanen, 1999). Hietanen (1999) argued that this is because when both head and eyes are averted, the agent is perceived to be oriented away from the participant and gazing straight ahead. As a result, the behaviour of the perceived agent is perceived as being less socially engaged or relevant to the task at hand, and the direction of attention has less signal value. The findings of the study by Hietanen offer insight as to why gaze cueing effects in response to the robot and human faces might have been reduced, rather than amplified, relative to the arrow.

With that being said, an emerging number of studies exploring gaze cueing effects with avatars and physically embodied robots demonstrate that full shifts in visual attention (as demonstrated by a head and gaze shift) by embodied and virtual agents can indeed have powerful impacts on human behaviour when performed *in situ* (e.g. Schilbach et al., 2010; Wilms et al., 2010; Kompatsiari et al., 2018). Perhaps the most compelling research to combine the gaze cueing paradigm with embodied robots comes from work with the iCub robot, which indeed demonstrates that an physically co-located robot’s gaze cues can be followed and used to induce prosocial feelings in a human interaction partner towards a robot (e.g., Ghiglino et al., 2020; Kompatsiari et al., 2019). Taken together, our failure to replicate the general human > arrows gaze cueing effect decreases what we might be able to read into any lack of effect we report with the heart rate synchrony manipulation on following Pepper’s gaze. However, we would suggest that probing more deeply into reasons why this is the case might not be as instructive or useful as further work looking to understand the role of eye gaze

and prosociality with embodied artificial agents, as these, and not screen-based images, are the focus of research and development into socially assistive robotics.

6. Conclusions

Here we report that cardio-visual synchrony had no significant effect on prosocial attitudes and behaviours towards the Pepper robot. The subtlety of the manipulation offers one explanation as to why no difference emerged, and it would be beneficial for our understanding of perceived synchrony to conduct further research using a more obvious signal or a robot that illuminated more uniformly (such as the MiRo robot developed by Consequential Robotics). Although the current study failed to yield affirmative answers regarding the utility of cardio-visual synchrony for inducing prosocial thoughts and behaviours toward a robot, it nonetheless highlights potential issues with certain methods used to measure perceptions of, and behaviour towards, robotic agents. Through qualitative data we show that the hesitance to hit paradigm is not only probing liking, but instead it measures other factors such as perceived cost and a person's concern for their own safety. We also find huge variation between participants with regards to how they refer to Pepper, and perceived feelings of connection towards the system. Based on current and other recent findings (e.g., Henschel & Cross, 2020), we would also argue that measures such as the Godspeed Liking scale are overly specific - focussing on social liking and neglecting other forms of liking in the process. In addition to advising the use of embodied human-robot interaction for probing other cognitive measures of social perception, such as gaze cueing (e.g., Kompatsiari et al., 2018; 2019; see also Henschel, Hortensius & Cross, 2020), we also suggest that future studies consider adding more might reveal about their research questions and robotic system of interest. In doing so it is possible to gain new insights regarding measurement techniques and individual differences, and generate research questions which are highly relevant to our future with robotic technologies.

Acknowledgements

The authors wish to thank members of the Social Brain in Action laboratory and #TeamSoBots for valuable assistance with and feedback on the research, especially Bishakha Chaudhury for her programming wizardry, and Kohinoor Darda, Te-Yi Hsieh and Rebecca Smith for feedback on a previous version of the manuscript. This work was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement number 677270), the Leverhulme Trust (PLP-2018-152) and an industrial strategy PhD studentship from SGSSS/RCUK.

References

- Admoni, H., Bank, C., Tan, J., Toneva, M., & Scassellati, B. (2011). Robot gaze does not reflexively cue human attention. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 33, No. 33).
- Aron, A., Aron, E. N., & Smollan, D. (1992). Inclusion of other in the self scale and the structure of interpersonal closeness. *Journal of personality and social psychology*, 63(4), 596.
- Bartneck, C. and Hu, J., 2008. Exploring the abuse of robots. *Interaction Studies*, 9(3), pp.415-433.
- Bartneck, C., Kanda, T., Ishiguro, H., & Hagita, N. (2007, August). Is the uncanny valley an uncanny cliff?. In *RO-MAN 2007-The 16th IEEE international symposium on robot and human interactive communication* (pp. 368-373). IEEE.
- Bartneck, C., Kanda, T., Ishiguro, H., & Hagita, N. (2009, September). My robotic doppelgänger-A critical look at the uncanny valley. In *RO-MAN 2009-The 18th IEEE International Symposium on Robot and Human Interactive Communication* (pp. 269-276). IEEE.
- Bartneck, C., Kulić, D., Croft, E., & Zoghbi, S. (2009). Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International journal of social robotics*, 1(1), 71-81.
- Bartneck, C., Rosalia, C., Menges, R. and Deckers, I., 2005. Robot abuse-a limitation of the media equation.
- Bartneck, C., Van Der Hoek, M., Mubin, O., & Al Mahmud, A. (2007, March). "Daisy, daisy, give me your answer do!" switching off a robot. In *2007 2nd ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (pp. 217-222). IEEE.
- Batchelder, L., Brosnan, M., & Ashwin, C. (2017). The development and validation of the Empathy Components Questionnaire (ECQ). *PloS one*.
- Bayliss, A. P., Di Pellegrino, G., & Tipper, S. P. (2004). Orienting of attention via observed eye gaze is head-centred. *Cognition*, 94(1), B1-B10.
- Benefeld, L. E., & Holtzclaw, B. J. (2014). Aging in place: merging desire with reality. *Nursing Clinics*, 49(2), 123-131.

- Bharucha, A. J., Anand, V., Forlizzi, J., Dew, M. A., Reynolds III, C. F., Stevens, S., & Wactlar, H. (2009). Intelligent assistive technology applications to dementia care: current capabilities, limitations, and future challenges. *The American journal of geriatric psychiatry*, 17(2), 88-104.
- Broekens, J., Heerink, M., & Rosendal, H. (2009). Assistive social robots in elderly care: a review. *Gerontechnology*, 8(2), 94-103.
- Cacioppo, J. T., Hughes, M. E., Waite, L. J., Hawkley, L. C., & Thisted, R. A. (2006). Loneliness as a specific risk factor for depressive symptoms: cross-sectional and longitudinal analyses. *Psychology and aging*, 21(1), 140.
- Carpinella, C. M., Wyman, A. B., Perez, M. A., & Stroessner, S. J. (2017, March). The Robotic Social Attributes Scale (RoSAS) Development and Validation. In *Proceedings of the 2017 ACM/IEEE International Conference on human-robot interaction* (pp. 254-262).
- Davis, F. D., Bagozzi, R. P., & Warshaw, P. R. (1989). User acceptance of computer technology: a comparison of two theoretical models. *Management science*, 35(8), 982-1003.
- Darling, K., Nandy, P., & Breazeal, C. (2015, August). Empathic concern and the effect of stories in human-robot interaction. In *2015 24th IEEE international symposium on robot and human interactive communication (RO-MAN)* (pp. 770-775). IEEE.
- Durkheim, E. (1995). The elementary forms of religious life, translated by Karen E. Fields. New York: Free.
- Driver IV, J., Davis, G., Ricciardelli, P., Kidd, P., Maxwell, E., & Baron-Cohen, S. (1999). Gaze perception triggers reflexive visuospatial orienting. *Visual cognition*, 6(5), 509-540.
- Fosch-Villaronga, E., & Albo-Canals, J. (2019). "I'll take care of you," said the robot. *Paladyn, Journal of Behavioral Robotics*, 10(1), 77-93.
- Foster, M. E., Alami, R., Gestranis, O., Lemon, O., Niemelä, M., Odobez, J. M., & Pandey, A. K. (2016, November). The MuMMER project: Engaging human-robot interaction in real-world public spaces. In *International Conference on Social Robotics* (pp. 753-763). Springer, Cham.
- Fozard, L., Jan Rietsema, Herman Bouma, JAM Graafmans, J. (2000). Gerontechnology: Creating enabling environments for the challenges and opportunities of aging. *Educational Gerontology*, 26(4), 331-344.
- Friborg O, Martinussen M, Rosenvinge JH (2006) Likert-based vs. semantic differential-based scorings of positive psychological constructs: A psychometric comparison of two versions of a scale measuring resilience. *Pers Individ Differ* 40:873–884
- Friesen, C. K., & Kingstone, A. (1998). The eyes have it! Reflexive orienting is triggered by nonpredictive gaze. *Psychonomic bulletin & review*, 5(3), 490-495.
- Gayzur, N. D., Langley, L. K., Kelland, C., Wyman, S. V., Saville, A. L., Ciernia, A. T., & Padmanabhan, G. (2014). Reflexive orienting in response to short- and long-duration gaze cues in young, young-old, and old-old adults. *Attention, Perception, & Psychophysics*, 76(2), 407-419.
- Ghiglinio, D., Willemse, C., De Tommaso, D., Bossi, F., & Wykowska, A. (2020). At first sight: robots' subtle eye movement parameters affect human attentional engagement, spontaneous attunement and perceived human-likeness. *Paladyn, Journal of Behavioral Robotics*, 11(1), 31-39.
- Greenwood, P. M., Parasuraman, R., & Haxby, J. V. (1993). Changes in visuospatial attention over the adult lifespan. *Neuropsychologia*, 31(5), 471-485.
- Gurley, K., & Norcio, A. F. (2009, July). A systematic review of technologies designed to improve and assist cognitive decline for both the current and future aging populations. In *International Conference on Internationalization, Design and Global Development* (pp. 156-163). Springer, Berlin, Heidelberg.
- Hawkley, L. C., Thisted, R. A., & Cacioppo, J. T. (2009). Loneliness predicts reduced physical activity: cross-sectional & longitudinal analyses. *Health Psychology*, 28(3), 354.
- Henschel, A., & Cross, E. S. (2020). No evidence for enhanced likeability and social motivation towards robots after synchrony experience. *Interaction Studies*, 21(1), 7-23.
- Hortensius, R., & Cross, E. S. (2018). From automata to animate beings: the scope and limits of attributing socialness to artificial agents. *Annals of the New York Academy of Sciences*, 1426(1), 93-110.
- Hortensius, R., Hekele, F., & Cross, E. S. (2018). The perception of emotion in artificial agents. *IEEE Transactions on Cognitive and Developmental Systems*, 10(4), 852-864.
- Hove, M. J., & Risen, J. L. (2009). It's all in the timing: Interpersonal synchrony increases affiliation. *Social cognition*, 27(6), 949-960.
- Kompatsiari, K., Ciardo, F., Tikhonoff, V., Metta, G., & Wykowska, A. (2018). On the role of eye contact in gaze cueing. *Scientific reports*, 8(1), 1-10.
- Kompatsiari, K., Ciardo, F., Tikhonoff, V., Metta, G., & Wykowska, A. (2019). It's in the eyes: the engaging role of eye contact in HRI. *International Journal of Social Robotics*, 1-11.
- Langton, S. R., & Bruce, V. (2000). You must see the point: automatic processing of cues to the direction of social attention. *Journal of Experimental Psychology: Human Perception and Performance*, 26(2), 747.
- Lehmann, H., Saez-Pons, J., Syrdal, D. S., & Dautenhahn, K. (2015). In good company? perception of movement synchrony of a non-anthropomorphic robot. *PloS one*, 10(5).
- Leslie, G., Ghandeharioun, A., Zhou, D., & Picard, R. W. (2019, September). Engineering Music to Slow Breathing and Invite Relaxed Physiology. In *2019 8th International*

Conference on Affective Computing and Intelligent Interaction (ACII) (pp. 1-7). IEEE.

Liu, J., Luo, H., Zheng, P. P., Wu, S. J., & Lee, K. (2018). Transdermal optical imaging revealed different spatiotemporal patterns of facial cardiovascular activities. *Scientific reports*, 8(1), 1-10.

Louie, W. Y. G., McColl, D., & Nejat, G. (2014). Acceptance and attitudes toward a human-like socially assistive robot by older adults. *Assistive Technology*, 26(3), 140-150.

Lundqvist, D., Flykt, A., & Öhman, A. (1998). The Karolinska Directed Emotional Faces - KDEF, CD ROM from Department of Clinical Neuroscience, Psychology section, Karolinska Institutet, ISBN 91-630-7164-9.

Macrae, C. N., Duffy, O. K., Miles, L. K., & Lawrence, J. (2008). A case of hand waving: Action synchrony and person perception. *Cognition*, 109(1), 152-156.

Martini, M. C., Buzzell, G. A., & Wiese, E. (2015, October). Agent appearance modulates mind attribution and social attention in human-robot interaction. In *International Conference on Social Robotics* (pp. 431-439). Springer, Cham.

Milgram, S. (1965). Some conditions of obedience and disobedience to authority. *Human relations*, 18(1), 57-76.

Mithen, S. (1996). The Prehistory of the Mind: A search for the origins of Art. *Religion and Science*.

Mogan, R., Fischer, R., & Bulbulia, J. A. (2017). To be in synchrony or not? A meta-analysis of synchrony's effects on behavior, perception, cognition and affect. *Journal of Experimental Social Psychology*, 72, 13-20.

Morgan, E. J., Freeth, M., & Smith, D. T. (2018). Mental state attributions mediate the gaze cueing effect. *Vision*, 2(1), 11.

Mörtl, A., Lorenz, T., & Hirche, S. (2014). Rhythm patterns interaction-synchronization behavior for human-robot joint action. *PLoS one*, 9(4).

Nie, J., Park, M., Marin, A. L., & Sundar, S. S. (2012, March). Can you hold my hand? Physical warmth in human-robot interaction. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (pp. 201-202). IEEE.

Niemelä, M., Heikkilä, P., & Lammi, H. (2017, March). A social service robot in a shopping mall: expectations of the management, retailers and consumers. In *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction* (pp. 227-228).

Nomura, T., Kanda, T., Suzuki, T., & Kato, K. (2008). Prediction of human behavior in human-robot interaction using psychological scales for anxiety and negative attitudes toward robots. *IEEE transactions on robotics*, 24(2), 442-451.

O'luanaigh, C., O'connell, H., Chin, A. V., Hamilton, F., Coen, R., Walsh, C., ... & Lawlor, B. A. (2012). Loneliness and cognition in older people: the Dublin Healthy Ageing study. *Ageing & mental health*, 16(3), 347-352.

Ozcana, B., Caligiorea, D., Fosch-Villarongab, E., Speratia, V., Morettac, T., & Baldassarrea, G. (2016). Designing socially assistive robot (sar) for cognitive child-robot interaction (crici) with children with autism

spectrum disorder-the case of “+ me. In *2nd International Conference on Social Robots in Therapy & Education*.

Perissinotto, C. M., Cenzer, I. S., & Covinsky, K. E. (2012). Loneliness in older persons: a predictor of functional decline and death. *Archives of internal medicine*, 172(14), 1078-1084.

Piau, A., Campo, E., Rumeau, P., Vellas, B., & Nourhashemi, F. (2014). Aging society and gerontechnology: A solution for an independent living?. *The journal of nutrition, health & aging*, 18(1), 97-112.

Pinquart, M., & Sörensen, S. (2003). *Risk factors for loneliness in adulthood and old age--a meta-analysis*. Nova Science Publishers.

Posner, M. I. (1980). Orienting of attention. *Quarterly journal of experimental psychology*, 32(1), 3-25.

Rantz, M. J., Marek, K. D., Aud, M., Tyrer, H. W., Skubic, M., Demiris, G., & Hussam, A. (2005). A technology and nursing collaboration to help older adults age in place. *Nursing Outlook*, 53(1), 40-45.

Rennung, M., & Göritz, A. S. (2016). Prosocial consequences of interpersonal synchrony. *Zeitschrift für Psychologie*.

Sel, A., Azevedo, R. T., & Tsakiris, M. (2017). Heartfelt self: cardio-visual integration affects self-face recognition and interoceptive cortical processing. *Cerebral Cortex*, 27(11), 5144-5155.

Riek, L. D., Adams, A., & Robinson, P. (2011, March). Exposure to cinematic depictions of robots and attitudes towards them. In *Proceedings of International Conference on Human-Robot Interaction, Workshop on Expectations and Intuitive Human-Robot Interaction*.

Rosalia, C., Menges, R., Deckers, I., & Bartneck, C. (2005). Cruelty towards robots.

Schilbach, L., Wilms, M., Eickhoff, S. B., Romanzetti, S., Tepest, R., Bente, G., ... & Vogeley, K. (2010). Minds made for sharing: initiating joint attention recruits reward-related neurocircuitry. *Journal of cognitive neuroscience*, 22(12), 2702-2715.

Suzuki, K., Garfinkel, S. N., Critchley, H. D., & Seth, A. K. (2013). Multisensory integration across exteroceptive and interoceptive domains modulates self-experience in the rubber-hand illusion. *Neuropsychologia*, 51(13), 2909-2917.

Syrdal, D. S., Dautenhahn, K., Koay, K. L., & Walters, M. L. (2009). The negative attitudes towards robots scale and reactions to robot behaviour in a live human-robot interaction study. *Adaptive and emergent behaviour and complex systems*.

Tanaka, F., Cicourel, A., & Movellan, J. R. (2007). Socialization between toddlers and robots at an early childhood education center. *Proceedings of the National Academy of Sciences*, 104(46), 17954-17958.

Tanioka, T. (2019). Nursing and Rehabilitative Care of the Elderly Using Humanoid Robots. *The Journal of Medical Investigation*, 66(1.2), 19-23.

Teufel, C., Alexis, D. M., Todd, H., Lawrance-Owen, A. J., Clayton, N. S., & Davis, G. (2009). Social cognition modulates the sensory coding of observed gaze direction. *Current Biology*, 19(15), 1274-1277.

- Teufel, C., Alexis, D. M., Clayton, N. S., & Davis, G. (2010). Mental-state attribution drives rapid, reflexive gaze following. *Attention, Perception, & Psychophysics*, 72(3), 695-705.
- Waytz A, Cacioppo J, Epley N. 2010 Who sees human? The stability and importance of individual differences in anthropomorphism. *Perspect. Psychol. Sci.* 5, 219 – 232. (doi:10.1177/1745691610369336)
- Wei, J., Luo, H., Wu, S. J., Zheng, P. P., Fu, G., & Lee, K. (2018). Transdermal optical imaging reveal basal stress via heart rate variability analysis: a novel methodology comparable to electrocardiography. *Frontiers in psychology*, 9, 98.
- Whelan, R. (2008). Effective analysis of reaction time data. *The Psychological Record*, 58(3), 475-482.
- Whitaker, J. L., & Bushman, B. J. (2012). “Remain calm. Be kind.” Effects of relaxing video games on aggressive and prosocial behavior. *Social Psychological and Personality Science*, 3(1), 88-92.
- Wiese, E., Wykowska, A., Zwickel, J., & Müller, H. J. (2012). I see what you mean: how attentional selection is shaped by ascribing intentions to others. *PloS one*, 7(9).
- Wilms, M., Schilbach, L., Pfeiffer, U., Bente, G., Fink, G. R., & Vogeley, K. (2010). It's in your eyes—using gaze-contingent stimuli to create truly interactive paradigms for social cognitive and affective neuroscience. *Social cognitive and affective neuroscience*, 5(1), 98-107.
- Wykowska, A., Chaminade, T., & Cheng, G. (2016). Embodied artificial agents for understanding human social cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1693), 20150375.
- Yang, G. Z., Nelson, B. J., Murphy, R. R., Choset, H., Christensen, H., Collins, S. H., ... & Kragic, D. (2020). Combating COVID-19—The role of robotics in managing public health and infectious diseases.
- Young, J. E., Hawkins, R., Sharlin, E., & Igarashi, T. (2009). Toward acceptable domestic robots: Applying insights from social psychology. *International Journal of Social Robotics*, 1(1), 95.