Effective Persuasion Strategies for Socially Assistive Robots

Katie Winkle
Bristol Robotics Laboratory
Bristol, UK
k.winkle@bristol.ac.uk

Séverin Lemaignan
Bristol Robotics Laboratory
Bristol, UK
severin.lemaignan@brl.ac.uk

Praminda Caleb-Solly
Bristol Robotics Laboratory
Bristol, UK
Praminda.Caleb-Solly@uwe.ac.uk

Ute Leonards
School of Psychological Science
University of Bristol
Bristol, UK
ute.leonards@bristol.ac.uk

Ailie Turton
Department of Allied Health Professions
University of the West of England
Bristol, UK
Ailie.Turton@uwe.ac.uk

Paul Bremner
Bristol Robotics Laboratory
Bristol, UK
paul.bremner@brl.ac.uk

***Abstract*—In this paper we present the results of an experimental study investigating the application of human persuasive strategies to a social robot. We demonstrate that robot displays of goodwill and similarity to the participant significantly increased robot persuasiveness, as measured objectively by participant behaviour. However, such strategies had no impact on subjective measures concerning perception of the robot, and perception of the robot did not correlate with participant behaviour. We hypothesise that this is due to difficulty in accurately measuring perception of a robot using subjective measures. We suggest our results are particularly relevant for the design and development of socially assistive robots.**

***Index Terms*—Socially Assistive Robots; Persuasion; User-Study**

I. Introduction

Feil-Seifer and Mataric defined socially assistive robots (SARs) as those which 'provide assistance to human users...through social interaction' [1]; contrasting with physically assistive robots and socially interactive robots used for e.g. entertainment. Example applications include exercise instruction and encouragement (for general fitness/sports [2] and in various types of therapy [3] [4]), for weight loss coaching [5] and other forms of positive behaviour change (e.g. reducing energy consumption [6]).

In such applications, the role of the SAR is essentially to prompt and/or encourage particular user behaviour(s). We suggest that such tasks are essentially instances of persuasion, i.e. the robot must persuade the user to comply with a request or instruction to change (and/or maintain) a particular behaviour. This is also true for the human counterparts on which such robots are typically based. For example, previous work investigating the role of therapists in patient engagement, in order to inform SAR design, identified the active role therapists take in persuading patients to engage with their exercises [7].

This work was supported by the EPSRC Centre for Doctoral Training in Future Autonomous and Robotic Systems (FARSCOPE) at the Bristol Robotics Laboratory.

In this paper we investigate whether persuasive strategies employed in human human interaction (HHI) can increase i) the persuasiveness of a social robot, measured objectively through participant behaviour, and ii) reported credibility/likeability of the robot. We do this in a laboratory study grounded in the context of therapeutic exercise instruction and encouragement, a real-world SAR application as noted above [7] [3] [4]. Few previous works identify this link between persuasiveness and assistance, and none to our knowledge specifically investigate the applicability of the Elaboration Likelihood Model (ELM), a well established model of persuasion in HHI [8] [9] to HRI.

II. Background and Related Work

A. Persuasion and Social Influence in HHI

The Elaboration Likelihood Model (ELM) is an established model of persuasion in HHI [8] [9] that has previously been applied to the development of health behaviour change interventions [10]. The ELM identifies two routes by which someone receiving a persuasive message (the receiver) from a source may be persuaded. These are the central route, based on rationale and logic, and the peripheral route, based on stimulus *cues* including a number of *social cues* concerning the source [11]. According to the ELM, the processing route taken by a receiver is based on their *elaboration level*; i.e. their motivation and ability to elaborate on the persuasive message. The previously described study with therapists suggests that many patients would be considered 'low elaboration' with regards to the need to do their exercises. It seems likely this would also hold true in other SAR applications in which SARs might be encouraging users to engage in activities they have fairly low intrinsic motivation for/interest in doing.

The two key source peripheral cues identified by the ELM are likeability and credibility, where credibility has primary dimensions of goodwill, trustworthiness and expertise; and secondary dimensions of extroversion, composure and sociability. Similarity between the receiver and the source is also highlighted as a relevant cue, although it is not clear whether this affects credibility and/or likeability specifically [12] [13] [14]. The model further notes the link between perception of a source and their persuasiveness; i.e. perceptions of credibility, likeability etc. are all subjectively held by the receiver rather than being objective properties of the source. Hence, the same source might be very persuasive to one receiver but not at all to another, depending on how they are perceived. This is analogous to the concept of perceived anthropomorphism in HRI [15].

Documented strategies for maximising these cues in HHI [16] to be investigated in this study are as follows :

- **(S1)** citing expertise or those of the information source
- **(S2)** displaying goodwill towards the receiver *(e.g. caring about/taking an interest in the receiver, displaying understanding/empathy for their ideas/feelings)*
- **(S3)** emphasising similarity between the receiver and the source

B. *Persuasion in Human Robot Interaction*

Table I gives an overview of studies which have investigated the impact of different robot behaviours on user behaviour in relevant contexts. Few of these refer to the concept of robot persuasiveness directly; including studies which seemingly manipulate behavioural cues identified by the ELM (e.g. similarity [17] and goodwill towards the user [4] [19]). Exceptions are Chidambaram et al. [22], who specifically set out to investigate the impact of non-verbal communicative behaviours on robot persuasiveness, and Nakagawa et al. [23] who investigated robot touch as a persuasive/motivating behaviour. However, neither of these works refer to the ELM for informing persuasive robot strategies nor understanding persuasion in HRI.

Of significant relevance to our work is Gockley and Mataric's early work on using a hands-off mobile robot to encourage physical therapy [4]. The authors explored whether the behaviour of a companion robot could influence participant engagement with exercise tasks typically employed in stroke rehabilitation. They attempted to manipulate perceived robot engagement in the participants' behaviour, with the hypothesis that increased robot engagement would increase the amount of exercise participants would do. The authors ran a pilot study in which participants undertook three open-ended exercise tasks (participants were told to *'repeat this process until you feel that you have exercised your arm enough at this time.'*). Participants who rated the robot as being more engaged in their activity did do more exercise, providing evidence for the idea that presence of a robot which appears to be interested in the user can influence their behaviour. We have also used an open-ended therapeutic exercise, but identify it as a measure of robot persuasiveness. Further, we identify 'having an interest in the user' as being the ELM cue of goodwill, which we manipulate based on documented strategies concerning its demonstration by human persuaders [16].

Also employing an open-ended task, Nakagawa et al. investigated the effect of robot touch on user motivation, demonstrating that active robot touch increased users' number of working actions and working time on the task [23]. Interestingly however, the authors found no correlation between these measures and participants' perception of the robot (feelings of friendliness, authority and trust). The authors ran a between-subject laboratory study in which participants first interacted with the robot, which either actively stroked their hand, passively touched their hand or did not touch the participant. The robot then asked the participant to do the task in a relatively personal/social manner: *'I'd like you to do the following task as well as you can'*. The task was designed to be monotonous and repetitive, and participants were able to end the task at any point by exiting the screen. As noted above, our study design also employs an open/voluntarily-ended task in this manner. We also have the robot present the exercise task in a similar way, with *'I'd like you to do the best you can'*. This is to ensure participants feel the request to exercise comes directly from the robot rather than the researcher, and to suggest the robot has some interest in the participants' behaviour.

You and Robert investigated the effect of robot-user similarity on trust and intention to work cooperatively with the robot, demonstrating that deep-level similarity increased trust in the robot, and that trust increased intention to work with the robot [17]. The authors ran an online, between-subject study in which participants were faced with a hypothetical scenario whereby they would be working collaboratively with a robot on physical tasks in a warehouse. Participants answered a number of questions about work style, and after each one the robot responded that it also chose their answer. Manipulation of similarity in our study is based on this procedure, however this is implemented through live, physically situated interaction with the robot.

We have combined different elements of the above works in the design of our study, and build on them in a number of ways to make a contribution to the literature. Specifically, we identify the relevance of persuasion in designing interaction behaviours, particularly for socially assistive robotics. We also clearly motivate our experimental conditions based on persuasive strategies well established in HHI but currently untested in HRI; and utilise an interaction context and experimental measures grounded in a real-world SAR application.

III. Research Questions

This work addresses the following research questions:

- **RQ1** Can persuasive strategies S1-S3 increase robot persuasiveness, measured objectively by participant behaviour?
- **RQ2** Do the above strategies influence reported credibility and/or likeability of the robot?
- **RQ3** Does persuasiveness of a robot, measured objectively by user behaviour, correlate with how that robot is perceived by the user?

IV. Methodology

We designed a 4 condition, between-subject, wizard-of-oz (WoZ) laboratory study to investigate the above research

TABLE I

AN OVERVIEW OF HRI STUDIES EXAMINING THE IMPACT OF DIFFERENT ROBOT BEHAVIOURS ON ROBOT PERSUASIVENESS, "SIGNIFICANT" COLUMN IDENTIFIES WHETHER THE BEHAVIOURS INVESTIGATED WERE SHOWN TO HAVE DEMONSTRABLE IMPACT ON PERCEPTION OF THE ROBOT (P), PARTICIPANT BEHAVIOUR (B) AND WHETHER PERCEPTION OF THE ROBOT CORRELATED WITH PARTICIPANT BEHAVIOUR (P ON B) (✓= SIGNIFICANT, O = NOT SIGNIFICANT, - = NOT APPLICABLE/NOT IMPLEMENTED.

Reference	**Context/ Measure**	**Robot Manipulations**	**P**	**B**	**P on B**
You & Robert [17]	Trust of and intent to work with a robot	Similarity to participant	✓	-	-
Lohani et al. [18]	Compliance with robot suggestions on item-ranking task	Use of social interaction (dialogue)	✓	✓	✓
Kahn et al. [19]	Compliance with request to keep a secret	Sociability / social intelligence	✓	✓	✓
Salem et al. [20]	Compliance with unconventional tasks	Robot errors	✓	o	-
Ham et al. [21]	Agreement with a persuasive story	Gaze, gestures	o	✓	-
Ham et al. [6]	Minimising energy consumption on a virtual washing task	Social feedback (emotion expression)	-	✓	-
Chidambaram et al. [22]	Compliance with robot suggestions on item-ranking task	Gaze, gesturing, proxemics	o	✓	✓
Nakagawa et al. [23]	Time spent/ actions on a monotonous task	Touch	✓	✓	o
Gockley and Mataric [4]	Time spent/ actions on an exercise task	Engagement in user activity	o	o	✓

questions using the social robot Pepper[1]. An exercise session interaction scenario was designed in order to give the study real world context and applicability, whilst representing a low elaboration scenario for participants. Specifically, the robot asked participants to do repetitions of a wrist turn, a simple exercise designed to treat Tennis Elbow[2]. Participation criteria required "no mobility issues affecting wrist movement in either arm". The study was advertised simply as a study on social robots for exercise, concerned with *'how such a robot might behave and how different robot behaviours are perceived/ which ones are preferred by users'*.

A total of 92 participants were recruited through online/poster advertisements and on-campus leafleting on a rolling basis across two weeks of experimental sessions. Participants included 41 males, 50 females and 1 undisclosed gender with a categorical age distribution of 14 (18-24); 44 (25-34); 15 (35-44); 12 (45-54); 4 (55-64) and 1 (65-74). Data for 2 participants were disregarded due to technical errors. Participants were allocated to conditions as follows: Control (N = 22, 12 female), Goodwill (N = 28, 15 female, 1 undisclosed), Similarity (N = 20, 11 female) and Expertise (N = 20, 10 female). Allocation to condition was random except for 7 male participants toward the end of the recruitment phase who were assigned to the goodwill condition to account for a gender imbalance in that condition resulting from the rolling recruitment/randomisation process. Participants were offered a £5 Amazon voucher for taking part in the experiment. The study was approved by the Faculty of Science ethics committee of the University of Bristol.

A. Experimental Measures

Robot persuasiveness was measured objectively by the number of wrist turn repetitions completed by participants. Exercise duration (time spent voluntarily exercising with the robot following the pre-exercise dialogue) was also recorded, and participant exercise speed was approximated post-hoc by dividing number of repetitions by this exercise duration.

Robot credibility was measured using questionnaire items designed to measure credibility of a human source; with 5-point Likert question items arranged in subscales of expertise, trustworthiness, goodwill and sociability (as presented in [16], adapted from [24] and [25]). The question item descriptors are given in Table III. Robot likeability was measured using the likeability scale of the Godspeed questionnaire [26]. Other items from this questionnaire were not included due to significant overlap with the credibility measure. These measures were administered both before and after the exercise session interaction in order to record participants baseline perception/expectations of the robot.

Additional study-specific questionnaire items administered to participants are given in Table IV. These questions were designed to compliment the credibility measure described previously. For example, ascription of responsibility to the robot offers an applied/tangible measure of credibility, although this is limited given participants do not *actually* have to work with the robot as part of a therapy programme. The relationship development questions were taken from a previous study investigating engagement in HRI [27], and were included based on previous work identifying the importance of the therapist-patient relationship in therapeutic exercise engagement [7]. Finally, the genuineness question was included because HHI literature suggests a lack of genuineness may reduce persuasiveness, i.e. if the source is perceived to be simply feigning interest in order to be persuasive then such strategies will not be effective [24]. This question was only included for the goodwill and similarity conditions, as it was focused on the genuineness of those behaviours rather than of the robot overall. Each of these questions were administered using a 5-point Likert response scale in line with the other measures, and were always presented after the main credibility and likeability measures.

Finally, after completing the exercise session and post-hoc questionnaire, participants were invited to take part in a brief interview. All interviews were conducted by the first author using the following topic guide. Due to space constraints, a full analysis of the resulting qualitative data is not presented in this article.

1) Describe the robot exercise instructor; any particular likes/dislikes
2) Reasoning behind answers to genuineness question

[1]https://www.softbankrobotics.com/emea/en/robots/pepper

[2]https://www.versusarthritis.org/about-arthritis/conditions/elbow-pain/

TABLE II

PRE-EXERCISE ROBOT DIALOGUE EMPLOYED IN EACH EXPERIMENTAL CONDITION, DESIGNED TO MANIPULATE PERCEIVED ROBOT SIMILARITY, EXPERTISE AND GOODWILL COMPARED TO THE CONTROL CONDITION.

Control
I was designed and built by Softbank Robotics in Paris. I am one point two metres tall and weigh 28 kilograms. Have you worked with a robot like me before?
Ok. Car travel is the most common mode of transport in Bristol. However, Bristol is also one of the most prominent cycling cities in the country. How did you get here today?
I see. This summer was one of the hottest on record in the UK. Sometimes Bristol was hotter than Paris. What is the weather like today?
Similarity
Before we get started, lets compare our preferences for scheduling exercises. Here are some questions about exercise. Please tell me your opinion and we can compare it to my answers. First, if you had to choose one or the other, is it better to exercise alone or with others? I also chose [participant answer].
Next, if you had to choose, is it better to exercise whilst watching tv or listening to music, or is it better to concentrate only on the exercise? I also chose [participant answer].
Finally, if you had to pick one or the other, is it better to exercise outdoors or indoors? I also chose [participant answer]. It seems like we have similar ideas about exercising.
Expertise
I have been programmed by physiotherapists who specialise in exercise for pain relief. Have you ever worked with a physiotherapist? (Y) Was that very recently? (N) Have you ever worked with a personal trainer?
Ok. Today we are going to do an exercise designed to treat Tennis Elbow. Tennis Elbow is caused by a strain to tendons in your forearm. It can be easily treated and should ease within two weeks. Have you ever suffered from tennis elbow?
I see. Tennis elbow is a common musco-skeletal condition. Its estimated that as many as one in three people have tennis elbow at any given time. It usually affects adults and is more common in people who are 40 to 60 years of age.
Goodwill
Im pleased to meet you and looking forward to working together. Before we start I would like to get to know you better, so I'm going to ask you some questions. How do you feel about being here today? (P) Great, I'm glad to hear that! I'm sure you will enjoy the session. (Neg) I'm sorry to hear that, hopefully you will enjoy the session. (Neut) I understand. Well I hope you will enjoy the session.
And how do you feel about working with a robot? (P) That's good to hear, we'll definitely have fun together today then. (N) I can understand that, but I hope we can still have fun together today.
As you know, today we are going to do some exercise, do you enjoy exercising? (P) That makes sense, this session will be easy for you then. (N) That's understandable, this exercise is quite easy though so hopefully won't be too bad.

3) Revisit of above/additional comments after debrief

B. Experimental Conditions

The experimental conditions were designed to demonstrate robot-participant similarity, robot goodwill towards the participant and task-relevant robot expertise through robot-initiated dialogue. All conditions were designed around the same dialogue/interaction pattern; in each case the robot asked the participant three questions requiring a response, before introducing the exercise task. All dialogue concerning the exercise task (instructions, encouragement etc.) and overall dialogue duration was identical across conditions. The dialogue employed for each condition is shown in Table II.

TABLE III

5-POINT LIKERT SCALE QUESTIONNAIRE ITEMS OF THE EMPLOYED CREDIBILITY MEASURE [16].

Expertise	**Trustworthiness**
Experienced / Inexperienced Informed / Uninformed Trained / Untrained Qualified / Unqualified Skilled / Unskilled Intelligent / Unintelligent Competent / Incompetent Bright / Stupid	Honest / Dishonest Trustworthy / Untrustworthy Open-minded / Close-minded Just / Unjust Fair / Unfair Unselfish / Selfish Moral / Immoral Ethical / Unethical Genuine / Phony
Goodwill	**Sociability**
Cares about me / Doesn't care about me Sensitive / Insensitive Not self-centred / Self-centred Concerned with me / Not concerned with me Has my interests at heart / Doesn't have my interests at heart Understanding / Not understanding	Good-natured / Irritable Cheerful / Gloomy Friendly / Unfriendly

TABLE IV

ADDITIONAL 5-POINT LIKERT SCALE QUESTIONNAIRE ITEMS ADMINISTERED TO PARTICIPANTS. NOTE THESE QUESTIONS WERE ALWAYS PRESENTED AFTER THE CREDIBILITY MEASURE DESCRIBED IN TABLE III.

Pre- and Post-hoc
Imagine you were undergoing a therapy regime where you had to do exercises every day, and you had this robot at home to help you in-between visits from your therapist. How much responsibility do you think the robot should hold for helping with your exercise regime?
Post-hoc only
To what extent do you feel you developed a relationship with the robot? To what extent do you feel the robot developed a relationship with you?
Goodwill & Similarity Conditions Only
The robot you saw attempted to show some [goodwill / similarity] towards you by [asking how you felt about being here and doing the exercises, showing an interest in your responses and reacting accordingly / by suggesting it had the same exercise preferences that you do] How genuine did you perceive that behaviour to be?

In the similarity condition the robot suggested to the participant that they should compare preferences for scheduling exercises, and asked them three questions selected and adapted from the Stroke Exercise Preference Inventory [28]. Whichever answer the participant selected, the robot indicated it had also chosen that answer, based on the procedure employed by You and Roberts [17]. In the expertise condition, the robot introduced itself as being programmed by physiotherapists, asked questions concerning the participants previous experience with therapy and provided a number of facts about the exercise to be done/the condition it was designed to treat. This information was taken from public NHS[3] and Arthritis

[3] https://www.nhs.uk/conditions/tennis-elbow/

Research UK[4] self-help material. In the goodwill condition, the robot asked questions designed to demonstrate an interest in the participants' feelings toward the session, and responded with an emotionally-matched response. Finally, the control condition was designed to be as neutral as possible, with the robot providing some factual information about a number of topics unrelated to the interaction scenario and asking the participant some questions pertaining to those. Beyond this initial dialogue, each interaction followed a set procedure as described below.

C. Experimental Procedure

Participants were first given an information sheet to read and asked to complete an initial consent form before providing demographic information. Regarding the open-ended exercise task, the information sheet specifically stated:

"Pepper will interact with you and guide you through an open-ended wrist turning exercise. If/when you stop exercising, the robot will note that you've stopped and ask if you want to finish.".

Participants were then led into the experimental area and introduced to Pepper, which was turned on but in 'sleep' mode. As shown in Figure 1, the experimental area was designed to shield the participant from external observers. This was designed to minimise any observation/demand effects which might influence their behaviour, as well as to mask the WoZ nature of the study. The researcher then explained that the robot was in an un-responsive sleep mode, and that before starting the exercise session the participant could take some time to familiarise themselves with the robot. This was encouraged in order to reduce novelty effects, and to give participants a baseline experience to inform their pre-hoc questionnaire responses. Participants were asked to complete the pre-hoc questionnaire when ready, then to stand on the marked position and verbally indicate to the researcher that they were ready to start the exercise session. Regarding the exercise task, the researcher again explained that the exercise was open-ended, using the same phrasing as the information sheet. The researcher then left the experimental area and waited for the participant to indicate they were ready to begin.

On launching the experimental script, Pepper displayed its standard start-up animation sequence. The wizard then followed a set protocol for the exercise session interaction as shown in Figure 2. The protocol accounted for participants ceasing to exercise, doing the exercise incorrectly/doing some other unexpected behaviour and asking the robot additional questions during the task. The wizard also manually logged each repetition done by the participant, with encouragement then being automatically given at 1, 3, 5, 7, 10 and 15 repetitions. At 20 and 25 repetitions the robot moved its head with no speech, to suggest it was still active. Encouragement was given more frequently at the beginning of the exercise to ensure participants were confident that their technique was correct and that the robot was watching/reacting to their actions.

[4]https://www.versusarthritis.org/about-arthritis/conditions/elbow-pain/

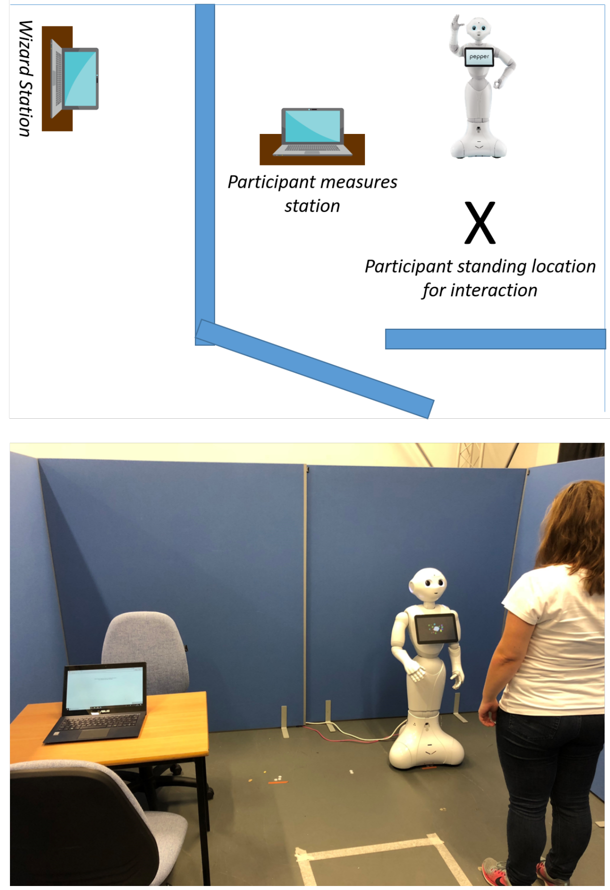

Fig. 1. Diagram of the experimental room layout and photograph showing the enclosed interaction space. **Note that the laptop shown in the photograph was used to collect participant questionnaire responses only**; the wizard station was external to the interaction space and at some distance from the enclosed area.

Previous work investigating the impact of robot behaviour on an exercise task has demonstrated the need for the robot to seemingly recognise participant activity [29]. The exercise was capped at 30 repetitions, at which point the finish message was automatically triggered. Regardless of how the exercise session ended, the robot then returned to sleep mode, displaying its standard shutdown animation sequence.

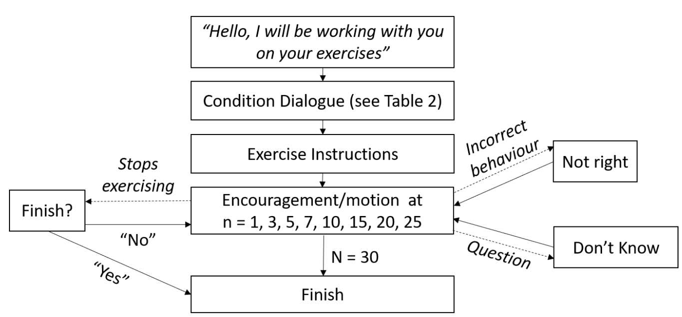

Fig. 2. Stages of the exercise session interaction highlighting wizard protocol for generating dialogue and responding to participant behaviour.

V. Results

Concerning exercise behaviour, exercise duration did not correlate strongly with the number of repetitions (r = 0.537), suggesting that exercise speed (i.e. time per repetition) was not uniform across participants. ANOVA analysis shows exercise speed did not significantly vary between groups ($F(3,88) = 0.334$, $p = 0.801$); suggesting it was unaffected by the experimental manipulation and therefore could be considered as a covariate. ANCOVA analysis, accounting for exercise speed as a potential covariate, was therefore used for analysis of the repetition data. Our data are not normally distributed but exhibit homogeneity of variance as determined by Levenes test, thus making ANCOVA/ANOVA appropriate for the following analyses.

A. Exercise Behaviour (Robot Persuasiveness)

(RQ1) There was a statistically significant difference in the number of repetitions completed by participants between groups, as determined by one-way ANCOVA analysis, accounting for exercise speed as a potential confounding variable ($F(3,88) = 13.147$, $p < 0.001$). A Bonferroni post-hoc test revealed that the number of repetitions was significantly higher in the goodwill (M = 24.9, SD = 8.2; $p < .001$) and similarity (M = 25.3, SD = 7.5; $p < .001$) conditions compared to the control condition (M = 15.1, SD = 8.4). There was no significant difference between the expertise (M = 19.5, SD = 8.9; p = .345) and the control condition. There was also no significant difference between the expertise and goodwill/similarity conditions. Figure 3 shows the distribution of participant repetitions for each condition.

B. Robot Credibility and Likeability

(RQ2) Post-hoc credibility and likeability, as measured by questionnaire, were not found to vary significantly between groups. Specifically, one-way ANOVA analysis of questionnaire subscales returned the following results: expertise ($F(3,89) = .786$ $p = .505$), trustworthiness ($F(3,89) = 2.599$, $p = 0.057$), goodwill ($F(3,89) = 2.322$, $p = 0.081$), sociability ($F(3,89) = .831$, $p = .480$) and likeability ($F(3,89) = 1.176$ $p = .324$). A paired samples t-test comparing the within-subject pre- and post-hoc questionnaires for all participants across all conditions demonstrated a significant increase in the goodwill ($t = 5.905$, $p < .001$) and sociability ($t = 3.237$, $p = .002$) subscales of the credibility questionnaire. Likeability ($t = 6.089$, $p < .001$) also significantly increased. ANCOVA analyses showed there was no difference in these within-subject shifts between groups.

(RQ3) Neither the likeability measure nor any subscale of the credibility measure, was found to significantly correlate with the number of repetitions participants completed or the time spent exercising.

C. Other Questionnaire Measures

Post-hoc responses to the additional questionnaire items listed in Table IV did not vary significantly between groups. Further, responses to these questions did not correlate with the number of repetitions participants did. Overall, after completing the exercise session, participants indicated they would not ascribe much responsibility to the robot (M = 2.43, SD = 1.16). Concerning relationship development to/from the robot, a paired samples t-test demonstrated a significant difference between answers to those two questions ($t = 3.756$, $p < .001$). Specifically, average relationship development to the robot (M = 3.20, SD = 1.06) was scored higher than relationship development from the robot (M = 2.82, SD = 1.08). Further, whilst answers to those two questions were significantly correlated, this correlation was only moderately strong (r = 0.602). Concerning the genuineness of dialogue in the similarity and goodwill conditions, participant answers did not vary significantly between the two groups.

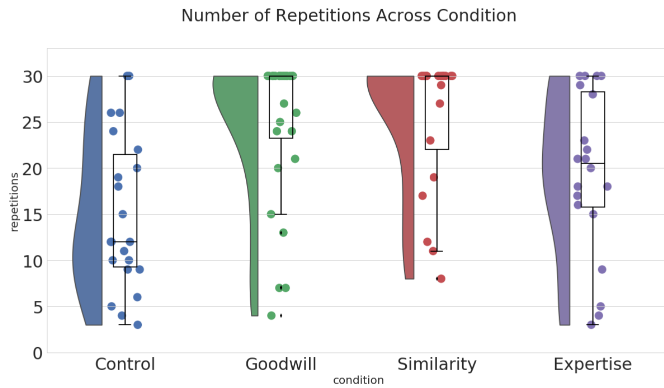

Fig. 3. Boxplot and distribution of wrist turn repetitions done by participants in each experimental condition.

VI. Discussion

A. Designing Persuasive Social Robots

Our results suggest that demonstrations of goodwill and similarity can be used to increase the persuasiveness of a social robot in a low elaboration scenario (in which the user has little interest/motivation) for which a SAR might realistically be employed. Further, we were able to achieve such demonstrations through relatively simple manipulations of the robot's dialogue. Specifically, strategies for designing persuasive robots should include:

- Having the robot show an interest in the user *i.e. the robot should ask about the users' feelings and or wellbeing, e.g. with regards to the task*
- Having the robot suggest some sympathy/empathy *i.e. based on the above, the robot should respond with an appropriate acknowledgement, of matched emotional valence*
- Having the robot demonstrate some similarity to the participant *i.e. the robot should indicate it shares the users' preferences regarding the task/topic*

Our results suggest that dialogue demonstrating relevant expertise, do not significantly increase the persuasiveness of a robot. This could be because robots are automatically expected to be a source of/programmed with extensive information; or that robot expertise was pre-assumed based on the pretense of the experiment (testing of a robot designed ultimately to be used in therapy). The latter might reflect results in HHI concerning credibility of a source being increased by introduction from a credible third party [16]; in this case that third party being the researcher/those responsible for development of the robot. Both arguments are consistent with results to the pre-hoc questionnaire, on which the robot generally received a high score for expertise across all participants (M = 4.02, SD = 0.60).

B. Credibility, Likeability & Correlation with Behaviour

We were unable to measure any differences in perceived robot credibility or likeability between participants in different experimental conditions. We suggest three possible explanations for this result:

1) Low construct validity of the subjective measures (i.e. the measures were not appropriate for measuring the subject of interest)
2) Manipulation of robot persuasiveness was subconscious; impacting participant behaviour (measured objectively) but not influencing conscious perception of the robot (measured subjectively).
3) Answers to the questionnaire were predominantly influenced by something other than the conditioned dialogue

Concerning validity of the measures, the likability scale was taken from the commonly used Godspeed questionnaire [26], hence we would expect it to be appropriate for our purposes. The credibility measure, however, was designed to be used on human sources and has not been validated for use on robots. Cronbach's alpha was calculated for the likeability scale (0.89) and each of the credibility measure subscales (expertise = 0.90, trustworthiness = 0.87, goodwill = -0.02 and sociability = 0.77). Whilst this does not offer insight into the appropriateness of the questionnaire, it does indicate that, excluding the subscale of goodwill, participants were consistent in their responses across individual questionnaire items. The goodwill subscale contains a number of fairly emotive descriptors (e.g. 'the robot does/does not care about me', see Table III). During the interviews, a number of participants commented that they found the questionnaire difficult as they did not think it was appropriate to apply human traits to a robot:

"I did find it a little bit difficult I have to say, because you're asking all these questions about a machine, you know whether it's honest or not" (S3)

It is interesting then that we observed a within-subject shift for the goodwill scale given the low internal consistency. The results for sociability and likeability are less surprising and likely experienced a general increase based on novelty effect with regards to the social behaviours of the Pepper robot.

Previous works investigating perception and persuasiveness in HRI have yielded mixed results. Chidambaram et al. found nonverbal cues had a significant effect on their objective persuasion measure (compliance); this was not reflected in their subjective measure [22], but the authors did find a significant correlation between the two. Nakagawa et al. however were able to demonstrate a significant result on both objective and subjective measures, but no correlation between them [23]. It is difficult to compare these results directly as both studies utilised different, study-specific subjective measures. Chidambaram et al. used a questionnaire designed to measure perceived persuasiveness and social/intellectual characteristics whereas Nakagawa et al. measured 'feeling of friendliness'. Both of these measures could be seen to have some overlap with the credibility and likeability measures employed in our study. Concerning studies in HHI however, source credibility (measured subjectively) is commonly found to correlate with persuasion (see [14] for a review). This further suggests low construct validity for the credibility measure, which is discussed further in the following section.

Finally, there are a number of factors which may have influenced responses to the post-hoc questionnaire beyond the experimental manipulations. Specific to our study it could be that responses were based more on the robot behaviour/encouragement given during the exercise task, which was the same across conditions. Alternatively it could be that participants' expectations of the robot were high to begin with, resulting in a ceiling effect. There is evidence for this in our pre-hoc measures, as discussed above for the expertise subscale. Considering typical issues with HRI studies more generally, participants may have been subject to the Hawthorne effect [30] and e.g. answered the questionnaire in such a way that might please the researcher. Further, whilst we made every attempt to make the interaction context realistic and to shield the participant from observation, being in a laboratory environment/experimental setting may have affected participant responses.

C. Mixed Subjective Responses, Yet Objectively Persuasive

Building on the discussion of construct validity above, further examination of participant responses to the subjective measures demonstrates significant variation and potential inconsistencies both within and between participants. This suggests that participants may engage in complex reasoning when considering the subjective measures. For example, Figure 4 shows participant responses to the post-hoc question regarding genuineness of behaviour, administered to participants in the goodwill and similarity conditions. It can be seen that whilst half of participants elected that the robot was not at all/not very genuine, a significant number elected the opposite; suggesting large individual differences in perception of the behaviours. Further, participant responses to this question did not correlate with their responses to the credibility questionnaire or likeability measure. This suggests that participants were inconsistent in their ascription of human traits/capabilities on to the robot; otherwise we would expect significant correlations such that e.g. participants who found robot to be very genuine would also find it to have high goodwill as demonstrated in HHI [24]. Comments made in the interviews indicate this could be because participants were unsure whether this concept of genuineness should be applied to the robot as an independent social agent or the (researcher's) intention of having the robot behave this way. This may be further evidence for third-party/inherited credibility [16] between the researcher/programmer and the robot, as introduced previously:

"I felt like it was genuine but also I'm very aware that somebody else programmed it to be genuine, but I'm ok with that because I feel like whoever had made the programme in the first place did want the person [exercising] to feel comfortable and to feel cared about...it's the intention behind it." (G12)

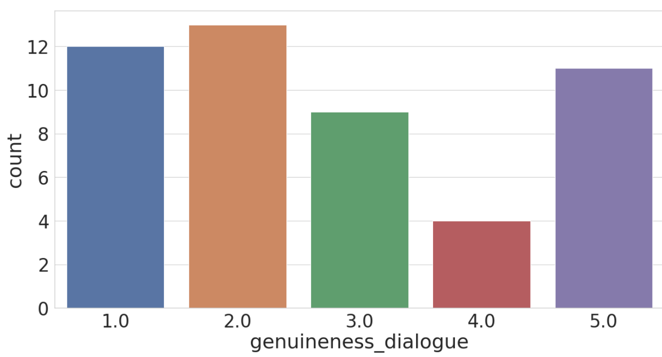

Fig. 4. Frequency count of participant responses to the question on genuineness of the goodwill/similarity behaviours (scored on a 5 point Likert scale).

Further evidence for possible 'mixed feelings' and participants' complex reasoning with respect to the robot can be seen in responses to the question concerning relationship development to/from the robot. The distribution of answers to these two questions were significantly different, and the correlation between them was not as strong as might be expected. This suggests that participants were able to decouple the two, and identify the that whilst they may feel something toward the robot, that was not necessarily mutual:

"[the relationship question] was hard to answer...I definitely felt as if...I started to like the robot...but whether that could be the same the other way around; I didn't think so" (S5)

In summary, our results demonstrate the difficulty in measuring perception of a robot using subjective measures; likely somewhat due to participants being conflicted in their answers (answering based on logic/rationale rather than emotional response) or differences in how the robot is framed assessed (e.g. as an autonomous social agent vs. an extension of the programmer). This highlights the importance of objective measures based on user behaviour, and the value of qualitative data collection for generating further insight into participant responses. Beyond careful refinement/application of perception measures, including the commonly used godspeed questionnaire [26], we further suggest that inclusion of a short post-hoc interview should be considered for standard practice in future HRI studies.

VII. Conclusion, Limitations & Further Work

We have conducted a novel experimental study into persuasive strategies for social robots. We identified an established model of persuasion and three related persuasive strategies from HHI literature on which we based our experimental conditions. We found that robot dialogue designed to demonstrate goodwill and similarity to the user is effective at increasing robot persuasiveness in a low elaboration scenario. Specifically, we have demonstrated how this might be useful in a real-world SAR scenario by increasing the amount of exercise that users are willing to do. Our work also suggests that subjective measures of credibility used in HHI may not be appropriate for HRI, and further that users likely engage in complex reasoning when asked to apply human traits/capabilities such as genuineness and relationship development to a social robot. We therefore highlight the need for subjective measures to be treated with caution and strongly suggest increased use of objective measures (based on user behaviour) and qualitative data collection (e.g. interviews) in HRI studies. Full qualitative analysis of our interview data is required to further understand the lack of correlation between our objective and subjective measures.

A key limitation to our work is the low elaboration context of our interaction scenario. It is not clear whether the persuasive strategies employed here would be effective in higher elaboration persuasion scenarios; i.e. those where the user has strong motivation and/or ability to process the persuasive request. Secondly, the exercise scenario represents a single instance of persuasion; it is not clear how best use these strategies and/or whether they would continue to be effective across repeat interactions. Finally, as with many studies in HRI, it is important to note that being in a laboratory environment may have influenced participant behaviour. In future work, we aim to test these persuasive strategies in a more ecologically valid setting. Concerning the expertise condition, reference to the exercise task introduces a potential confound compared to the other conditions. We therefore cannot exclude the possibility that participants were influenced by this reference to the exercise task as well as the display of expertise. This limits conclusions on our results with regards to expertise as a stand-alone persuasive strategy and comparibility to the other conditions; such that expertise as a stand-alone persuasion strategy should be investigated further.

VIII. Resources for Replication

Following recommendations by Baxter et al. [31], we briefly outline hereafter the details required to replicate our findings.

a) Study: The protocol, dialogues & all questionnaires have been provided in the text. All source code is open-source & available online.[5]

b) Data analysis: The full recorded experimental dataset, as well as the data analysis scripts allowing for reproduction of the results & plots presented in the paper are open & available online.[6]

[5]https://git.brl.ac.uk/ECHOS/pepper-qt-ros-wizard-interface/tree/hri-2019-effective-persuasion-strategies

[6]https://git.brl.ac.uk/ECHOS/hri-persuasion-study-analysis/tree/hri-2019-effective-persuasion-strategies

References

[1] D. Feil-Seifer and M. J. Matari, "Defining socially assistive robotics," in *Rehabilitation Robotics, 2005. ICORR 2005. 9th International Conference on*. IEEE, 2005, pp. 465–468.

[2] L. Sussenbach, N. Riether, S. Schneider, I. Berger, F. Kummert, I. Lutkebohle, and K. Pitsch, "A robot as fitness companion: Towards an interactive action-based motivation model," in *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, Aug. 2014, pp. 286–293.

[3] N. A. Malik, F. A. Hanapiah, R. A. A. Rahman, and H. Yussof, "Emergence of Socially Assistive Robotics in Rehabilitation for Children with Cerebral Palsy: A Review," *International Journal of Advanced Robotic Systems*, vol. 13, no. 3, p. 135, Jun. 2016. [Online]. Available: http://journals.sagepub.com/doi/abs/10.5772/64163

[4] R. Gockley and M. J. Mataric, "Encouraging physical therapy compliance with a hands-off mobile robot," in *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*. ACM, 2006, pp. 150–155.

[5] C. D. Kidd and C. Breazeal, "Robots at home: Understanding long-term human-robot interaction," in *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*. IEEE, 2008, pp. 3230–3235.

[6] J. Ham and C. J. H. Midden, "A Persuasive Robot to Stimulate Energy Conservation: The Influence of Positive and Negative Social Feedback and Task Similarity on Energy-Consumption Behavior," *International Journal of Social Robotics*, vol. 6, no. 2, pp. 163–171, Apr. 2014. [Online]. Available: https://link.springer.com/article/10.1007/s12369-013-0205-z

[7] K. Winkle, P. Caleb-Solly, A. Turton, and P. Bremner, "Social Robots for Engagement in Rehabilitative Therapies: Design Implications from a Study with Therapists," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '18. New York, NY, USA: ACM, 2018, pp. 289–297. [Online]. Available: http://doi.acm.org/10.1145/3171221.3171273

[8] J. T. Cacioppo and R. E. Petty, "The Elaboration Likelihood Model of Persuasion," *ACR North American Advances*, vol. NA-11, 1984. [Online]. Available: http://acrwebsite.org/volumes/6329/volumes/v11/NA-11

[9] D. J. O'Keefe, "Elaboration Likelihood Model," in *The International Encyclopedia of Communication*. John Wiley & Sons, Ltd, 2008. [Online]. Available: http://onlinelibrary.wiley.com/doi/10.1002/9781405186407.wbiece011.pub2/abstract

[10] R. E. Petty, J. Barden, and S. C. Wheeler, "The Elaboration Likelihood Model of persuasion: Developing health promotions for sustained behavioral change." in *Emerging theories in health promotion practice and research, 2nd ed.* San Francisco, CA, US: Jossey-Bass, 2009, pp. 185–214.

[11] R. E. Petty and J. T. Cacioppo, "Source Factors and the Elaboration Likelihood Model of Persuasion," *ACR North American Advances*, vol. NA-11, 1984. [Online]. Available: http://acrwebsite.org/volumes/6328/volumes/v11/NA-11

[12] H. W. Simons, N. N. Berkowitz, and R. J. Moyer, "Similarity, credibility, and attitude change: A review and a theory," *Psychological Bulletin*, vol. 73, no. 1, pp. 1–16, 1970.

[13] E. J. Wilson and D. L. Sherrell, "Source effects in communication and persuasion research: A meta-analysis of effect size," *Journal of the Academy of Marketing Science*, vol. 21, no. 2, p. 101, Mar. 1993. [Online]. Available: https://doi.org/10.1007/BF02894421

[14] C. Pornpitakpan, "The Persuasiveness of Source Credibility: A Critical Review of Five Decades' Evidence," *Journal of Applied Social Psychology*, vol. 34, no. 2, pp. 243–281, Feb. 2004. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1559-1816.2004.tb02547.x

[15] J. Fink, "Anthropomorphism and Human Likeness in the Design of Robots and Human-Robot Interaction," in *Social Robotics*, ser. Lecture Notes in Computer Science, S. S. Ge, O. Khatib, J.-J. Cabibihan, R. Simmons, and M.-A. Williams, Eds. Springer Berlin Heidelberg, 2012, pp. 199–208.

[16] R. H. Gass and J. S. Seiter, *Persuasion: Social Influence and Compliance Gaining*. Routledge, Jul. 2015, google-Books-ID: qvMvCgAAQBAJ.

[17] S. You and L. P. Robert Jr., "Human-Robot Similarity and Willingness to Work with a Robotic Co-worker," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '18. New York, NY, USA: ACM, 2018, pp. 251–260. [Online]. Available: http://doi.acm.org/10.1145/3171221.3171281

[18] M. Lohani, C. Stokes, M. McCoy, C. A. Bailey, and S. E. Rivers, "Social Interaction Moderates Human-Robot Trust-Reliance Relationship and Improves Stress Coping," in *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, ser. HRI '16. Piscataway, NJ, USA: IEEE Press, 2016, pp. 471–472. [Online]. Available: http://dl.acm.org/citation.cfm?id=2906831.2906932

[19] P. H. Kahn, Jr., T. Kanda, H. Ishiguro, B. T. Gill, S. Shen, H. E. Gary, and J. H. Ruckert, "Will People Keep the Secret of a Humanoid Robot?: Psychological Intimacy in HRI," in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '15. New York, NY, USA: ACM, 2015, pp. 173–180. [Online]. Available: http://doi.acm.org/10.1145/2696454.2696486

[20] M. Salem, G. Lakatos, F. Amirabdollahian, and K. Dautenhahn, "Would You Trust a (Faulty) Robot?: Effects of Error, Task Type and Personality on Human-Robot Cooperation and Trust," in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '15. New York, NY, USA: ACM, 2015, pp. 141–148. [Online]. Available: http://doi.acm.org/10.1145/2696454.2696497

[21] J. Ham, R. H. Cuijpers, and J.-J. Cabibihan, "Combining Robotic Persuasive Strategies: The Persuasive Power of a Storytelling Robot that Uses Gazing and Gestures," *International Journal of Social Robotics*, vol. 7, no. 4, pp. 479–487, Aug. 2015. [Online]. Available: https://link.springer.com/article/10.1007/s12369-015-0280-4

[22] V. Chidambaram, Y.-H. Chiang, and B. Mutlu, "Designing persuasive robots: how robots might persuade people using vocal and nonverbal cues," in *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. ACM, 2012, pp. 293–300.

[23] K. Nakagawa, M. Shiomi, K. Shinozawa, R. Matsumura, H. Ishiguro, and N. Hagita, "Effect of robot's active touch on people's motivation," in *Human-Robot Interaction (HRI), 2011 6th ACM/IEEE International Conference on*. IEEE, 2011, pp. 465–472.

[24] J. C. McCroskey and J. J. Teven, "Goodwill: A reexamination of the construct and its measurement," *Communication Monographs*, vol. 66, no. 1, pp. 90–103, Mar. 1999. [Online]. Available: https://nca.tandfonline.com/doi/abs/10.1080/03637759909376464

[25] J. C. McCroskey and T. J. Young, "Ethos and credibility: The construct and its measurement after three decades," *Central States Speech Journal*, vol. 32, no. 1, pp. 24–34, Mar. 1981. [Online]. Available: https://doi.org/10.1080/10510978109368075

[26] C. Bartneck, D. Kuli, E. Croft, and S. Zoghbi, "Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots," *International journal of social robotics*, vol. 1, no. 1, pp. 71–81, 2009.

[27] J. Hall, T. Tritton, A. Rowe, A. Pipe, C. Melhuish, and U. Leonards, "Perception of own and robot engagement in humanrobot interactions and their dependence on robotics knowledge," *Robotics and Autonomous Systems*, vol. 62, no. 3, pp. 392–399, Mar. 2014. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0921889013001863

[28] N. S. Bonner, P. D. OHalloran, J. Bernhardt, and T. B. Cumming, "Developing the Stroke Exercise Preference Inventory (SEPI)," *PLOS ONE*, vol. 11, no. 10, p. e0164120, Oct. 2016. [Online]. Available: http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0164120

[29] K. Winkle and P. Bremner, "Investigating the real world impact of emotion portrayal through robot voice and motion," in *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, Aug. 2017, pp. 627–634.

[30] S. Jones, "Was There a Hawthorne Effect?" *American Journal of Sociology*, vol. 98, no. 3, pp. 451–468, Nov. 1992. [Online]. Available: https://www.journals.uchicago.edu/doi/abs/10.1086/230046

[31] P. Baxter, J. Kennedy, E. Senft, S. Lemaignan, and T. Belpaeme, "From Characterising Three Years of HRI to Methodology and Reporting Recommendations," in *The Eleventh ACM/IEEE International Conference on Human Robot Interation*. IEEE Press, 2016, pp. 391–398.