

Exploring Minimal Nonverbal Interruption in HRI

Paul Saulnier, Ehud Sharlin, Saul Greenberg

Abstract—Designing robotic behaviours capable of initiating an interruption will be extremely important as robots increasingly interact with people. Consequently, we explore the social impact of a *minimal* set of physical nonverbal cues that can be exhibited by a robot to initiate robot-human interruption: (a) speed of motion, (b) gaze, (c) head movement, d) rotation and (e) proximity to the person. We present two related studies evaluating this set. First, for requirements gathering, we observed the behaviour of interruption between humans, with a human actor attempting to interrupt other humans while being constrained to use only a set of behavioural cues that could be mimicked by a simple nonverbal robot. Next, we programmed a robot to exhibit similar social physical nonverbal cues, and tested their feasibility in a user study of robotic nonverbal interruption across interruption scenarios. Our results show that people were able to interpret interruption urgency from robot behaviour using only minimal nonverbal behavioural cues. These findings contribute to informing future designs of social human-robot interfaces.

I. INTRODUCTION

For robots to work in social settings, both robots and humans must understand each other's behaviours and respond accordingly. This is not yet something that we, as interaction designers, fully understand how to do. We do know that interpersonal behaviour is a complex phenomenon that includes language, tone of voice, gesture, posture, body movements, spatial orientation, physical proximity, eye contact, and facial expression amongst other attributes [7]. The problem is that it is unrealistic (at least for now) for a robot to exhibit this behavioural richness. Thus our general research goal is questioning whether there are *minimal non-verbal behavioural cues* that robots can exhibit to communicate their internal state, and are those cues understandable by people? By *minimal*, we mean that we are interested in determining behaviours that rely on only a few simple physical capabilities present in (or that can be easily added to) most robots.

In this paper, we address a narrowed down subset of the above goal: exploring the process of interruption. For people, interruptions are a normal part of daily life. We change our behaviours to initiate an interruption, where the particular behaviour we exhibit informs the person being interrupted about the importance and urgency of a situation.

Manuscript received March 17, 2010. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC), the Alberta Informatics Circle of Research Excellence (iCORE), and the University of Calgary.

The authors are with the University of Calgary, 2500 University Drive NW, Calgary, Alberta, T2N 1N4 (phone: +1 (403) 210-9502, e-mail: p.saulnier@ucalgary.ca, ehud@cpsc.ucalgary.ca, saul.greenberg@ucalgary.ca).

Our actions are based on our expectations of how others will understand, interpret and ultimately respond to our interruption behaviours.

Designing comprehensible robotic behaviours that are capable of initiating and tuning an interruption will be extremely important as robots increasingly interact with people. We can expect classes of robots that will be capable of using verbal communication to interrupting users. Yet many robots will be non-verbal, and there are likely many situations where robotic voice conversation would be inappropriate. Even in situations when the use of robotic voice is appropriate and feasible, we argue that robots could benefit from mastery of non-verbal interruption cues which can augment verbal cues and make the overall experience more acceptable for the person who is being interrupted. Thus we are interested in exploring a fundamental layer of social interruption which involves physicality, movement, interpersonal distance, gaze, etc. We argue that most human to human interruptions, verbal or not, include this physical layer, and that people use this information to help not only interpret another's action as an interruption, but also the urgency of the interruption. Consequently, we believe that designers of social robotic interfaces will need to effectively harness this physical layer if robots are to interrupt humans in a socially acceptable manner.

This view leads us to our main research question: are there *minimal non-verbal behavioural cues* that robots can exhibit to communicate interruption urgency, and are those cues understandable by people?

II. RELATED WORK

Interruption and its effects have been explored extensively in human-computer interaction (HCI) as well as other technical domains. For example, Horvitz et al. [4] found that decreased performance resulted from inappropriate interruptions to complex tasks due to higher demands on cognitive capacity. Chapanis and Overbey [2] found that while interruptions changed the way that participants chose to accomplish a task, the actual performance time generally was not affected. Storch [9] explored whether the style of computer user interface used by a person affects their performance following an interruption, uncovering lessons in human-computer interaction that we believe can be generalized to HRI. In particular, Storch showed that on-screen interruption messages were very disruptive; in-person visitors were somewhat disruptive, and telephone calls were not disruptive at all.

Yet very few studies have touched on the nonverbal interruption aspects of human-robot interaction. Of those

that do, most consider interruption tangentially and not explicitly. Bethel and Murphy [1] explored nonverbal and non-facial expressions in appearance-constrained robots, including those that might be used to approach people in search and rescue operations. Satake et al. [8] and Hayashi et al. [3] considered how robots approach people in train stations and shopping malls, respectively. Yamaoka et al. [10] described a model for a robot to appropriately control its proximity to the person when presenting information. A study by Mutlu and Forlizzi [5] revealed the possible problems that can be caused in a medical working environment by a robot that is not designed with interruptibility in mind.

While these examples do not explicitly explore interruption, they do suggest that the robot's effectiveness at coordinating with people could be better if they were able to use nonverbal robotic interruption cues more effectively.

III. METHODOLOGY

Our methodology focuses on designing and evaluating minimal robot behaviours for social interruption. To do so, we created a three-step methodological process.

1. Identify, through observations, human interruption behaviours within a particular situation that could be mimicked by a robot.
2. Based on these observed behaviours, design and critique potential robotic non-verbal behavioural cues, where those behaviours are based upon a minimal amount of the robot's physical capabilities (see below).
3. Implement these interruption behaviours using an actual robotic interface and use it to evaluate people's ability to interpret the robotic behaviours.

We have already stated that we are exploring interruption conveyed by *minimal non-verbal behavioural cues*. By this we mean that we are interested in determining interruption behaviours that rely on only a few simple physical capabilities present (or that can be easily added to) most robots, ignoring any verbal interaction layer.

Our methodology was limited to the following five non-verbal interruption parameters, which we deemed would provide a solid starting point to our explorations: (a) speed of motion, (b) gaze, (c) head movement, (d) rotation, and (e) proximity to the person. By looking at minimal non-verbal behavioural cues, we are trying to determine the lower bound of robotic interruption behaviours that could be understood by people.

IV. HUMAN-HUMAN OBSERVATIONAL STUDY

Our first question was: can we identify human interruption behaviours within a particular situation that could be mimicked by a robot? We focussed on identifying ways of interrupting a person based on different levels of urgency and importance, where the person is able to correctly interpret those interruptions. We argue that answering this question can be valuable when designing robot interruption behaviour that is minimally disruptive to people, while still

adequate to convey urgency and importance appropriately within a particular context.

Our approach - which we believe can be generalized to other situations - was to constrain human behaviour to particular physical capabilities that we believed could be reasonably mimicked by a robot, and to then have a person, an informed actor, try to interrupt others using only those capabilities. While these constraints disallow some natural human behaviours (e.g., speech), our goal was to capture how people manifest these particular behaviours, use these to model robot behaviour, and then see how these robotic behaviours were interpreted by people.

Specifically, three people were recruited as "robot actors" and asked to act through five interruption scenarios. (e.g., Fig. 1) These scenarios ranged from time-insensitive non-urgent matters, to very time-sensitive emergency situations. The emergency scenario was deliberately chosen to illicit extreme behaviour from the "robot actors" in order to elicit the widest range of possible behaviours.

Within each scenario, the actor had to interrupt two people who were engaged in a meeting inside an office with an open door. The actor was asked to improvise interruption behaviour appropriate to the urgency of the situation. The robot actors were constrained to show only the five previously mentioned parameters that our target robot could replicate, i.e., speed of motion, gaze, head movement, rotation and proximity to the people being interrupted. Robot actors were not allowed to speak or to make sounds. Furthermore, we instructed them to leave if no acknowledgement of their actions was provided after 10 to 15 seconds. An element kept secret from the robot actors was that the people they were interrupting were instructed to ignore the robot actors for at least 10 seconds, to allow the experimenters to have enough time to observe the robot actor's behaviour.

We videotaped the robot actors' actions, and identified characteristic behavioural trends (see Video Figure for three examples). We saw that our robot actors improvised with a range of body and head movement to match the given scenarios.

In the less urgent scenarios, actors used non-disruptive behaviour. In these cases, the robot actor would 'peak into' the office from a distance to see if the people inside the office were busy, checking whether it would be possible to interrupt without disrupting a more important task. The actors' behaviours allowed the people seated in the office to notice them, while at the same time providing what seemed to be a socially appropriate possibility to ignore him or her if desired.

As the task urgency and importance increased, our robot actors used more disruptive behaviour, maintaining close interpersonal proximity, and not leaving until their interruption was acknowledged and addressed. One robot actor entered the office running and stood between the two people inside the office until acknowledged. In these cases, the behaviour was certainly more noticeable to the people



Fig. 1. An actor looks inside the office from a distance to interrupt unobtrusively.

inside the office, who were often unable to continue their conversation because of the interruption.

The range of behaviour used by our actors shows that people can improvise their behaviour using our five basic characteristics to interrupt in different yet understandable ways. This suggests [6] that a robot designed to use similar interruption behaviour based on the limited set of physical non-verbal cues could also be understood by people.

V. USER STUDY

Based on the findings of our human-human observational study, we designed a user study to test the degree to which particular *minimal non-verbal behavioural cues* used by robots to communicate interruption urgency (Table 1) are understandable by people.

A. Equipment

Our equipment comprised four major components: (a) the robot platform (b) robot control, (c) a controller workstation and (d) controller software.

The robot platform we used was a Mobile Robots Inc Pioneer 3-DX base with a custom body added on top (shown in Fig. 2). The base was capable of moving faster than human walking speed and can carry loads up to 22 kg. Our custom body consists of a plastic container used to increase the height of the robot, covered by a t-shirt to reduce the robot's mechanical appearance without going as far as completely anthropomorphizing it.

The robot's 'head', used to portray head movement and gaze, is just a small cardboard box; depending on the scenario, we rotated this head left/right and up/down. The head does not include any facial markings such as eyes, though it does convey a clear directional 'forward' position, i.e., gaze. We used this minimalist design because, as mentioned earlier, we wanted to rely on only a generic shape and a few simple physical capabilities of movement that are present in (or that can be easily added to) most robots.

Robot control was done through direct serial connection to an on-board hidden laptop, where the laptop runs custom C#/C++ software. All sensory monitoring (including obstacle avoidance) occurs on the laptop. Commands to control the robot were sent from the study administrator via

a controller station. The robot also sent back timestamps and descriptions of high level events that it was performing back to the controller station.

The controller station, which also ran custom software, served two purposes: (a) controlling the robot's behaviour using a Wizard of Oz methodology, and (b) recording and logging relevant participant comments. It comprised a standard laptop with a second monitor, and a wireless router that linked the controller station with the robot. The controller station was positioned so that the study administrator could hear and record participant comments, and to keep the robot in view except when it entered the office. During the study, participants were not aware of the study administrator's actions as they could not see the controller station from within the office.

Controller software was primarily used by the study administrator to issue high level commands to the robot that triggered predefined macros, each executing a particular robot behaviour sequence. Normally, the robot macros automatically returned the robot to its starting position after the completion of a behavioural episode. If needed, the administrator could also use manual positioning controls to return the robot to its starting location. The software also supported logging and time-stamping high level events that were continuously communicated back to the station by the robot.

B. Participants

Twenty participants were recruited through mailing lists at the University of Calgary. Although no particular groups were targeted, the vast majority of participants were a nearly equal mix of male and female graduate students with varied ethnic backgrounds, many of whom were members of the Faculty of Engineering, with ages ranging from 20 to 30. Participants received \$15 in compensation for a study session which was approximately 45 – 60 minutes long.

C. Designing Robot Interruption Behaviours

We used our observations of interruption behaviour expressed by the robot actors to design and program the robotic interruptions. In particular, we created ten robotic interruption behavioural episodes using combinations of the parameters observed in the behaviour of the human actors. These are summarized in Table 1 (the gaze variable captures both gaze and head movement). We also included an 11th 'null base case' where the robot would use behaviour that had nothing to do with interruptions: slow movement outside the office without any direct interaction with the participant.

The ten episodes are labelled 1A-F and 2A-D. They were designed so that any one episode had another matching episode that differed only by a single difference in one of the behavioural properties. For example, episodes 1A and 1B are both based on non-urgent robotic interruption via observing the participant from outside the doorway, but they differ in the speed of motion used. This approach thus enables any differences in the participant's interpretations between two episodes to be feasibly attributed to the single



Fig. 2. The robot interrupts a participant inside the office

variable that changed its value.

These episodes were arranged to be presented in a scrambled order to the participant, following the null base case. The design of the episodes made it difficult to objectively argue for a precise ordering of interruption magnitude. However, we expected that episodes with fast speed, close, direct gaze, and rotation (Table 1, bottom right) would be higher in interruption magnitude than those with opposite values (Table 1, top left).

TABLE 1. BEHAVIOR EPISODES BY FACTOR (SPEED, PROXIMITY, GAZE, ROTATION) USED IN THE USER STUDY.

proximity position	rotation	slow speed		fast speed	
		direct gaze	erratic gaze	direct gaze	erratic gaze
far from doorway	none	1A		1B	
	rotating				
at the doorway	none	1C		1D	
	rotating				
next to participant	none	1E		1F	
	rotating	2A	2C	2B	2D

To ensure that all participants would observe nearly identical robotic behaviour, we created fully automated pre-programmed behaviour macros corresponding to each behavioural scenario that ran autonomously.

D. Experimental Procedures

The study consists of two phases. Both were qualitative, while the second also included a quantitative component. In both phases, the participant and the interviewer were seated in an office (e.g., Fig 2), having a conversation about topics unrelated to robots or the user study. The participant had a clear view of the doorway to his or her right. While the conversation was occurring, the robot underwent attempts to interrupt the participant across a series of behavioural episodes described in Table 1. All episodes began with the robot out of view outside the office. The two phases differed in the particular episodes used (and thus in the robot behaviour exhibited), and also in whether a verbal interview or ranking by the participants occurred once the robot had completed its behaviour.

The primary purpose of Phase 1 was to gather qualitative and unbiased reactions to interruption. Phase 1 comprised four pre-programmed behavioural episodes initiated by the study administrator that exhibited a wide range of robot



Fig. 3. A partial view of the Interruptedness Metre used by participants to rank interruptedness in the study’s second phase.

behaviour. The order of episodes used following the null base case (i.e., Null, 2D, 1A, and 2A) was randomly generated and the same across different participants. To reduce predictability of when an interruption would occur, each episode was separated by a short delay of a few minutes in which the chat between the interviewer and the participant continued. During each attempted interruption, the interviewer encouraged participants to talk about their reaction following a constructive interaction methodology (a variant of the think-aloud methodology). The participant reflections and statements were all recorded in real time by the study administrator. After the episode ended, the interviewer asked further questions about the interruption.

The primary purpose of Phase 2 was to gather additional reaction to the robot’s behaviour, and to quantitatively rank the level of ‘interruptedness’ of each episode. In this phase, the robot progressed through all eleven pre-programmed episodes, with little delay between them (i.e., Null, 1F, 2C, 2B, 2D, 1C, 2A, 1D, 1E, 1B, 1A). Following each episode, the participant was asked to rank how interrupted he or she felt by the robot’s behaviour by placing a marker corresponding to the robot’s behaviour scenario on our ‘Interruptedness Metre’ (Fig. 3). The participant would order the sequences they saw from least interruptive (left) to most interruptive (right). Rankings, which were translated from their relative position on the metre to a continuous scale from 0 to 100, gave a subjective but quantitative measure of how each participant interpreted the robot’s behaviours.

After the two phases, participants were interviewed and ask for their final impressions and thoughts.

VI. QUANTITATIVE RESULTS

The quantitative data comprised eleven rankings (one for each behavioural episode) collected from each participant using the Interruptedness Metre during Phase 2. The rankings form the participant’s subjective measure of how they interpreted the observed robot’s behaviour. Rankings were collected from all twenty participants, although one participant’s results were discarded due to corruption of their data. Based on the nature of the collected ranking data, a *linear mixed statistical model* was used for our statistical analysis. The model was configured to use the null base case data as the covariate, ensuring that all data analysis took that into account as the baseline.

A. Identifying Significant Behavioural Factors

The robot’s behavioural episodes (Table 1) were designed to enable statistical analysis that identified which of the robot’s behavioural factors (e.g., speed of motion, head movement,

etc.) had a statistically significant impact on the interruptedness felt by a person due to the robot's behaviour.

Tables 2 through 6 summarize the statistical significance of each individual factor of the robot's behaviour used in the study as well as the interaction between factors. P-values are considered statistically significant based on a threshold of $p < 0.05$. Statistically significant findings are distinguished using bold text in the tables.

Table 2 summarizes the effect of speed of motion, interacting with gaze and proximity. Robotic speed and its head movement were both classified as either slow or fast. As Table 2 indicates, speed of motion was significant only when the robot was situated next to the participant. When the robot was located at the doorway of the office or outside the doorway, no significant impact was observed.

TABLE 2: SIGNIFICANCE OF SPEED OF MOTION

For Gaze	at participant					erratic
Proximity	far from doorway	at doorway	next to participant			
Episodes	1A & 1B	1C & 1D	1E & 1F	2A & 2B	2C & 2D	
P-Value	0.139	0.360	0.025	0.010	0.006	

Table 3 summarizes the effect of gaze interacting with speed. The gaze suggested by the robot's head movement was either directly focused on the participant, or erratic movement where the head was constantly moving in all directions. The data in the table indicates that gaze had no statistically significant impact.

TABLE 3: STATISTICAL SIGNIFICANCE OF GAZE

For Speed	at slow speed	at fast speed
Episodes	2A & 2C	2B & 2D
P-Value	1.000	0.996

Table 4 summarizes the effect of proximity, interacting with speed. The robot used three proximity positions: next to the participant, at the doorway, and outside the doorway. The data shows that there was no significant statistical difference between being at the doorway or far from the doorway. However, there was a significant difference between being far from the doorway and being next to the participant. When comparing positions at the doorway and next to the participant, there was only a significant difference at fast speed, but not at slow speed.

TABLE 4: STATISTICAL SIGNIFICANCE OF PROXIMITY

Speed/Proximity	at slow speed	at fast speed
Far from Doorway vs. At doorway	1.000	1.000
At Doorway vs. Next to Participant	0.050	p<0.001
Far from Doorway vs. Next to Participant	0.012	p<0.001

Table 5 summarizes the effect of rotation, which interacts with speed. For some behavioral episodes, the robot rotated its body in place while stopped, while it used no body movement when stopped in other episodes. The data shows that this factor was not statistically significant.

TABLE 5: STATISTICAL SIGNIFICANCE OF ROTATION

Speed	at slow speed	at fast speed
Episodes	2A & 1E	1F & 2B
P-Value	1.00000	1.00000

B. Means

The means presented in Table 6 shed light on the magnitude and direction of differences for the robot behavioural factors that proved significant.

TABLE 6: INTERRUPTEDNESS MEANS BY FACTOR

proximity	rotation	at slow speed		at fast speed	
		direct gaze	erratic gaze	direct gaze	erratic gaze
far from doorway	none	19.3		34.6	
	rotating				
at the doorway	none	23.8		33.4	
	rotating				
next to participant	none	43.8		64.0	
	rotating	46.2	51.2	63.7	72.2

For speed of motion, the differences in interruptedness between slow and fast when the robot is next to the participant are both statistically significant and large: around 20% each (see Table 6, bottom row is close proximity, 46.2 slow vs. 63.7 fast, and 51.2 slow vs. 72.2 fast). For proximity, the significant differences in interruptedness are around 25-30% when comparing positions of far from the doorway to next to the participant (see Table 6, 19.3 far from doorway vs. 43.8 next to participant, and 34.6 far from doorway vs. 64.0 next to participant).

As mentioned, our statistical analysis indicates that the differences for gaze and rotation are not significant.

VII. QUALITATIVE OBSERVATIONS BY SCENARIO

The qualitative data consisted of verbal comments made by the participants during both phases, and during the verbal interview after phase 2. Comments were captured in video recordings, as well as in notes taken by the study administrator in real time. The notes were used to assist processing the full video recordings.

We now summarize and discuss the qualitative comments received for both phases of the study. We begin with presenting the participants impressions of each Phase 1 episode. Following that, we talk more generally about particular perception trends participants had across both phases

A. *The Null Base Case: Impressions of the Robot*

The null base case occurred first, after the participant was in the office for a few minutes, i.e., the robot passed by the office door without any head movement, and did not gaze into the office. During this episode, about half of the participants commented on the robot's behaviour just as it began moving past the doorway; the others just kept talking to the interviewer, and mentioned the robot only when asked by the interviewer when the episode ended. Most said they first detected the robot because of its noise, even before it was visible through the doorway. Many described the details of how they observed the robot's behaviour using phrases like "it just passed by", "it's coming", and "it disappeared". The behaviour was "calm", and "not disturbing". One participant said the robot looked as if "it could move faster" than it was. Although all participants clearly noticed the robot, one said it was not "super distracting" and that it got his attention in a "polite way". Another said the robot was "minding its own business" and that it "didn't affect the flow of conversation". One found the whole behaviour to be "pretty weird".

Participants were asked what they thought the robot was trying to do, or what its intentions were. None felt that the robot was trying to interact directly with them, but opinions of what it was doing during the null base case varied. Some assigned social presence to the robot similar to a co-worker just passing by on the way to some other location, or pacing about with no specific mission, or in the middle of accomplishing a task such as delivering messages or moving objects around the office. One even compared the robot to a "child waiting to be noticed".

Participants were quite generous in the social abilities they afforded to the robot, despite it lacking any form of eyes, ears, cameras, microphones, or speakers. Many participants felt the robot was "curious" about their presence, even "spying" or "eavesdropping" on the conversation, because the robot did not "know" who the participant was. One even implied that the robot felt territorial because it was approaching for a "sense of security". Others felt that it was responding to them and the interviewer being loud in their conversation. Some were more specific, saying that the robot "heard its name" (despite the robot not having a name during the study) or the word "robot", and wanted to hear more of the conversation. Many noticed the robot's lack of active behaviour in the null base case (other than moving by the doorway), and said the robot was not interacting because they were "not paying attention".

B. *Fast, erratic gaze, close proximity, rotating (2D)*

In this episode, the robot directly approached participants with its most extreme behaviour, where it was active and fast-paced (Table 1, episode 2D).

Participants initially described this behaviour using active words such as "weird", "big", "racing", "scared", "frantic", "hard to ignore", "in a rush", etc. Many participants said

they were "annoyed", "distracted", "disturbed" and "interrupted" by the robot's behaviour, and unable to continue their conversation with the interviewer. The entrance of the robot into the room was described as "forceful" or comparable to banging on a door. Because of the robot's faster movement, its motors made more noise, which one participant described as "different" and "huge".

The behavioural factor mentioned most often was speed. Head movement was also mentioned, but to a much lesser extent. Many participants also noted that the robot came into the room (referring to the factor of close proximity) vs. the previously discussed scenario where the robot just passed by the doorway. Very few commented specifically on the robot's body movement while inside the room during the whole study, even though it was persistently rotating back and forth for 15 seconds. One specifically said that the closeness of the robot felt more significant than its movement. Another said it was "kind of weird" that the robot was communicating with body language only, and no verbal communication.

Almost all participants viewed this behaviour as representative of an "emergency", "something [being] wrong", "someone hurt", or something having "happened". Several participants even identified the emergency as a possible "fire", one saying "probably a fire". One said the robot's behaviour indicated that it was necessary to stop the conversation and move out of the room.

Almost all used words such as "important" or "urgent" to describe the potential reasoning behind the interruption. One said this behaviour would be "rude" if it was used to interrupt an important meeting, though perhaps not as rude if the meeting was casual.

In summary, it is clear that this behaviour was largely associated with "fire" or "emergency", which matched our predictions. Indeed, one participant said the behaviour would be "inappropriate" for a non-urgent interruption such as a greeting.

C. *Slow, direct gaze, far from doorway (1A)*

In this Phase 1 episode, the robot stood outside the doorway and did not enter the room (Table 1, episode 1A). Generally, this behaviour was seen as non-interruptive. In all but one case (where the robot was not even noticed), participants noticed the robot in part due to the noise it was making. Comments described how non-disruptive it seemed, for example it was "not interrupting" because it "did not approach too close, but from a distance".

Many participants felt that the robot was acknowledging their presence and "noticing" them, e.g., "this time I'm sure it's noticing us", because of the "head movement". A few said the robot was going by, but was stopping to "listen to the conversation", and that it was "paying attention". Another said it was "curious" and that it was "eavesdropping a bit" because it "overheard the conversation and was interested".

Other participants interpreted robot behaviour as something other than interruption-based. One said it was

doing “periodic checking, in case we need something”. Another said the robot was “peeking inside the room” and then “reporting back to someone else”.

D. *Slow, direct gaze, close proximity, rotating (2A)*

In this episode, the robot moved slowly, with direct gaze at close proximity (Table 1, episode 2A). Participants had varied impressions. One participant noted that the robot, like a person, was more interruptive when it entered the room, compared to when it did not enter. Another said the robot seemed to be acting with more “maturity” due to the eye contact, and that it was respectful and more “accustomed to social rules”. One said he was “surprised by the smooth motion”, and that it was “not going crazy”.

Many participants surprisingly expressed how they felt emotionally about this interaction, often contrasting it to the feeling they had regarding the previous ‘urgent’ behavioural episode (2D). Two participants said that this behaviour “didn’t scare” them. One said that the previous one had lots of “shaking” and required some “getting used to”. Another said the robot was “not very annoying” whereas it was previously “making a lot of noise” and “bothering” him. One participant felt more comfortable, whereas they had been previously worried that the robot might have hurt them (in episode 2D). One participant “preferred” this behaviour.

VIII. OVERALL QUALITATIVE OBSERVATIONS

We now turn to more general impressions of the robot, across all episodes in both phases.

A. *The Robot as a Social Being*

Many participants made comments about the robot as if it were a person. One participant said the robot was like a “real being” because it was showing interest in things, going away and then coming back. Another said it moved and tried to gain attention by “barging” in and moving its head. One felt that the robot was “annoyed” that its space was being intruded on. Another suggested that the robot was actually trying to annoy him or do something funny.

One participant compared the robot to a dog running up to a visitor when entering a house. Another compared the robot to a child entering the room, in a manner that a child might approach his or her parent, to say that someone was annoying them.

B. *The Robot as a Machine*

A few participants described the robot as a machine. One said its procedure was “smooth”, because of the “mechanics or software”. Another suggested that the robot was exhibiting certain behaviour because it was “broken” or “damaged”. Yet another felt that the robot was “examining the perimeter, becoming familiar with its surroundings, and mapping out objects”. One said the robot seemed to be “analyzing” them, collecting data, taking pictures, and recording audio.

A couple of participants suggested the robot was running through “programmed” behaviour or being controlled by the

study administrator during the study.

C. *Politeness when Interrupting*

A common theme used in describing the robot’s behaviour in some episodes was politeness. Many participants felt the robot had some intention, but that it had chosen to defer that intention when it noticed that a conversation was in progress. One participant thought the robot wanted to “say something” that was “not important”, but that it “changed its mind” because of the conversation, and that it would “come back later”. Another felt the robot was coming for a “scheduled meeting”, but that it was “waiting outside”. Many participants defined this behaviour as either “normal”, “better” and “more gentle” than episodes where the robot had used much more extreme behaviour (e.g., episode 2D). Another said the robot was “looking for someone” on behalf of someone else, and that it was “trying to say something”, but did not say anything because it “didn’t want to interrupt” the conversation. Similarly, another participant said the robot was “trying to look for an opening in the conversation” so that it could “add to it”.

A couple of participants interpreted authority as a factor in some episodes, and compared the behaviour to someone who “is waiting for a superior to finish” and that the matter was not “urgent”, as the robot was not “actively catching attention”. Another participant saw the robot as a “messenger” or “servant”. Another participant felt the robot was acting as a servant, but for someone else (that is, not for the participant or the interviewer).

While most participants used comments that implied some element of politeness, one participant said the robot was “impolite” because it was just “staying there and staring”, though this person later noted that the robot “didn’t want to interrupt”.

IX. DISCUSSION

Based on the results of the user study, we now have several insights about robot interruption behaviour.

First, we have demonstrated that robots can convey urgency about an interruption situation using only basic elements of its physical behaviour. Our quantitative results statistically show that speed of motion, and proximity to the person can both express a range of interruption levels. While this is a simple result, we believe it does provide an important insight demonstrating that basic physical behaviours can be used by practically any robot (e.g., a Roomba) to convey interruption context. Our findings indicate that interruption is feasible even for robotic interfaces that lack gaze or precise body rotation movement, as we found (a) these factors are statistically insignificant in our study, and (b) the motion and proximity factors by themselves suffice to convey interruption.

While other factors did not exhibit statistical significance, anecdotal comments from participants do suggest that some form of eyes, head, or indication of forward direction can benefit the interpretation of interruption (our study likely did

not have enough power to reveal this statistically). This too can be easily added to simple robots, e.g., by ‘painting eyes’ onto its front. When the robot was distant or not gazing directly, participants often did not feel that they were necessarily the ones that the robot wanted to interrupt. Instead, they seemed confused when the robot was close, moving frantically but not making eye contact. Participants also felt that the robot was searching for someone else when far away. In these cases, it is helpful to equip the robot with some method of identifying who it wishes to address in an interruption scenario. In summary, while gaze is not statistically significant, it appears to play a role in disambiguating who the robot is addressing.

Second, participants were almost eager to think of the robot as a social entity, although nothing in our study explicitly encourages this. Many participants saw the robot as more than just a machine and referred to it as a social being with its own desires, goals, and thought process. We believe these results illustrate the ‘Media Equation’ effect uncovered by Nass and Reeves [6], where people were found to be inclined to treat media items (such as computers) as living things.

Third, the interpretation of politeness in the robot’s behaviour provides confirmation that a robot can communicate interruption urgency in a way that minimizes the perception of disruption by people. Clearly, there are cases, emergencies for example, where being ‘polite’ may not be important, so long as the person understands the message. However, minimizing disruption could be very important in many common cases where a robot is attempting to interrupt a busy person for a non-urgent matter without running the risk of annoying that person.

Finally, our results reveal that people may have preconceived notions of certain robotic behaviour that is inappropriate in all but the most urgent scenario. For example, frantic and fast movement close to the participant was mostly seen as just plain annoying by many participants, or at the very least associated with a fire or other extreme emergencies, thus making it inappropriate for any other less critical scenarios.

X. CONCLUSION AND FUTURE WORK

Following from our research question, our study results have shown that it is possible to use only a *minimal* set of nonverbal behaviour cues to convey different levels of interruption. Our findings demonstrate that an interrupting robot can work simply by using motion and proximity. The robot may not need to be equipped with features such as verbal communication, as our robot interrupted without it, or even gaze and head movement. We are not suggesting these other behaviours are unnecessary; rather, they could be used to add other kinds of interruption information such as who is being addressed (by gaze) and the message of the interruption (by speech). We believe our findings contribute to informing the design of various minimalistic robotic interfaces where robots are perhaps incapable of much

beyond locomotion, but still need to convey interruption in a socially appropriate manner. Furthermore, we do feel that even robots that are designed with additional richer capabilities can use our findings to provide further context about an interruption. For example, motion and proximity can set the stage for a person to acknowledge the interruption, after which the robot (if equipped) could use speech to convey the message content. This is normally what happens between people.

Of course, much is left to do. We need to examine how different interruption behaviours translate into working environments beyond the limited office setting we used, e.g., public places, high stress working environments, and various domestic settings. We need to explore differences in interpretation of interruption behaviour across different cultures with varying societal norms and familiarity with robots. Exploration could move beyond interpretation and into ways to design robots that are capable of learning appropriate interruption behaviour based on past experiences and their own observations of human-human interactions.

REFERENCES

- [1] C.L. Bethel and R.R. Murphy, "Survey of Non-facial/Non-verbal Affective Expressions for Appearance-Constrained Robots," *Systems, Man, and Cybernetics, Part C: Applications and Reviews*, IEEE Transactions on , vol.38, no.1, pp.83-92, Jan. 2008
- [2] A. Chapanis and C. Overbey. *Studies in Interactive Communication: III. Effects of Similar and Dissimilar Communication Channels and Two Interchange Options on Team Problem Solving. Perceptual & Motor Skills*, 38(2):343–374,1974.
- [3] K. Hayashi, D. Sakamoto, T. Kanda, M. Shiomi, S. Koizumi, H. Ishiguro, T. Ogasawara, and N. Hagita. Humanoid robots as a passive-social medium: a field experiment at a train station. *Proc. Human-robot interaction (HRI'07)*, 137–144, 2007. ACM.
- [4] E. Horvitz, J. Apacible, and M. Subramani. Balancing awareness and interruption: Investigation of notification deferral policies. *UM*, 2005:433–437, 2005.
- [5] B. Mutlu and J. Forlizzi. Robots in organizations: the role of workflow, social, and environmental factors in human-robot interaction. *Proc. Human-robot interaction (HRI'08)*, 287–294, 2008. ACM.
- [6] C. Nass and B. Reeves. *Media Equation: How people treat computers, television, and new media like real people and places*. CSLI Publications, Center for the Study of Language and Information Leland Stanford Junior University, Cambridge, UK, first paperback edition, 1996P. Riley. Discursive and communicative functions of non-verbal communication. *ED143217/FL008785, ERIC (Educational Resources Information Center)*, Sewell, NJ, 1976.
- [7] P. Riley. Discursive and communicative functions of non-verbal communication. *ED143217/FL008785, ERIC (Educational Resources Information Center)*, Sewell, NJ, 1976.
- [8] S. Satake, T. Kanda, D. Glas, M. Imai, H. Ishiguro, and N. Hagita. How to approach humans? strategies for social robots to initiate interaction. *Proc. Human-robot interaction (HRI'09)*, 109–116, 2009. ACM.
- [9] N. Storch. Does the user interface make interruptions disruptive? A study of interface style and form of interruption. In *Posters and short talks, Proc. ACM CHI*, p14, 1992, ACM.
- [10] F. Yamaoka, T. Kanda, H. Ishiguro, and N. Hagita. How close? model of proximity control for information-presenting robots. In *Proc. Human-robot interaction (HRI'08)*, 137–144, ACM.