1-1-2011

How Can Robots Adapt To A Social Human World? A Study Into The Role Gestures Can Play In Human-Robot Relations

Bradley Matherne

WASHINGTON UNIVERSITY IN ST. LOUIS
School of Engineering and Applied Science
Department of Computer Science and Engineering

Thesis Examination Committee:
William D. Smart, Chair
Ron K. Cytron
Christopher Gill

HOW CAN ROBOTS ADAPT TO A SOCIAL HUMAN WORLD? A STUDY INTO THE ROLE GESTURES CAN PLAY IN HUMAN-ROBOT RELATIONS

by
Bradley M. Matherne

A thesis presented to the School of Engineering of Washington University in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

August 2011
Saint Louis, Missouri

ABSTRACT OF THE THESIS

How Can Robots Adapt to a Social Human World?
A Study into the Role Gestures Can Play in Human-Robot Relations
by
Bradley M. Matherne
Master of Science in Computer Science
Washington University in St. Louis, 2011
Research Advisor: Professor William D. Smart

There is no doubt that robots are now starting to increasingly be integrated into mainstream society, and as they do the amount of contact and the number of interactions they have with humans will increase at a similar rate. These interactions open a new set of issues for robot designers and programmers. What is the best way for robots to interact with humans? This study tested the importance of gestures in creating effective human-robot interactions. Conducted using the PR2, this study explored the role of gestures in two basic kinds of communications: the robot communicating a need (low power) to an unsuspecting human, and a robot building trust with a human participant on an instruction reading task. We predicted that gestures would lead to more effective interactions than the non-gesture controls. We also used the opportunity to explore a largely unresearched area in proxemics: the idea that loose, "bouncing" arms led to lower attributions of dominance than stiff, fixed arms. Our research highlighted the importance of gestures in communication, particularly among people who tended to look at robots as more than machines.

Acknowledgements

First and foremost this research would not have been possible without the support and resources of Willow Garage. I would especially like to acknowledge Leila Takayama, Irene Rae, and Bianca Soto for their assistance in resolving issues pertaining to experimental design and assistance with recruiting and experiment execution. One of the experiments in this study was inspired in part by work done with former classmate Eva Dooley, so I would like to thank her for her previous work with me.

I would like to recognize and extend gratitude to the members of my committee, Dr. Chris Gill, Dr. Ron Cytron, and Dr. Bill Smart for taking the time to review me for this degree.In particular I would like to recognize my degree advisor Dr. Ron Cytron for his ever reliable guidance throughout my time with the computer science department at Washington University. Above all else I would like to extend my most sincere thanks to Dr. William Smart without whom none of this would have happened, as he was the primary liaison between Willow Garage and me, and also managed the technical aspects of hardware and software preparation.

Bradley M. Matherne

Washington University in St. Louis

August 2011

Contents

List of Tables

List of Figures

Preface

Can robots ever hope to be human? Not in a literal sense, but in a behavioral sense; can a robot ever behave as if it were human? What would be the implications of such a machine that behaved like a human? The concept of robots that behave in human-like ways is actually not a foreign concept in the slightest. Throughout popular science fiction we can find examples of robots that are almost human. Whether it is WALL-E's adorable antics in trying to find his beloved EVE (Morris, Collins, Lasseter & Stanton 2008), C-3PO's futile pleads trying to coax his friends away from dangerous situations (Kurtz & Lucas 1977), or Marvin's humorous self-pitying rants (Adams, 1978), our culture has long embraced the idea of robots that were more than just metal parts and wires. The examples from above were all social personalities existing in and interacting with the social worlds around them.

Modern robots, however, can claim nowhere near the level of sophistication that any of these fictional robots can. The vast nuances of human social interactions are still largely mysteries to psychologists and sociologists alike, so it is no wonder that there has been limited work done in implementing human-like social behavior in robots. Much like the evolutionary pressures of living in social environments that may have led to the development of our own social behaviors, robots are now increasingly finding themselves in social situations, which many are ill-equipped to handle. A classic example would be Lewis the wedding photographer robot (Byers, Dixon, Smart & Grimm 2004). Lewis was a robot programmed to drive around a wedding and take pictures of people, much as a professional wedding photographer would. Although it may not seem like it at first glance, photography is a very social endeavor with a fairly nuanced social interaction occurring between the photographer and the subject. Lewis was unable to indicate to people when it was taking pictures, so subjects were left to guess if it was being a good photographer or merely malfunctioning, resulting in a fairly low appraisal of Lewis' abilities by those in attendance at the wedding.

Lewis' story is not unique among robots today. In fact, robots are becoming more mainstream every day, but they are bound to face numerous hardships and hurdles as they become available for commercial use that requires an increasing amount of human interaction. Obviously, robots like iRobot's Roomba (the robot vacuum cleaner) will continue to be successful regardless of the fact that they possess no social skills, but as robots take on more complex jobs, they will inevitably need to be able to interact with people in a social world. People need to be able to communicate to robots easily what they need the robot to do and feel confident in the fact that the robot will do it. Likewise, the robot needs to communicate its needs and goals with people. As a large percentage of what we communicate is nonverbal, simply adding a voice box to the robot with some crucial prerecorded phrases will not be enough for a natural interaction. The robot must be endowed with social gestures.

This is where our work comes in. This study consists of experiments to test if giving robots gestures both facilitates a robot in sharing its needs and goals with humans and helps instill confidence in humans that robots understand their instructions when they share them with the robot. We will also be taking advantage of the opportunity to explore a new factor that could affect proxemics distance between humans and robots: the robot's arms rigidly held in a fixed position or in a looser position, more able to move slightly with the robot's movement. We are hoping that the evidence gained here will help pave the way for the future generations of social robots.

1 Background

1.1 Gestures Facilitating Communication

One interesting observation that clearly distinguishes robots from other machines in the world is that people are generally more apt to treat these machines as people than they would a typical automobile, washing machine or even standard computer. Many are frequently treated like dogs or other small pets, and sometimes even as children. Take, for instance, the example of Sparky produced by Interval Research Corporation (Scheef et al. 2000), where the researchers created a very social teleoperated robot named Sparky. Sparky was endowed with gestures and mobility that allowed him to have fairly rich social interactions with the humans it encountered, and since it was teleoperated, it allowed a human to make the difficult interpretations of the social scenarios and choose the appropriate response. The results showed that when naïve humans (that is to say humans without prior experience with robots) interacted with Sparky, they overwhelmingly treated it as if it were a real living creature. The humans (especially children) displayed a wide variety of emotions and reactions to Sparky based on different behaviors it displayed. Some would tease Sparky and many would show great compassion treating it as they would a puppy. Behaviors like this do not manifest for interactions with more conventional machines, and it appears to be the nonverbal communication that brings about these behaviors.

It seems that, largely, robots get an automatic upgrade from machines when naïve people meet them. Perhaps this is the result of generations of science fiction movies showing us what we hope one day robots could be; at any rate it seems readily apparent that people tend to approach robots with very different expectations than they approach washing machines or computers. Often times the robots fail to meet the lofty expectations of the naïve human and the expectation drops quickly (as was seen with the wedding photographer robot), usually to the more appropriate expectation level of a glorified mannequin. However, if, as was the case with Sparky, the robot is able to engage the human in a social way, we hypothesize people will continue to place these machines on a pedestal above all others. The key lies in the machine's ability to be perceived as a social actor, rather than a functional object.

Another example of a robot working to break into the social realm is the artificial emotional pet robot created by Shibata, Yoshida, and Yamato (1997). This robot did not possess gestures as rich as Sparky- in fact most of its gesturing came from tail wagging, but it still demonstrates the same idea that the appropriate social gestures can elicit social responses from humans. In this case the robot displayed social gestures more commonly associated with dogs than people (unlike Sparky which attempted to use more human-like gestures), however, because dogs are so deeply ingrained in our lives and have been for thousands of years, canine gestures also prove to elicit social responses from humans. The same was true in the case of this robot, especially when humans were primed with canine images; they overwhelmingly responded to the robots gestures with behaviors typically displayed towards dogs. The humans would pet or stroke the robot- behaviors certainly not common in other human-machine relationships. Some humans even went as far as appraising the robot's intelligence as high based merely on these interactions, suggesting that perhaps nonverbal behavior is the key to perceptions of intelligence.

The idea that nonverbal behavior can lead to perceptions of intelligence is a bit murky, largely because the concept of intelligence itself is a bit murky, left largely to the interpretation of the person ascribing the intelligence. Intelligence is a broad, vague term meaning many different things to many different people. Certainly proper use of social cues does not suggest mathematical prowess (though perhaps merely being a robot does); what it does seem to suggest is a certain level of social intelligence, which in and of itself is a vague term, but it is clearly narrower and easier for us to work with. Social intelligence is known by many names including social competence, emotional intelligence, and social skill to name a few. For the purposes of this paper we will use the word social competence to help maintain a better distinction from the messier notion of intelligence. The relationship between nonverbal behavior and social competence is well documented across many studies for human-human interactions. Feldman, Phillipot, and Custrini (1991) summarized the results of several of these studies noting that there was a clear positive correlation between social competency and use of nonverbal behavior skills. In other words, people who were evaluated to have high social competence also displayed a high level of nonverbal behavior skills. The researchers did not go as far to as to suggest a direction of the correlation or suggest a causation between the variables. This evidence reflects the reports by the participants of the Shibata et al. (1997) study that the robot displaying social behavior was "intelligent." Perhaps "intelligent" was not the most accurate word to describe it; as we have discussed previously social competence seems to be more applicable for this situation. However, unlike Feldman et al., Shibata et al. can comment on the relationship in a causal sense. Since the robot has no intrinsic social competence, it is impossible to say that the social competence caused the nonverbal behavior. Clearly in this case the nonverbal behavior caused the participants to ascribe them with social competency, a key point this study aims to confirm as well.

1.2 Nonverbal Behavior and Dominance

As it turns out, social competence is not the only skill that can be ascribed to objects given the object's nonverbal behavior. Consider cartoons and 3D animation films. Similar to robots, cartoons and other animated characters only loosely represent humans or other autonomous agents in physical appearance, yet when done well, very strong feelings and thoughts can form around these characters. The reason for this is because these characters speak to us on a nonverbal level. As Johnston and Thomas (1995) noted in their landmark book, *Illusions of Life: Disney Animation*, animators seek to create deep, rich personalities for their characters without explicitly verbalizing those personalities. That means their personalities must be communicated through their dress, their mannerisms, and their gestures. The idea of conveying dominance and power, for instance, always seems to be associated with a quiet stillness or stiffness. Johnston and Thomas used the example that it would be a major break of character if the wicked witch, while slowly and deliberately walking down the dungeon stairs, stumbles and goes tumbling down to the bottom. Stumbling and falling portrays weakness; she would clearly not be in control of her environment in that scenario. There is no reason why these animation principles cannot be applied to robots to try to convey similar ideas in our studies and as robots continue to grow more social in the future.

So, it is easy to say that if we apply these animation techniques to our robots we can convey these deep social messages, and quite another to prove that these messages are indeed successfully conveyed. Merely asking participants to report their perceptions might be adequate for some standard of certainty, but we aim to have harder metrics not tied to subjective participant perceptions. In order to do this, we turn our attention to the field of proxemics, the study of personal space. In 1966, Edward T. Hall published his findings on proxemics in which he claimed that we all have personal bubbles that get smaller and smaller the more familiar we are with people. For instance, our most intimate of companions we permit to enter our most intimate circle, whereas strangers are kept at the furthest distance from us. If we take this notion a step further we can say that we permit people with whom we are more comfortable to be closer to us than those we are not. This does not give us a perfectly clean metric for determining whether messages have been successfully conveyed, but it will give us a better idea of what participants actually think. For instance, dominant personalities are often perceived as intimidating, so we are not comfortable around them and therefore try to keep them at greater distances from us. Using this idea we can test if one posture or stance is more submissive than another.

Similar studies have been conducted already testing a wide range of parameters that affect proxemics between human and robots. A study done by Takayama and Pantofaru (2009) tested to see what role eye contact and the locus of control (human approaches robot versus robot approaches human) played in proxemics distance. They did find that eye contact played a role in distance. Strangely, it only played a role with women however. Women tended to keep the robots at a further distance when the robots were looking at them than when they were not. Also strangely, they found the locus of control had no significant bearing on distance. They concluded noting that these differences needed to be investigated more. Another study by Van Oosterhaut and Visser (2008) found that like the findings of Takayama and Pantofaru, people largely adhere to similar standards and rules with robots as they do with humans. They also found that women displayed further distances than men. They did test for height disparity, but were unable to find any conclusive trends. The fact that these two studies both found larger distances with women as opposed to men is very interesting and something that this study will hope to explore further.

1.3 Gestures and Trust

It is not just enough for people to be comfortable around robots; all machines require a function (a task for which they were designed to complete) in order to be useful for users. Since we already know that people tend not to interact with robots as they interact with other machines, we would not expect task assignment for an autonomous robot to work the way task assignment works for a microwave or washing machine. No doubt the long term goal for robots is for them to process voice commands from users, which seems a reasonable goal given similar voice command technology is already deployed in several makes of automobiles. Voice recognition software is inherently unreliable and many times makes mistakes, leading to confusion. If we take a moment however to think about it we may quickly realize that our speech recognition errs frequently as well. We mishear or misinterpret what people say on a daily basis; therefore it is not surprising that our machines struggle with similar issues. We get around it using subtle cues executed in a pseudo-ritualistic fashion showing the instructor that the instructions were received and providing assurance they were understood. So the question then becomes how do we assure users that robots have understood tasks that are given to them? In other words we must explore ways that a rudimentary level of trust can be established between human and machine. What cues are necessary to create a pseudo trust establishment ritual between human and robot?

Lee and See (2004) defined trust as "the attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability." This is exactly what we hope to influence with our studies. Lee and See went on to argue that trust is built and evaluated in three ways: analytical, analogical, and affective. Analytic trust is built using rational principles and is built up over time. Someone builds trust with someone else if they consistently prove to be trustworthy. This is true and fairly obvious for our interactions, but certainly not all trust is built that way; we can all think of those people we've interacted with that instantly gain our trust. In analogical construction of trust, people use characteristics they observe of the person and the context to decide whether that person is worthy of their trust. They also largely rely on outside sources such as gossip and previous experiences. The affective evaluation of trust does not rely on reason to the extent that the previous two do. Affective relies most heavily on emotion; as trust is betrayed, negative emotions will be associated with the violator and when trust is not abused, it is rewarded with positive emotions. Thus the decision of trust is made by deciding whether they feel the person deserves trust based on the emotions they feel towards the person in question. For the purposes of this paper, we will need to focus on analogical construction of trust, as both affective and analytical involve developing a bit richer of a history between truster and trustee in order to be effective. We can however most easily manipulate context and robot characteristics in a controlled setting.

There is no doubt that trust will have to play an important role in any kind of interaction between human and machine. A study by Freedy, DeVisser, Weltman and Coeyman (2007) highlighted this fact in their study involving an autonomous targeting system for robotic military platforms. Participant trust in the targeting system was primarily affected by cues that suggest the system was not competent to handle targeting on its own. If the machine was displaying low competence, then participants would help out, but if it was displaying medium-level competence, people were more inclined to not interfere, because, the researchers theorized, the level of competence could not easily be determined. They also noted that the first impressions participants had with the autonomous system greatly affected their levels of trust in the system. For example if participants were originally tested with a high competence system, and then tested with a low competence system, their previous experience with a high competence system seemed to lead them to not interfere with the low competence system. This provides us with keen insight into critical moments when trust can be fostered or lost, but it does not unfortunately give us an idea of what kind of cues can be used for more domestic applications.

A panel of human-robot interaction experts (Bruemmer et al. 2004) shared their ideas on what they believed would be the best way to build trust in robots. Their views came from a wide range of disciplines, and not all the approaches were applicable to this research, but several provided some interesting ideas that proved to be valuable in designing the present study. Donald Norman believes the key to getting people to trust robots lies in the robot's ability to be "human;" to have emotion, personality, and rich interactions with people. This idea is similar to the ideas of the Disney animators Johnston and Thomas, when they were explaining how to create "life" through animation, and making interactions richer certainly seems an easy way to build trust. In fact another panelist, William Smart, argues that the key lies in giving social cues indicating to others what the robot's internal state is. Much of this theorizing came from the earlier work he did with the wedding photographer robot discussed earlier. While the ideas of Smart and Norman seem to fundamentally get to the core idea of making interactions richer, the ideas that Smart proposes at least in the short run are more feasible because making a robot "human" as Norman suggests is an incredibly lofty goal, that likely will not see fruition for several decades. The views presented by these two researchers are the most applicable to this research, because the primary variable we are trying to manipulate is the participant perception. Building off these ideas, we can use human-like gestures and cues to help foster a sense of trust.

This study tries to take all of these ideas discussed above and test each of them in an effort to determine the role that gestures can play in human-robot interaction. The first experiment aims to test to see if gestures can help communicate ideas about the robots internal states ("needs"). This is accomplished by having the robot attempt to communicate that it needs help with an unsuspecting participant. The robot will communicate in one of two ways either in a vocoded voice or using gestures. The second experiment will draw on the idea of dominance discussed in section 1.2, and rely on proxemics measures to interpret whether or not dominance has been conveyed. The robot will approach the participant with either rigid (unmoving) arms, or slightly bouncing arms; first the robot will approach the participant, then the robot will back up, and the participant will approach the robot. Finally we will assess what role gestures can play in establishing trust. The participant will read a set of instructions to the robot and the robot will either nod or do nothing, and the participants level of trust will be measured using a questionnaire.

2 Methods

2.1 Participants

Participants for this study consisted of 18 adult, 9 of which were female, and 9 of which were male. All participants resided in the Greater San Francisco Bay Area, and were notified using a variety methods including Craigslist and Facebook market advertisements, mailing lists forwarded to Stanford University students, personal email to friends and acquaintances of Willow Garage staff, and participants sharing information with their friends. All participants were screened to ensure that they were all at least 18 years old and were native English speakers. We also attempted to ensure that participants had minimal exposure and experience with robots. While we largely succeeded in this specific filtering method, a few individuals who possessed familiarity with robots did participate in the experiment. Their data was flagged, but did not seem to have significantly affected the results.

2.2 Materials

This study was conducted entirely using Willow Garage's primary robotic platform, the PR2 (Figure 1). As can be seen the robot has two arm-like actuators and cameras arranged on the head that give the robot an approximately human-like appearance, though clearly distinct enough so as can clearly be identified as a robot. In the event of a malfunction in the autonomous system, the experimenter always had a teleoperated

Figure 1: A PR2

controller nearby during all experiments to take control manually if necessary. A remote kill switch was also kept on hand in the unlikely event the robot malfunctioned in a way that could put someone at risk.

2.3 Procedure/Data

The experiment is broken up into four smaller experiments. For the purposes of this paper, we will only focus on the first three parts of the experiment, as the fourth part was designed and administered by researchers at Willow Garage. Sometimes the Willow Garage study was run before our study and sometimes it was run after, but again since it was a different robotic platform, we believe the pollution will be minimal. For descriptions of these platforms see appendix page (29). Before the study begins, each participant was given a brief tour of the facilities, and introduced to a few PR2 robots (to ensure that their measured reactions to the robot are not merely reflective of them just marveling at the technology). The complete experimental protocols are found on appendix page 30.

2.3.1 Experiment 1: Robot needs help

Figure 2: Robot Gesturing, (left) calibration imitation, (center) waving cord, (right) pointing at plug

The first experiment intends to ascertain if a robot that uses gestures can more effectively relay information regarding a robot's current "internal states" or, in lay terms, if gestures can help robots convey to humans what their "needs" are and see if the robot can successfully solicit help from the participants. This is accomplished by having the participant sit alone next to a PR2 for two minutes while the robot performs one of the two experimental behaviors, which will indicate to the participant that the battery is low and needs to be recharged. Some of the participants (n= 6) encountered a robot issuing an explicit verbal command ("Low battery, help me"), and the remainder (n=12) encountered a robot that gestured at the participant by holding out its power plug and pointing to the power outlet (See Figure 2). To start, the participant enters a room and sits at a desk near a robot seemingly engaged in a sensor calibration routine (See figures 3 and 4 for room layout and setup). Immediately after the participant signs the consent form, a camera starts rolling that will be used to track eye contact. The experimenter then apologizes for being disorganized and says he must run and go print another document before the study can proceed. Shortly after the experimenter leaves the room, the robot will stop the calibration imitation and begin one of the experimental conditions. After approximately 2 minutes, the experimenter will return with the "missing" document and take the participant to another room to complete the post experiment 1 questionnaire. A copy of the questionnaire is included on page 31 of the appendix. At times participants were reminded which

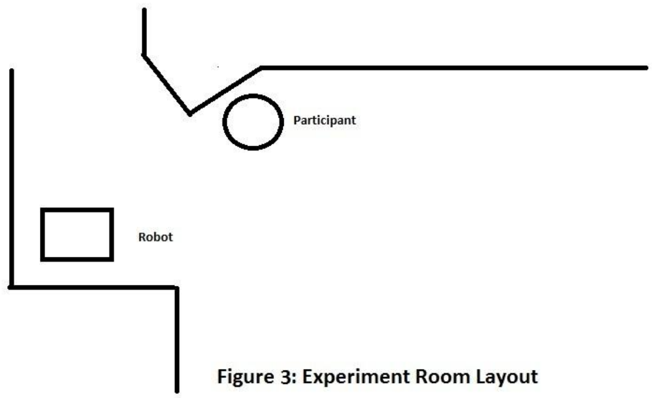

Figure 3: Experiment Room Layout

Figure 4: Experiment 1 room setup

robot was the one they had the interaction with. For this experiment it was important that the participant be left alone so as to minimize diffusion of responsibility and maximize the chances that the robot's actions were perceived as directed towards the participant. The participants' eye gaze at the robot in the presence of these socially engaging stimuli was measured by using a stop watch to time how many seconds the participant looked at the robot in the video of the interaction. We believe that a higher amount of eye gaze reflects an attempt by the participant to discern what the robot is trying to convey.

2.3.2 Experiment 2: Proxemics

The goal of the second experiment is to find out what a comfortable interaction distance is for humans and robots, and how different ways the robot appears affects those distances. There have already been extensive studies on the effects of height on comfortable interaction distance, so for this experiment we will attempt to control for that by setting the robot height to approximately 85% of the participant's height, and focusing our attention on arm tension (firm or loose). The motivation behind this comes from the idea posited by Johnston and Thomas (1995) where characters that are stiff and rigid are perceived as more dominant than those that are not. Half (n = 9) of the participants were in the firm arms condition and the other half (n = 9) were in the loose arm condition. We ran each test twice, once where the robot approached the participants and once where the participants approached the robot. After the questionnaire was completed for experiment 1, the experimenter leads the participant through a maze of hallways and shows off other operations at Willow that weren't introduced in the original tour.

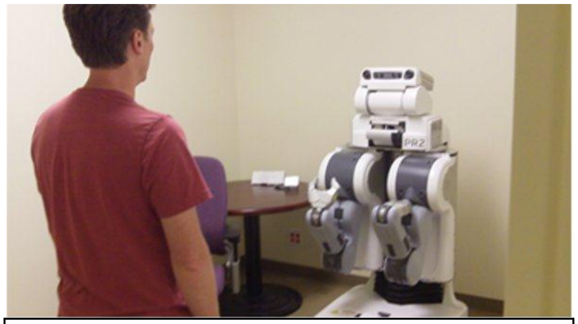
Figure 5: Experiment 2 room setup

Meanwhile, a confederate moves the robot to the location of experiment 2 & 3. The experimenter and participant arrive in the room shortly after the room is set up (See figure 5 for room setup). The experimenter then directs the participant where to stand, and explains that the robot will approach him, and when the robot is on edge of a comfortable distance for the participant (i.e. if the robot were to get any closer it would be uncomfortable for the participant), the participant should clearly say stop, and the robot's progress will be halted. A measurement from the robot's laser range finder will be taken at that point. The robot then backed up several feet and the participant was told to approach the robot before beginning the third and final experiment. The laser range finder reading was again taken at this point. Data from the laser range finder was compared between conditions (loose and stiff arms), and within conditions (human approach and robot approach) to see if there were any significant correlations. The experiment immediately proceeded into experiment 3 with no interference from the experimenter.

2.3.3 Experiment 3: Head nodding

Like Experiment 1, Experiment 3 involved robots using human-like social gestures, and in this case the gesture was head nodding while a participant read instructions. At the beginning of experiment 2, the experimenter will have given the participant a set of instructions to read to the robot (for a copy of the script, see appendix page 32) and asked the participant to read the instructions to the robot after they approached the robot at the end of experiment 2. Upon stopping when reversing into position for the second half of experiment 2, the robot transitioned into a finite state machine that will have its gaze stare at the participant but break every so often as if thinking approximately 35% of the time. After each instruction piece is read, the robot either nodded or did nothing, depending on what condition it has been assigned to. Half of the participants were in the head nod condition (n = 9) and the other half (n = 9) was in the no head nod experiment. We measured how quickly the participant read through the instructions. We are hoping to ascertain or at least approximate the level of trust or confidence that the participant has successfully imparted the instructions to the robot, and we feel that how quickly they are able to read the instructions is a good indication that they think the robot understands the instructions. This evidence will be compared with more qualitative measures that will be taken in a post experiment questionnaire, which will hope to use participants' perceptions to validate the metric. When the experimenter finished reading the instructions the robot will leave the room to give the participant the impression that the robot is performing the instructed task. After the robot left, the experimenter returned to the room and the two of them left the room together to administer the final questionnaire, while Willow researchers prepped the room for their study.

3 Results

3.1 Experiment 1

In the first experiment, we tested the gesture condition and verbal condition and evaluated the interaction on several measures (See Table 1 for averages). We measured 7 different items:

1) how much eye contact the participant gave the robot measured using a video tape and stop watch (Eye Contact),

2) whether or not they had the proper appraisal for the idea the robot was trying to communicate (Appraisal),

3) to what extent they felt the robot was trying to communicate with them (Attention),

4) how comfortable they felt with the communication (Comfort),

5) how unnatural the communication felt to them (Unnatural),

6) how intelligent the robot was (Intel.), and

7) how socially competent the robot was (Social Comp.)

(see appendix page 31 for a copy of the questionnaire). All scores were taken using a seven point Likert scale (a rating of 1 to 7) except for eye contact which was measured in seconds, and appraisal which was simply a binary representation of whether or not they correctly identified what the robot was communicating. Using R Statistics, we ran an analysis of variance (ANOVA) on each of the items separately, and unfortunately found no statistical significance across any of the variables. We also evaluated the data in terms of the gender of the participants but again found no significant results.

Table 1: Exp 1 Means							
Condition	Eye Contact (sec)	Appraisal (Binary)	Attention (Likert-7)	Comfort (Likert-7)	Unnatural (Likert-7)	Intel. (Likert-7)	Social Comp. (Likert-7)
Gesture	30.7	0.67	4.25	5.83	3.73	4.25	3.92
Verbal	28	0.86	4.86	5.57	3.29	3.86	3.86

3.2 Experiment 2

In the second experiment we did both a between subjects and within subjects proxemic test. The between subjects test was the loose/stiff arms condition and the within subjects test was the approach locus (robot approach human vs human approach robot). All measurements were taken and presented in meters using the data gathered from the robots laser range finder (See Table 2 for a complete list of averages). Using R statistics, we ran a 2x2 ANOVA on the data and found a few notable results.

Table 2 Experiment 2 Means		
Condition	Robot Approach (m)	Human Approach (m)
Loose arms	0.47	0.59
Stiff arms	0.54	0.62

We found significant results within subjects ($p < 0.001$) for approach distance. On average, participants let the robot approach them 10 centimeters closer than they would approach the robot. There was also a mildly significant trend ($p < 0.1$) suggesting that people averaged 15 cm distance over the robot approach distance when they approached the robot in the stiff arms condition, versus the 5 cm on average in the loose arms condition. This is exciting news because the bouncing arms were only visibly "bouncy" when the robot's acceleration changed after the initial approach, stopped and then backed up. Perhaps running a few more participants would have yielded significant results. We again checked for gender differences and again found nothing.

3.3 Experiment 3

In the third experiment we tested what role having the robot nod can play in building confidence that the participant-read instructions were understood and the task described was completed. We timed how long the participants took to read the instructions and recorded that in seconds; the rest of the questions were taken from a questionnaire comprised of seven point Likert scale questions:

1) How confident they were the robot understood the task read to them (Understood),

2) how confident they were the robot successfully completed the task (Completed),

3) how intelligent they felt the robot was (Intelligence)

4)how socially competent the robot was (Soc. Comp.)

(For a copy of the questionnaire see appendix page 33).

The results are summarized in table 3 below.

Table 3: Experiment 3 Means

Condition	Time (sec)	Understood (Likert-7)	Completed (Likert-7)	Intelligence (Likert-7)	Soc. Comp. (Likert-7)
Nod	58.38	5	4.44	4.88	4.44
No Nod	46.29	3.89	4.33	5.22	4.00

The results of this study reflected a mildly significant result ($p<0.1$) for the timing. Participants in the nod condition averaged approximately 12 seconds slower on reading the instructions than participants in the control condition. While this ran counter to our predictions, it become apparent very early on why this was the case. Participants in the nod condition slowed their reading and would not begin reading the next step until the robot had nodded, whereas participants in the control condition did not have such a pause. No other notable results were found, and again we also checked for gender factors but found none. It is also important to note that two participants timed out on this experiment in the no nod condition because they were waiting for the robot to give some indication (like a nod) it was ready to receive the instructions.

Table 4: First Impression		
Experiment	Intelligence	Social Competence
1	4.12	3.90
3	5.00	4.22

Also, at the end of this study we compared the appraisal of intelligence and social competence from experiment 1 and experiment 3 and we found a significant trend ($p<0.05$) reflecting that the participants had a significantly higher appraisal of the robot's intelligence the more they interacted with it. We had expected to see a first impression affect where the initial appraisals significantly correlated with later appraisals, so this was an interesting surprise, and will be explored more in the discussion section.

4 Discussion

4.1 Implications

The most significant finding we discovered in this study was that in the proxemics study, participants let the robot approach them closer than they would approach the robot. Previous work done by Takayama and Pantofaru (2008) were unable to find significant results in this area. The surprising issue was that this finding runs completely counter to both our hypothesis and the hypothesis of Takayama and Pantofaru. This may have happened because of a minute delay in the participant's request for the robot to stop and the robot actually stopping, although the p-value was very low, and the robot's stopping distance was not 10 cm, so this is certainly something that could be explored further in future studies.

There were also additional mildly significant findings. Of particular interest was the finding that participants would not approach the robot as close in the stiff arm condition as they approached the robot in the loose arms condition. This is particularly exciting, largely because the robot's loose arms were generally not apparent on the approach, but they were on the stop and reverse. This puts some empirical evidence to theories already utilized in the worlds of animation and acting. This suggests participants were more likely to perceive the tight-armed robot as intimidating and therefore would not approach it as closely as the loose armed robot. While these results are far from conclusive, they are very provocative and suggest that there is something very tangible that can be gleaned from the art world and applied to great effect in the world of robotics. The results from this study point to a rich and largely unresearched area of robotics, opening up a new area for researchers to explore.

While we were unable to confirm many of the other hypotheses we had at the start of this study, we were able to take away several valuable lessons that we can learn from and build upon in the future. Most important among these is that it seems everyone possesses a different concept of a robot, and these concepts can vary wildly from person to person. These variations seemed to play a significant impact on our data, as we could not find many significant results. This vital piece of evidence is humbling and shows us how much we underestimated the skeptical nature of the human mind. We cannot trick our participants with static, scripted gestures; our participants will know when they were being duped.

While this is very interesting, it does highlight a significant hurdle that robots must overcome in order to be able to adopt gestures that fulfill the functionally communicative roles we have set for them. We obviously want robotic gestures to appeal to both sexes, seeing as robots will not just be interacting with women. This means that our gestures must be rich enough to reliably get both men and women to suspend what they know about robots and other computing machinery, and merely look at it as a social agent. We know this is possible; our robotic friends of the science fiction realm have shown us that it is possible to get people of all walks of life and all backgrounds to suspend the concept of mere machine and build upon that further- the concept of social agent (Think WALL-E). In fact, we found that the more the participants interacted with the robot the more intelligent they appraised it. This promising evidence allows us to see that we are making headway in creating robots perceived as intelligent agents. The question then turns back on the researchers and it is three-fold: Do we have a deep enough understanding of how social gestures work in human-human interactions; do the social gestures reliably translate from human-human interaction to human-robot interaction; and do we have a robotic platform sufficiently sophisticated to adequately implement these gestures?

The answer to the first of these questions is both yes and no. There is a myriad of research across several disciplines of social science that identify, define, and even interpret a vast majority of the social gestures we use every day. Still I feel it would be foolish of us to claim we understand them all. Human-human interaction is so deeply nuanced and has so many levels that we frequently miss even obvious cues in our own personal interactions. Certainly researchers on the outside looking in on interactions can more easily see some things in an objective light, but so much is subjective that the researcher cannot see and accurately interpret everything. So in a sense, no we do not possess a complete understanding of how gestures work; however, the vast amount of data we do have ought to be able to give us something tangible at this point. Even if we fail to understand the entire picture in micro-fine detail, surely we possess enough understanding to make gestures work in a rudimentary sense.

So then this leads us to the second question: do social gestures from human-human interactions reliably translate to human-robot interactions? The answer to this is a resounding yes. Think of the robot named Sparky we discussed early on (Shibata, Yoshida, and Yamato, 1997). This robot displayed gestures reflecting in wide range of human emotion and behavior and the gestures were concluded to be largely interpreted as such. Importantly, Sparky was controlled by a human, who could quickly interpret the interactions and seamlessly choose an appropriate response to make the gesture fit with participant expectations. In fact we need not rely on empirical work to know that this is possible. Once again we can point to the world of science fiction and the social robots. Even robots as non-anthropomorphic as R2D2 are still looked at as autonomous social agents and unlike his partner C-3PO, he cannot even talk; he merely beeps and whistles. Under these circumstances R2D2 is able to pull off humor and even sarcasm as he treks all over the universe; however, like Sparky, he was non-human only in presentation. Underneath the hood, however, there was a human calling the shots, creating the right combination of bells and whistles to create these advanced interactions. So given this, we know for certain that the gestures can translate to human-robot interactions.

We then turn to the platform itself: Do we have a platform adequately sophisticated to implement these gestures? The answer here is no, at least for the time being. According to a computational theory made famous by Alan Turing (1950), the only thing that separates human from machine is the computer technology has not yet been advanced enough to rival that of human brain; the computational complexity, he argues, is the same. Therefore in theory, there are no permanent barricades preventing us from creating a machine as sophisticated as a human. The plain and simple reality is we just are not there yet. We tried to fake it in our experiment using the PR2, by implementing what is tantamount to a simple song and dance, and our participants were largely able to see through it. Little tells were everywhere, and if the participant deviated even marginally from our expectations, they could easily see the man behind the curtain. The level of sophistication required to pull off a rouse of this magnitude just was not there. All is not lost, however, as new advances in robots and computation are made every day. In the meantime we should not give up on the research, and if that means we must run experiments with a man behind the curtain a bit longer, then so be it. The data will be of significant value when the robots reach the level of sophistication to make these gestures feasible.

4.2 Future Work

Of course many areas for improvement become readily apparent after the studies have been completed. Among these, the first that seems to come to mind is a brief pre-experiment. This study seemed to suggest a prima facie link between peoples' responsiveness to the robot's gestures and any preconceptions they had about robots before they came to the study. We tried to control for peoples' preconceptions by screening for participants with experience with robots, but as the preface alluded to, robots in several forms are already very prolific in modern day society so the idea that we could find people that were blank slates on the concept of robots seems a bit unlikely after the fact. We know some participants came in with very strong preconceptions because of comments they made and/or questions they asked during the study; for instance one participant said robots are just machines doing what they are programmed to do, and therefore cannot be deemed as intelligent. A questionnaire that would provide us with more quantitative insight into those preconceptions would likely yield some interesting results and allow us to evaluate participants' responses and behaviors in terms of their pre-conceptions. Furthermore, a pre-experiment questionnaire would be useful in determining to what extent the participant's preconceptions of the robot mattered more than their first encounter with the robot in the first impressions part of the study. It is possible, though we believe unlikely, that the participant's preconceptions weighed so heavily that they colored both the first impression and the second impression. A future study would ideally be better able to control for this kind of variance utilizing this pre-questionnaire.

The first experiment also seemed to have one keen issue that would be worth exploring in the future. Oftentimes, participants were not sure when the transition occurred between when the robot stopped doing the calibration task and when it started gesturing at them in the gesture condition. Frankly, there is the potential to have a whole other study wrapped around this issue. One concept that is strongly worth exploring is the role eye contact, and more specifically interactive eye contact plays on these interactions. Much of human nonverbal communication comes from facial gestures and in particular the eyes. Perhaps if the robot did some attention-getting activity like waving at the participant, and then somehow loops that attention-getting activity until the robot senses it has the participant's eye contact; then it could begin the gesture with increased certainty that the participant has noticed the robot. In our own human on human interactions, this fits perfectly with what we would expect. A person does not just approach a stranger on the street and say, “Could you tell me the time?” There is usually some sort of introduction, such as, “Excuse me” or “I’m sorry to bother you.” Both of these introductory phrases secure the person’s attention, and usually we do not proceed into the next part of the question until the introduction has been acknowledged (often times merely with eye contact). The experiences of this study suggest usefulness for this sort of interaction to be implemented and utilized by a robot.

Experiment 1 also suggested the importance of the robot keeping and maintaining eye contact with the participant throughout the gesture and using it to reinforce the idea that the robot was gesturing at the participant. The eye contact the robot did make in the first experiment was just having the robot’s head turn to face where the participant was supposed to be. Some of our more tech savvy participants really wanted to test this to see to what extent the robot was interacting with them and got up and moved around the robot to see if it tracked them, and of course it did not. This for many of them was sufficient to reveal the man behind the curtain, and their evaluation of the machine was significantly lower. Had the robot been more dynamically engaged with the participant, the robot would have followed the person with its head movement. The drawback is that the head tracking would have been “jerky” and, for lack of a better term, robotic. This could have resulted in lower appraisals of the interaction because it would also effectively reveal the man behind the curtain. The result needs to rely on a smooth implementation of face tracking and if that can be handled successfully, we believe the robot will make great strides in its effectiveness in interacting with humans.

Experiment 2 had one major fault that we were unable to address and remedy while simultaneously maintaining the integrity of the experimental design within the given time frame. The loose arms condition in the experiment did not yield as much bouncing or movement as we would have liked to have seen, and a large part of that is that the robot has a naturally smooth ride. Peoples’ arms bounce when we walk because we have two legs and therefore a certain level of bounce is required to move and that can be either exaggerated or minimized in given situations. The PR2 has wheels on a stable platform and therefore was lacking that natural level of bounce. Combine that with level, smooth floors and the fact that proxemics studies require the robot to move in a straight line towards the participant, there are simply not many opportunities to have outside forces (such as momentum, etc.) act upon the robot's arms. One interesting observation we did make, however, is that with the arm controllers off, if the robot does not move in a straight line, the arms will drift and bounce quite eccentrically. The movement of the arms in those instances was vaguely reminiscent of the movement of Captain Jack Sparrow's (Johnny Depp's) arms in the *Pirates of the Caribbean* series. While we were intrigued by this movement, it was discovered too late to incorporate into our study in any form, but was something we certainly wanted to note here if we hoped to get more concrete results involving the loose arm/stiff arm question for increasing levels of comfort in human interactions in future work.

The final experiment was difficult in many senses. First and foremost, the notion of trust means many different things to many different people. While we tried to get around that by not explicitly referring to the notion of trust, the participants' methods for evaluating the questions we had them answer regarding their confidence in the robot's performance varied greatly depending on how their particular sense of trust is formulated. Some participants noted that they had no confidence in the robot's performance because they did not see any tangible results. Others were perfectly willing to claim with absolute certainty that the robot performed the task properly because the robot merely looked at them while they were reading instructions. One idea to help control for this extreme variance is to make the instructions much simpler. One reason why the variance could have been so great is that as noted above, participant's came in with widely varying preconceptions of robots. If any of the participants were technically inclined in the slightest (and given our population sample was mostly Stanford students and other residents of Silicon Valley, it is fairly safe to assume many were) they may have quickly realized that if the robot successfully performed the instructed task, it would be a major technological breakthrough that would have likely been wildly publicized. So perhaps this is what gave our rouse away to some of our more jaded and cynical participants, and resulted in them demanding proof of task completion before trust would be granted to the robot. The thought is, however- if we make the robot's task more achievable, and ergo more plausible, people may be more willing to be trusting of the robot's performance. This can be done merely by simplifying the task. We made the task long by design to be sure that the robot's head nodding would not be overlooked, but participants seemed to readily notice the nodding, so it may be safe to sacrifice the number of nods in order to simplify the instructions a bit, and perhaps rein in some of the drastic variability of the participants' willingness to grant trust to the robot.

5 Conclusions

We started this research hoping to learn something valuable about robots and in the end we rediscovered something valuable about ourselves that we had underestimated. We are all individuals living in a social world and not one of our interactions with any person is ever the same twice. Our interactions are highly dynamic and very responsive to even the smallest changes in the persons or the environment. Our robot studies were unable to keep the interactions dynamic enough to fool the ever-critical human brain, but even this humble reminder of our robot's shortcomings provides us with valuable insight. Now our focus needs to be how do we take this newly rediscovered information about ourselves and turn it in to something of value as we continue to strive for advancements in the field of human-robot interaction.

As mentioned previously, it will remain important for us to continue to do research on what roles these gestures can play in facilitating human-robot interactions, but until we are able to have robots interact with humans in a more dynamic fashion, we will need to leave the interpretation of the social scenarios and decision-making to the mind of another human. The area is still rich for mining valuable data on human-robot interaction, and if we continue to explore and broaden our understanding of the roles these gestures can play, when robot technology is ready to hold dynamic interactions with humans, we will have an impressive arsenal of tested and proven gestures ready to be implemented and tested autonomously. We know technology will advance to that point; now it is only a question of how long will we have to wait to see it.

How long until the robots of science fiction can become reality? Any definitive answer to that question would merely seem a mirage, constantly moving farther into the distance the closer we get. Eventually is the best and most accurate approximation we can give right now. Humans did not evolve into the social animals we are today overnight, and neither should we expect robots to develop in a similar matter. It took billions of years for humans to go from nothing to what we see today. It has only taken robots the better part of a century to achieve the sophistication we have today, and the future possibilities seem virtually limitless. So perhaps the most appropriate answer to the first question posed in the preface- "Can robots ever hope to be human?", the answer is yes, eventually.

Appendix

Willow Garage Study platforms

The platform used for the other half of the study (conducted by Willow Garage staff) used a different robotic platform called the Texai (Figure 6). As can been seen the robot is merely a mobile webcam/LCD screen combination. This helps insure that there is no pollution between studies. There were also times when Willow Researchers conducted a different study prior to the research discussed here. Again in those cases a different robot platform was used, in this case the Turtlebot (See Figure 7). The robot is essentially an iRobot Roomba with an Xbox Kinect on top.

Figure 6: Texai

Figure 7: Turtlebot

Experimental Protocols

1. Participant enters office, and receptionist has him wait in lobby for experimenter.
2. Experimenter enters lobby and greets participant; takes them on a brief introduction to the PR2.
3. Tour ends in pool room where experimenter says, "I have a consent form for you to fill out."
4. Experimenter gives the participant consent forms to fill out and shows the participant to a table near the PR2 that is going through pseudo-calibration motions .
5. If asked by participant, Experimenter should reply: "This PR2 is currently calibrating its cameras. It's a long and involved process."
6. Experimenter leaves forms with participant and goes to other side of the room, and appears to be engrossed in paperwork, but really is starting a camera.
7. After the participant completes the consent form, experimenter says he forgot to print a document, and asks the participant to wait at table while he goes to get it.
8. Shortly after the experimenter leaves, the robot stops the calibration task and starts one of the experimental conditions.
9. After the experimenter returns with the forms the forms, Experimenter asks if everything went OK.
10. If the participant says something about the PR2 needing help or needing something, they should be prompted to guess what exactly they needed (e.g. What's wrong? What does it need?)
11. Participant is then led to a different room and given a post Experiment 1 questionnaire, while experimenter notes what comments participant made, and turns off the camera. If necessary the experimenter can clarify which robot that they are answering questions about.
12. Experimenter leads participant on long route to primary study room with detours to the shop and burn room,, while a confederate drives PR2 to primary location room, and adjusts height according to height of participant. Also sets the arms to be either loose or firm as the pre-determined condition requires.
13. Participant enters room for next experiment; experimenter directs participant where to stand.
14. Experimenter will give a script to the participant
15. Experimenter: "In this study the robot will approach you. All you need to do is tell it to stop when it's on the edge of your comfort zone. Then take a step back, and the robot will too. Then approach the robot and read him these instructions. Be sure to speak clearly so the robot can process the instructions. The robot will then go off to perform the instructions. Wait in the room for further instructions."
16. Confederate will control robot remotely for the "stop" command, and also for the head nodding if that is the condition to which the participant is in.
17. The experimenter returns to the room shortly after the robot leaves and takes the participant out of the room to administer final questionnaire.
18. Then the Texai experiment begins.

Post Experiment 1 Questionnaire

1. To what extent did you feel the robot was trying to get your attention?

(Not at All) 1 2 3 4 5 6 7 (Definitely)

Please describe what happened with the robot.

____

____

____

1. What do you think the robot was trying to communicate?

____

____

____

____

1. How comfortable were you with the gesture?

(Not at All) 1 2 3 4 5 6 7 (Definitely)

1. How **unnatural** did the gesture feel to you?

(Not at All) 1 2 3 4 5 6 7 (Definitely)

1. How would you rate the intelligence of the robot?

(Not at All) 1 2 3 4 5 6 7 (Definitely)

1. How would you rate the social competence of the robot?

(Not at All) 1 2 3 4 5 6 7 (Definitely)

Experiment 3 Participant Script

Note: Please try to speak as clearly as possible so that the robot can correctly process the instructions

I am going to read a list of steps for completing a ball fetching task that must be completed in the order I say them:

1. Exit the office.
2. Turn left and proceed down the hallway until another hallway opens on the right.
3. Turn right onto this hallway and drive forward past four offices on the left.
4. Enter the next office on the left and pick up red ball sitting on the desk.
5. Exit the office.
6. Turn right and drive forward down the hallway past two offices on the right.
7. Enter the next office on the right and place the red ball on the desk.
8. Exit the office.
9. Return to recharge station.

Post Experiment 3 Questionnaire

1. How confident were you that the robot understood the instructions you read?

(Not at All) 1 2 3 4 5 6 7 (Definitely)

1. Why do you feel so (un)confident?

____

____

____

____

1. How confident were you that the robot successfully completed the task?

(Not at All) 1 2 3 4 5 6 7 (Definitely)

1. Why do you feel so (un)confident?

____

____

____

____

1. How would you rate the intelligence of the robot?

(Not at All) 1 2 3 4 5 6 7 (Very)

1. How would you rate the social competence of the robot?

(Not at All) 1 2 3 4 5 6 7 (Very)

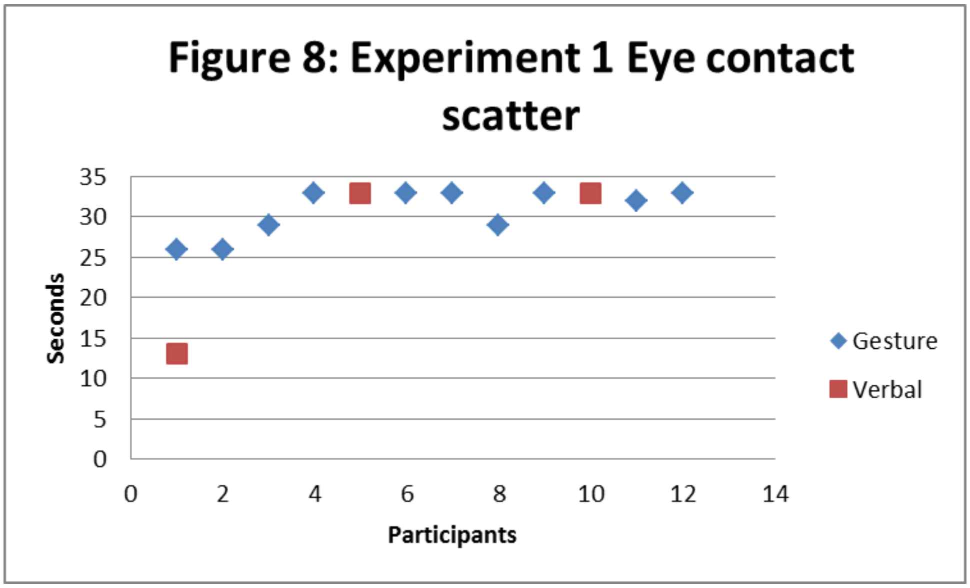
Figure 8: Experiment 1 Eye contact scatter
Seconds
35
30
25
20
15
10
5
0
0
2
4
6
8
10
12
14
Participants
Gesture
Verbal

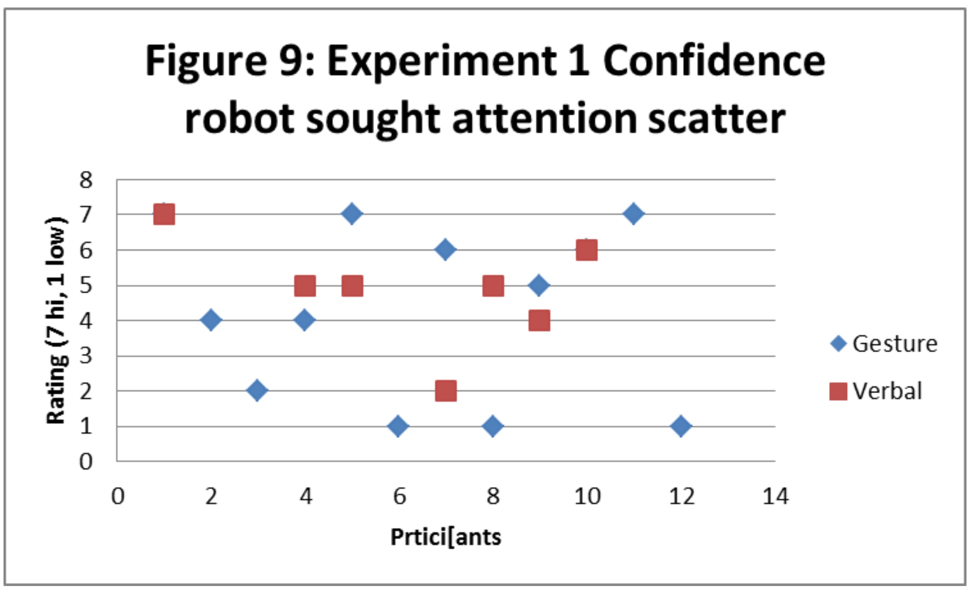
Figure 9: Experiment 1 Confidence robot sought attention scatter
Rating (7 hi, 1 low)
8
7
6
5
4
3
2
1
0
0
2
4
6
8
10
12
14
Prtici[ants
Gesture
Verbal

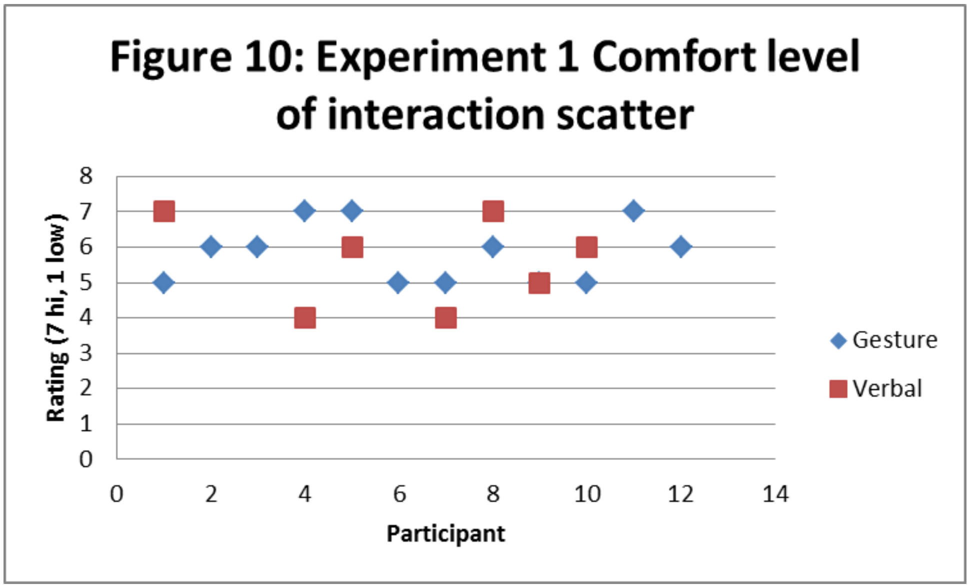

Figure 10: Experiment 1 Comfort level of interaction scatter

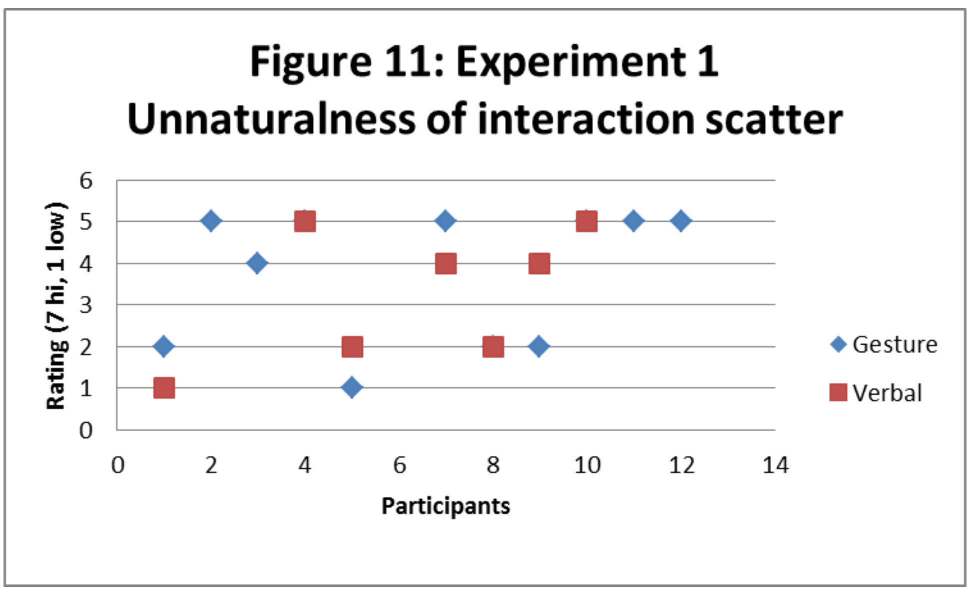

Figure 11: Experiment 1 Unnaturalness of interaction scatter

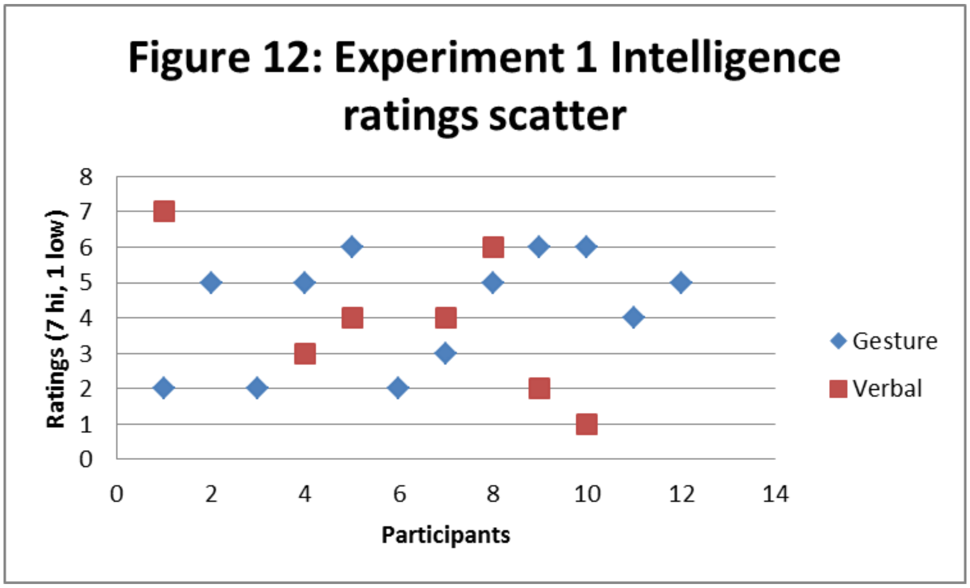
Figure 12: Experiment 1 Intelligence ratings scatter
Ratings (7 hi, 1 low)
0
1
2
3
4
5
6
7
8
0
2
4
6
8
10
12
14
Participants
Gesture
Verbal

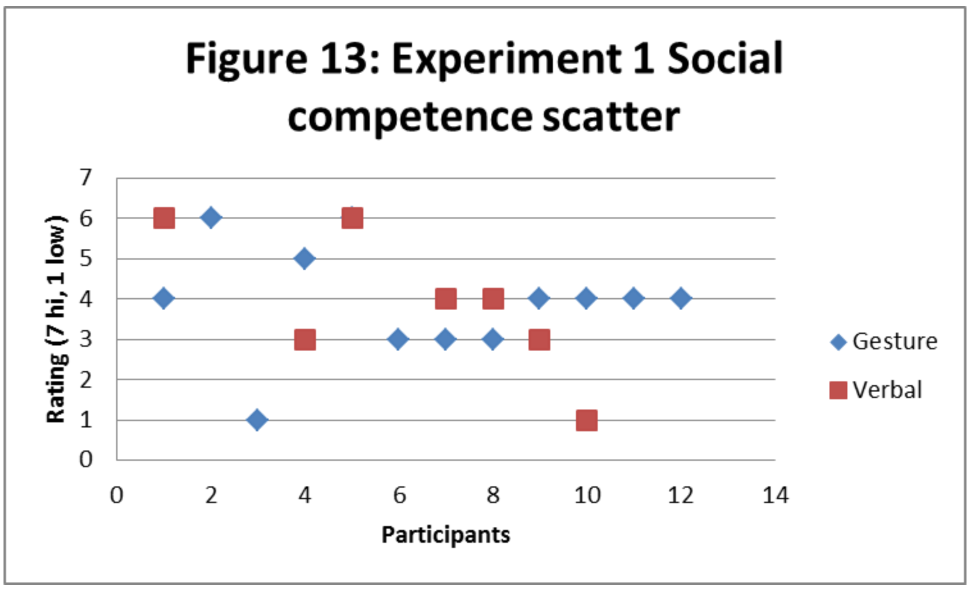
Figure 13: Experiment 1 Social competence scatter
Rating (7 hi, 1 low)
0
1
2
3
4
5
6
7
0
2
4
6
8
10
12
14
Participants
Gesture
Verbal

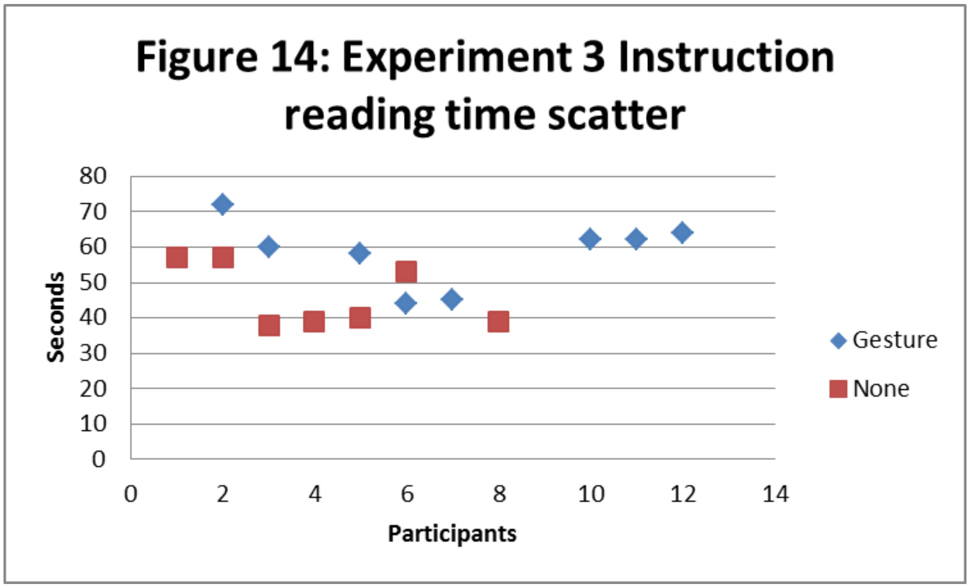
Figure 14: Experiment 3 Instruction reading time scatter
Seconds
80
70
60
50
40
30
20
10
0
0
2
4
6
8
10
12
14
Participants
Gesture
None

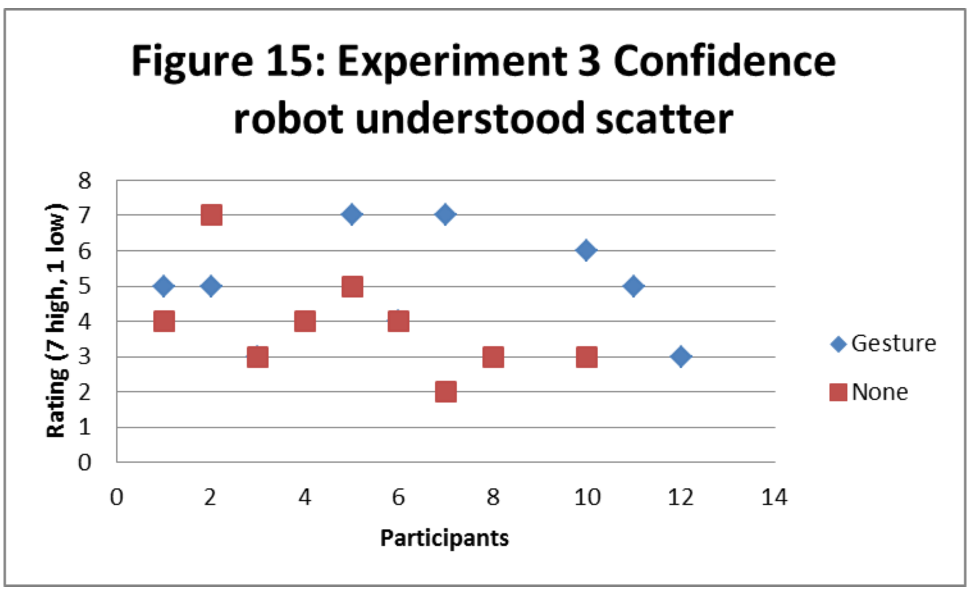
Figure 15: Experiment 3 Confidence robot understood scatter
Rating (7 high, 1 low)
8
7
6
5
4
3
2
1
0
0
2
4
6
8
10
12
14
Participants
Gesture
None

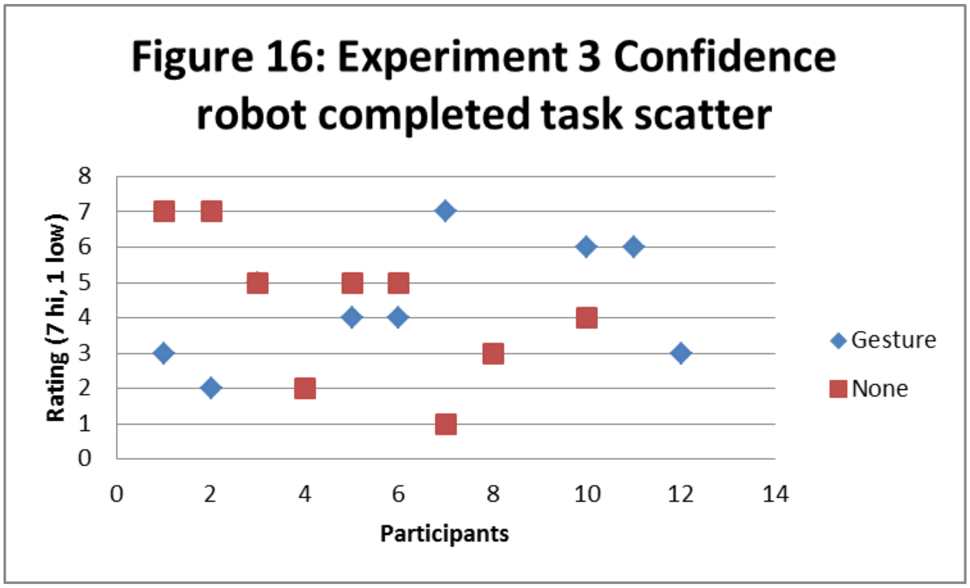
Figure 16: Experiment 3 Confidence robot completed task scatter
Rating (7 hi, 1 low)
8
7
6
5
4
3
2
1
0
0
2
4
6
8
10
12
14
Participants
Gesture
None

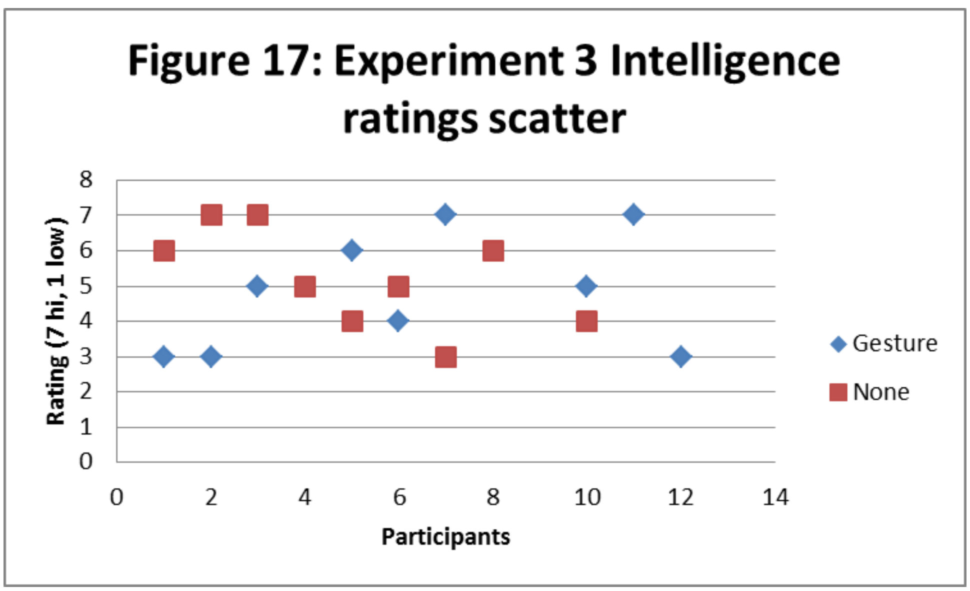
Figure 17: Experiment 3 Intelligence ratings scatter
Rating (7 hi, 1 low)
8
7
6
5
4
3
2
1
0
0
2
4
6
8
10
12
14
Participants
Gesture
None

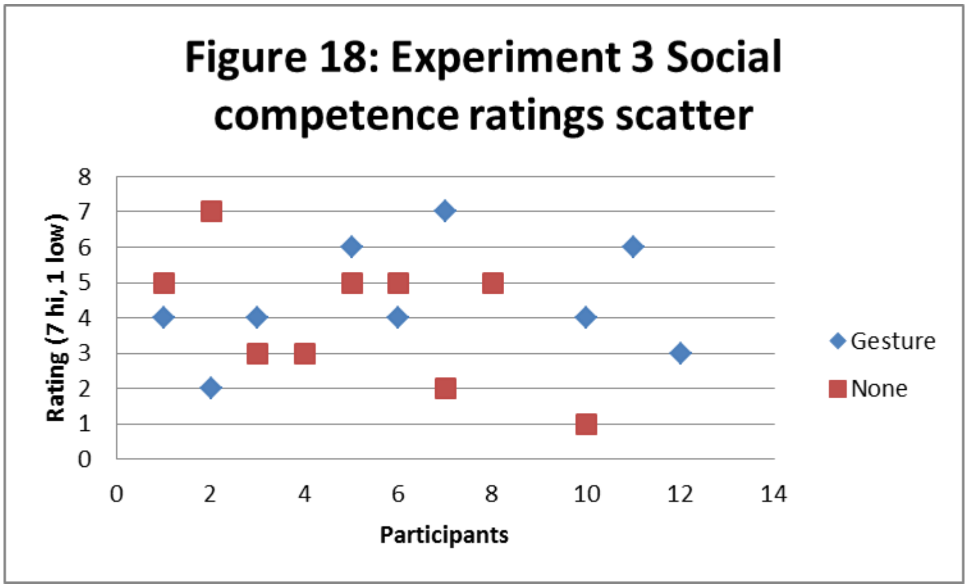

Figure 18: Experiment 3 Social competence ratings scatter

References

Adams, D. (1978). *The Hitchhiker's Guide to the Galaxy*. United Kingdom: Pan Books.

Bruemmer, D., Few, D., Goodrich, M., Norman, D., Sarkar, M., Scholtz, J., Smart, B., & Swinson, M. L. (2004). "How to Trust Robots Further than We Can Throw Them." *In CHI Extended Abstracts.* p. 1576-1577.

Byers, Z., Dixon, M., Smart, W. D., and Grimm, C. (2004). "Say Cheese! Experiences with a Robot Photographer." *AI Magazine.* 25(3). p. 37-46.

Feldman, R. S., Phillipot, & P., Custrini, R. J. (1991). Social competence and nonverbal behavior. In R. S. Feldman & B. Rimé (Eds.), *Fundamentals of Nonverbal Behavior* (p. 329-350). Cambridge: Cambridge University Press.

Freedy, A., DeVisser, E., Weltman, G. & Coeyman, N. (2007). "Measurement of Trust in Human-Robot Collaboration." *In Proceedings of the 2007 International Conference on Collaborative Technologies and Systems*. p.

Hall, E. T. (1966). *The Hidden Dimension.* New York: Anchor Books.

Johnston, O., Thomas, F. (1995). *The Illusion of Life: Disney Animation*. New York: Disney Editions.

Kurtz, G. & Lucas, G. (1977). *Star Wars Episode IV: A New Hope.* United States: LucasFilms.

Lee, J. D. & See, K. A. (2004). "Trust in Automation: Designing for Appropriate Reliance." *Human Factors.* 46. p. 50-80.

Morris, J., Collins, L., Lasseter, J., & Stanton, A. (2008). *WALL-E*. United States: Pixar Animation Studios.

Oosterhaut, T. V. & Visser, A. (2008). "A Visual Method for Robot Proxemics Measurements." *In Proceedings of 3rd ACM/IEEE International Conference on Human-Robot Interaction Workshop on Metrics for Human-Robot Interaction.*

Scheef, M., Pinto, J., Rahardja, K., Snibbe, S., & Tow, R. (2000). "Experiences with Sparky: A Social Robot." *Proceedings of the Workshop on Interactive Robot Entertainment.*

Shibata, T., Yoshida, M., Yamato, J. (1997). "Artificial Emotional Creature for Human-Machine Interation." *Proceedings of IEEE Intl. Conf.Systems, Man, and Cybernetics*. 3. p. 2269-2274.

Takayama, L. & Pantofaru, C. (2009). "Influences on proxemic behaviors in human-robot interaction." *In Proceedings of the 2009 IEEE/RSJ international conference on intelligent robots and systems*. p. 5495–5502.

Turing, A. (1950). "Computing Machinery and Intelligence." *Mind.* p. 433-460.