Experiences with Sparky, a Social Robot

Mark Scheeff, John Pinto, Kris Rahardja, Scott Snibbe, and Robert Tow
Interval Research Corp.
1801 Page Mill Rd., Bldg. C
Palo Alto, CA, 94304 USA
{scheeff, pinto, rajardja, snibbe, tow}@interval.com
contact@markscheeff.com

Abstract

In an effort to develop new ideas in human-computer interface and interactive robotics, we have a built a lifelike teleoperated robot. Our robot uses gesture, motion and sound to be social with humans in the immediate vicinity. We explored human/robot interaction in both private and public settings. Our users enjoyed interacting with Sparky and treated it as a living thing. Children showed more engagement than adults, though both groups touched, mimicked and spoke to the robot and often wondered openly about its intentions and capabilities. Evidence from our experiences with a teleoperated robot showed a need for next-generation autonomous social robots to develop more sophisticated sensory modalities that are better able to pay attention to people. We close by suggesting how our experience with Sparky might influence the design of future physical interfaces to computing.

Keywords

Personal robotics, emotion, affect, animation, gesture, physical interface, perceived intelligence, motion control.

Introduction

Our project began as an exploration in new human-computer interfaces, not robotics. As a thought experiment, we began by listing some properties of current interfaces under “workstation” and then listing their opposites under “anti-workstation”.

Workstation	**Anti-Workstation**
Fixed in space	Moves around
No gestural cues	Rich gesture
Symbolic communication	Affective communication
Few sensors	Rich sensing
Asocial	Social with others
Interface learned	Interface already known

Combining elements on the “anti-workstation” list, we imagined an interface with a dynamic, expressive body that communicated using affect and could interact naturally with others in its environment. In short, we imagined a device much like a living thing as the antithesis of a desktop workstation. We thus began building our robot, Sparky, as a way to explore radically different approaches to HCI.

Because we considered our work a broad ranging exploration, our robot’s goal was simply to be social with humans in its vicinity. Obviously, this is a fairly new goal for a research robot. Would people find it compelling or disturbing? What behaviors would people exhibit around the robot? What new skills does a robot need to develop when it is in a social setting (and what skills can it forget)? Lastly, how do experiences with this robot suggested new ideas in interface design?

Prior Work

In searching for inspiration in creating lifelike characters, we first looked towards the principles of traditional animation and cartooning. Thomas and Johnson of Disney Studios have made one of the best reference books for understanding the techniques of traditional animation and how to evoke emotional state through motion (Thomas and Johnston 1981). Jack Hamm has also written some good, simple books introducing the principles of the cartoonist (Hamm 1982).

The emerging field of *affective computing* has provided motivation and justification for our work. Picard (Picard 1997) defines affective interfaces as those which relate to, arise from, or deliberately influence emotions. We focus on a particular aspect of affective computing; emotional expression by computers for the purposes of social interaction with people.

The computer graphics community has explored many avenues towards creating realistic (screen-based) animated characters. Ken Perlin’s *Improv* system (Perlin and Goldberg 1996) provides the foundation of our approach to movement. Perlin ignores traditional kinematic and dynamic approaches to control (Badler, Barsky and Zeltzer 1991) and instead synthesizes movement using continuous periodic noise functions. Other approaches to driving computer graphic characters include Blumberg's ethologically inspired models for action selection (Blumberg and Galyean 1995), and Karl Sims' work on evolutionary feedback learning methods (Sims 1994). Recent work in using digital signal processing techniques to alter the quality of motion has also inspired us, and we have adapted some of these techniques for use in our hybrid system (Amaya and Bruderlin, 1996)(Bruderlin and Williams 1995).

Breazeal (Breazeal and Scassellati 2000) has built an animated head, called Kismet, that can sense human affect through vision and sound and express itself with emotional posturing. Her system uses a feedback loop of affective perception and affective behavior (in motion and speech) to explore social exchanges with humans which she models on those of the caretaker-infant dyad.

Lastly, Isbister (Isbister 1995) has written an excellent discussion on the difference between traditional notions of intelligence, which emphasize the construction of an accurate "brain", and the idea of perceived intelligence, which emphasizes the perceptions of those who experience these brains. This work helped us to understand how users saw intelligence in unfamiliar people or devices.

Design Goals

Pulling together this prior work along with our own experiences, we developed the following set of design goals to guide our creation of a lifelike robotic character for humans to interact with.

- *Interesting, eye-catching*: We wanted to make sure that whatever we built would attract and maintain attention.
- *Use an active, integrated body:* Typical research robots are not designed to display emotion or interact with people (e.g. The Nomad™ robots). They often consist of a static structure with some effector that reaches out to manipulate the world. Rather than simply adding an "affective output device" to an existing robot, we wanted a device where most of the body moved and shifted (much like a real creature) and which gave the impression of an integrated whole. We wished users to simply see the robot as happy, sad, angry, etc. rather than looking to some particular appendage to "read" its state.
- *Use a face:* Inspired by Darwin (Darwin 1872) and traditional animation techniques, and with a desire to portray clear, strong emotions, we decided to use a face (integrated with the body of course). In a user study on an early prototype, our design was found to be effective at communicating basic emotions (Schiano et al. 2000).
- *Use natural motion:* The term "robotic" is normally used to indicate stiff, unnatural movement. We wanted our robot to move in a natural manner and therefore decided to spend some of our time writing new algorithms for creating lifelike "robotic" motion. Our algorithms were based on work done in computer graphics (Perlin and Goldberg 1996).
- *Avoid the "Uncanny Valley":* Masahiro Mori has articulately pointed out that the progression from a non-realistic towards a realistic portrayal of a living thing is nonlinear (Adapted from Reichard 1978). One might naively begin with the assumption that the closer a robot resembled a living creature, the more likely humans would be to feel compassion and familiarity. However, Mori has observed that there is a striking "uncanny valley" (Figure 1), as similarity becomes almost, but not quite perfect. The subtle imperfections in an imperfect, realistic recreation of a human or other animate creature can be disturbing and revolting (the strongest example of this being a moving corpse). His work leads us to conclude that cartoons or other simplified representations of characters may be more acceptable than more complicated and realistic representations.

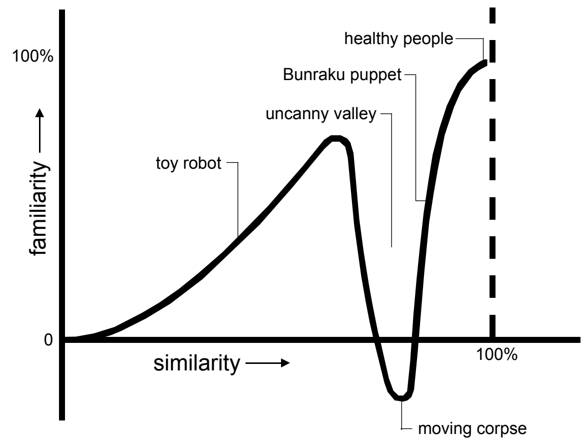

Figure 1. Masahiro Mori's uncanny valley. Familiarity increases with similarity until a certain point, when a realistic representation with subtle mistakes proves far more disturbing and far less believable then a less-realistic portrayal (Adapted from Reichard 1978).

- *Teleoperation, not Artificial Intelligence:* Originally, we wanted to build a fully autonomous creature and study people's responses to it. We decided to only study human response and to use a remote human operator as the robot's "brain". There is a significant amount of work being done in creating autonomous robots but little in understanding how people respond to them. We believed that teleoperation would allow us to generate and study complex social interactions that would be impossible or very time consuming to generate autonomously. Behaviors that were popular with participants could be developed to operate autonomously in future robots; conversely, unsuccessful or unnecessary capabilities could be eliminated without time consuming programming and testing. Lastly, because we wished to see successful behaviors incorporated into future devices, we avoided demonstrating capabilities that are clearly some years away, such as natural language processing and comprehension.

Our Robot, Sparky

Sparky is a small robot about 60cm long, 50cm high and 35cm wide (figure 2). It has an expressive face (figure 3), a movable head on a long neck, a set of moving plates on its back and wheels for translating around the room. Even though we send our robot high-level commands from a remote location, Sparky seems autonomous to those around it.

A remote operator manifests the personality we have constructed for Sparky. During operation, Sparky is usually a friendly robot, approaching anyone in the vicinity while

Figure 2. Our robot in action. In the top row, left to right: afraid and inquisitive. Below: happy and angry.

smiling and making an occasional happy utterance. Sometimes, though, our operator will command Sparky to act sad, nervous or fearful. If our robot suffers abuse, the operator can switch it into the “angry” emotion and, in extreme circumstances, even charge the abuser head on. Other emotions are expressed depending on the situation.

Because of the way we control our robot, Sparky makes extensive use of its body. Our operator will often track someone by moving Sparky’s head to keep the robot’s eyes locked on to the human’s. When curious, Sparky can crane its neck forward. If it’s very excited, or afraid or being touched, we can command Sparky to flair up the hackles on its back, a gesture reminiscent of a cat. Sparky can also be made to mimic people if they move their head in a “yes” or “no” motion. Lastly, we keep our robot moving around the room a lot, orienting himself freely to people.

Sparky is always moving and shifting its joints, much like a living creature. The type and amount of ambient motion is correlated to its current emotional state. We can also cue Sparky to make vocalizations, which sound something like muffled speech combined with a French horn. Just as in the case of ambient motion, the affective content of each sound is correlated to Sparky's emotional state. There are several sounds available in each state. Through a combination of facial expression, posture and the general quality of motion, Sparky can express nine different emotional states: neutral, happy, sad, angry, surprised, fearful, inquisitive, nervous, and sleepy. A short video of the robot can be downloaded from www.markscheeff.com.

Implementation Details

System Overview

Our system consists of a robot controlled by a base station (figure 4). The robot has all of its sensing, actuating and computing resources on board and is completely untethered. The base station has a modified game controller, a video monitor and speakers. Video and audio from the robot's head are fed wirelessly to the base station. A human operator at the base station uses the game controller to send commands to the robot. Controlling Sparky is similar to giving directions to an actor on a stage; the operator specifies some movements explicitly and then sets an emotional state that the robot uses to control the rest of its movements. The computer on board the robot interprets the operators commands and drives the motors.

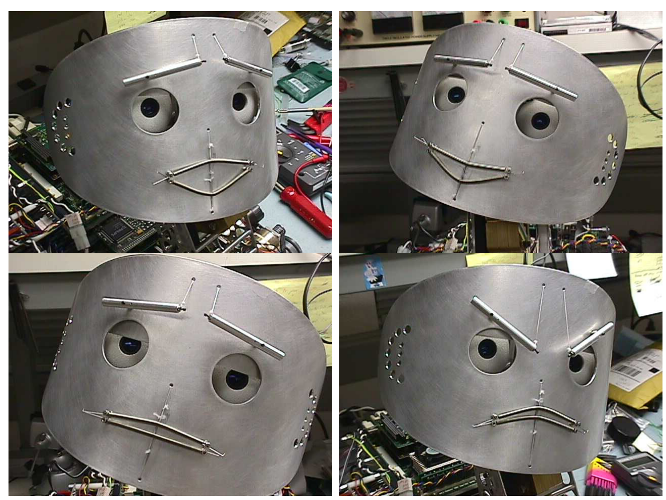

Figure 3. Range of facial expressions. Clockwise from upper left: surprise, happiness, anger, sadness.

Robot Design

Sparky has a Pentium 200 PC/104 stack with cards for networking, sound production and power regulation. Additionally, there are three motion control cards in the stack. These control cards, together with custom amplifiers, give us 10 channels of motion control on the robot. About 75% of the electronics are off the shelf.

The first four degrees of freedom (DOF) on the robot are dedicated to facial expression (figure 5a). The eyebrows and eyelids are each a single DOF. Each lip has a single DOF and can bend in both directions, allowing us to make, for instance, a true frown or a thin smile.

The remaining six DOF are dedicated to "whole body" motions (figure 5b). At the head, there are two DOF that allow for 180 degrees of yaw and 100 degrees of pitch. Where the long neck meets the body (a joint analogous to the human waist), there is a single DOF joint that can portray postures such as sitting up straight or leaning forward. Two DOF are used in the front of the robot, one on each drive wheel. Finally, there is a single DOF used on the back to allow it to rise up or flutter down. With the exception of the motors and some transmission components, the robot's body is completely custom made.
Lastly, the robot has two commercial wireless data links: one for receiving commands from operators and a second for transmitting video and audio from the robot's head to the operator's monitor.

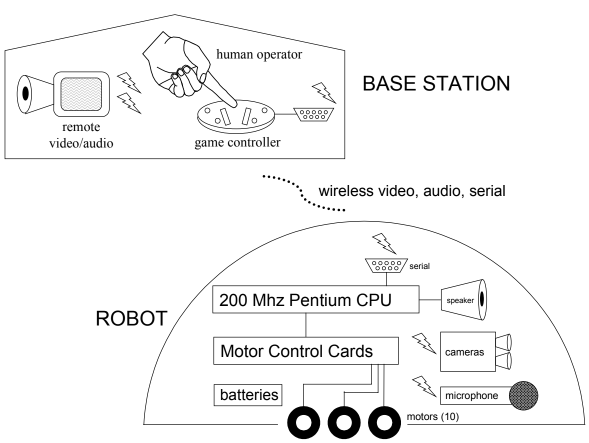

Figure 4: System Overview

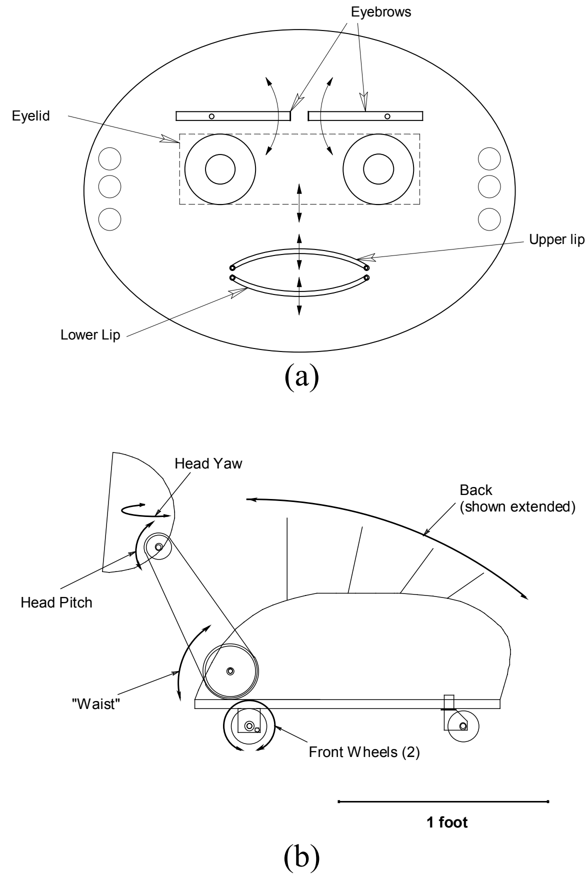

Figure 5. Schematic diagram of the robot. In (a) the four facial degrees of freedom are shown: eyebrow, eyelid, upper and lower lip. In (b) the six degrees of freedom for the body are illustrated: head yaw and pitch, neck, wheels and back. The scale is only accurate for the body illustration.

Software

Our software system consists of mechanisms for high level remote control. The components are a layered motion synthesis engine and an audio engine.

The principle of the motion control system is that of layered motion synthesis, inspired by Perlin's *Improv* System (Perlin and Goldberg 1996). Low-level periodic noise functions (Perlin 1985) are used as generation primitives, in a manner analogous to musical synthesis. From these primitives, *actions* are constructed which consist of a mapping from periodic functions to mechanical degrees-of-freedom. Multiple actions are combined by the *motion compositor* which generates final motion by weighting the actions appropriately. The varying of the weights for the active actions constitutes the behavioral aspect of the robot and in our system comes from a human. Finally, a signal-processing layer creates overall changes to affect with a simple bandpass filter (Bruderlin and Williams 1995). The output of the entire system is a continuously varying position signal for each degree-of-freedom. For the complete details on our motion architecture, and the details of implementation, please see our earlier publication on this topic (Snibbe, Scheeff, and Rahardja 1999).

Our audio system consists of an engine for creating nonsensical emotional speech. Sound clips for each emotion are generated by taking strong, affective, human speech and using Linear Predictive Coding (LPC) to extract a filter function. This filter is then used as the excitation function for a synthesized trombone. This results in trombone-like "speech," similar to the manner in which adults speak in Peanuts animated cartoons.

Robot Operation

Using a modified video game controller (figure 6), the operator can control four parameters remotely: gaze direction, platform motion, sound output and emotional state. The left joystick is used to control gaze direction (pitch and yaw for the head). The right joystick controls the two front wheels thus allowing the operator to move the platform around a space. One button cues the robot to make a sound appropriate to its emotion. Finally, the remaining buttons are each mapped to an emotional state The robot's emotional state completely determines facial expressions and the position of the waist and modifies the interpretation of the gaze and motion joystick commands to keep the emotional picture consistent. Our motion compositing software allows us to get realistic, lifelike motion without having to directly control all 10 DOF. By handling the low-level tasks, this software allows the operator to focus on creating interesting interactions with users.

Figure 6. Videogame controller. We use a retrofitted off-the-shelf video game controller for high level control of gaze, movement and emotion.

Other Devices

Several products, all toys, resemble some aspects of our robot. We introduce them here in order to make an immediate comparison with Sparky's features and capabilities.

Sony's robot dog, Aibo™, is an advanced toy that can walk, show emotions, play games and learn. Similar to our robot, it has a very expressive body and uses only simple sounds. Both the Furby™ and the Actimates Barney™ have basic sensing, simple motions and a number of playful interactions. Unlike Sparky and Aibo™, they rely heavily on vocalizations to communicate and use much simpler gestures.

All three of these robots are autonomous while Sparky is teleoperated. As mentioned above, teleoperation allows us to easily try out a variety of sophisticated interactions while concentrating most of our efforts on studying human response. Many of the interactions we observed (detailed below) go well beyond what even the most sophisticated robot, Aibo™, is capable of but are still realistic in a few years time. Ideally, our discoveries regarding human response would feed back into the next generation of these autonomous robots as well as future interface work.

Observing Sparky and People

To explore our research questions and to see how well our design goals were met, two venues were chosen in which to explore robot/human interaction, one in the lab and the second in public.

In the Lab. Thirty external subjects were recruited for 17 trials in our internal lab (singles and dyads). Participants ranged from age 6-81 and were not technically sophisticated. Subjects answered several background questions, interacted with the robot for about 15 minutes, and then discussed the experience with the interviewer in the room.

The robot operator was behind a one-way mirror. Interactions between the robot and the subject were necessarily chaotic and so we did not follow a strict protocol. In general, we started by paying attention to the subject and expressing happiness. After that, we tried simply to react reasonably to the subject's actions while still manifesting the personality we have described above.

In Public. Our lab testing allowed us to ask in depth questions, but the interactions, constrained to a small room, sometimes seemed a little unnatural. Therefore, we sought a public place where subjects could interact in whatever way and for whatever length of time they desired. We set up at The Tech Museum of Innovation in San Jose, an interactive, hands-on science museum.

The tests were conducted 2-3 hours a day for six days. The robot was released for an hour at a time to "wander" in an open area. There were no signs or explanations to indicate the purpose of the robot. We wanted to get unscripted, spontaneous reactions from a random sample of museum visitors. Unbeknownst to patrons, our operator was controlling the robot from a hidden room. Our operator manifested the same robot personality as in the lab. We videotaped interactions but did not ask any formal questions.

Reactions

Reactions are grouped into three categories. In "Observed behavior" we report on what users did with the robot. In "Interview response" we cover the feedback they gave to the interviewer in lab testing. Finally, in "Operating the robot" we report on what operator actions were successful in engaging users.

Observed Behavior

Two project members analyzed the same videotape from the public and lab settings. When a participant was interacting with the robot, we made notes on the behaviors s/he exhibited. Patterns quickly emerged and we sorted people's responses into several behavioral categories. We also commented on times when users declined to interact. Behaviors were generally similar in both public and lab testing. Differences are noted.

General responses

Kids

Children were usually rapt with attention and treated the robot as if it were alive. Young children (4-7ish) tended to be very energetic around the robot (giddy, silly, etc.) and had responses that were usually similar regardless of gender. They were generally very kind to Sparky. Occasionally, a group of children might tease or provoke Sparky and we would then switch into a sad, nervous, or afraid state. This provoked an immediate empathetic response. We never moved to the point of using the angry emotional state with young children.

Older children (7ish to early teens) were also engaged but had different interaction patterns depending on gender.

Older boys were usually aggressive towards Sparky. Boys usually made ugly faces at the robot's face and did such things as covering the eyes, trapping it, pushing it backwards and engaging in verbal abuse. Switching the robot to a sad, nervous or fearful emotional state, in hopes of provoking a pity response, actually increased the abuse. However, when we moved to an angry emotional state and drove directly at the boys (the robot is not large or powerful enough to be dangerous), they were usually genuinely surprised and seemed to feel a newfound respect. They often made comments indicating these feelings to their friends and treated the robot much better after this episode.

Older girls were generally gentle with the robot. Girls often touched the robot, said soothing things to it, and were, on occasion, protective of the robot. If an older girl did provoke Sparky a little and it switched into a sad emotion, response was similar to that with young children, mentioned above.

It should be noted that though the responses for older boys and girls were stereotypical, exceptions were rare.

Nearly all the children in our lab interacted readily with the robot. In public, some children, mostly older teens, would pass by the robot. This was fairly uncommon.

Adults

Most adult interaction was collected in our lab. Kids quickly dominated most interactions at the Tech Museum.

Adults tended to treat the robot like an animal or a small child and generally gave the impression that they were dealing with a living creature. Compared to children, they were less engaged. Though children almost always moved down to the robot level, adults only did this sometimes. Gender wasn't a significant factor in determining adult responses. Response to Sparky's emotional palette was similar to young children and older girls

In the lab, most adults quickly began to play with the robot. Some however, were clearly unsure what to do. They seemed to want a purpose for the robot. Many of these people eventually began to experiment with the robot (see below).

Behaviors

Touching

Many subjects touched the robot, principally on the back, face and top of head. This behavior was more prevalent in young people, but was still remarkably common in adults as well. Once again, older children had responses that varied with gender. On average, boys were rougher, more likely to push it or cover its face. Girls tended to stroke and pet the robot more with one girl going so far as to kiss Sparky. Adult touching was more muted and not dependent on gender.

Speaking

Subjects talked to the robot quite a bit, often using "pet speech". They sometimes interpreted the robot for other people and "answered" the robot when it made vocalizations. They often heard the robot saying things that it hadn't and assumed that its speech was just poor, rather than by design. Users often asked questions of the robot. Particularly in our public testing, subjects would sometimes ask question after question, even though the robot did not usually respond with gesture or sound. The most common question was "what's your name?"

Mimicry

It was very common for subjects to mimic some portion of the robot's motion. For instance, if the robot moved its head up and down in a yes motion, subjects often copied the gesture in time with it. They also mimicked the extension and withdrawal of the head and would even follow it in circles.

Orientation Style

When a subject first engaged with the robot, s/he usually did so in one of two ways. The active subject stood in front of the robot and did something that might attract attention (made a face, waved, said something). The passive subject stood still until the robot acknowledged the subject's presence. Only then would the subject respond with motion. Essentially, the passive subject waited to be acknowledged by the robot, while the active subject actively courted a response.

Experimentation/Theorizing

Some subjects, mostly adults, spent time trying to understand the robot's capabilities better. For instance, subjects would snap their fingers to see if the robot would orient to the sound, or they would move their hands/bodies to see if the robot could follow them.

Other Behaviors

There were a number of behaviors that were observed once to a few times but were nevertheless interesting. Unless noted, all of these behaviors were observed in our public testing.

- Some people would pose around the robot for photographs.
- Parents would occasionally place toddlers directly in front of the robot and leave them.
- Children would often try to engage the robot in games.

- One couple in the lab took turns walking the robot around the space. They seemed to believe that the robot wanted to keep moving.
- A couple children became so attached to Sparky that they began screaming when their parents tried to take them away from the robot.

Interview Response

Formal subject feedback was collected in the lab testing. Participants were asked what they thought of the robot and what they liked and disliked. They were also asked specifically about the robot's motion and it's vocalizations. Overall, subjects liked interacting with the robot and used such adjectives as "fun", "neat", "cool", "interesting" and "wild". The responsiveness of the robot in its movement and emotions was cited as compelling. In particular, subjects often mentioned that they liked how the robot would track them around the room and even look into their eyes. Subjects commented that the robot reminded them of a pet or a young child and also mentioned enjoying the mimicking behavior. All of this feedback was not surprising given the observations we have reported in the previous section.

For some, primarily adults, motivation was a confusing issue. Though they typically could understand what the robot was expressing, subjects sometimes did not know *why* the robot acted a certain way. Adults were also more interested in the purpose of the robot.

Vocalization of the robot was not generally liked. The sounds fell between simple beeps and chirps (like R2D2) and actual human speech and this seemed to confuse people. We believe our efforts in this regard fell directly into an "uncanny valley". Because the vocalizations sounded so nearly like speech, subjects expected that they would be able to understand. Because the speech was muffled, with no linguistic content, their expectations went unmet.

Because this project was a wide ranging exploration, we didn't probe with very specific questions (e.g. "What was more compelling, the eyelids or the eyebrows?"). These kinds of questions would best be asked with future devices placed in more concrete, narrow contexts. Our aim for this work was to get a broad picture of how people responded to a robot that was social with them.

Operating the Robot

One of our project goals was to understand what new skills a social robot would need to learn. We therefore noted what our operators did in this regard.

Though it was not surprising, operators consistently got the best engagement by orienting the robot to the person. The robot's face pointed to the human's face and, moreover, we consistently found it valuable to look directly into the human's eye. This eye to eye contact almost always provoked an immediate response in the subject. Also, being able to read the basic affect of human faces allowed us to react to it.

Operators also found themselves having to deal with the robot's close proximity to many quickly moving humans. Users expected Sparky to not only not harm them, but also to know that they were there. For instance, if they touched Sparky somewhere, they expected it to know that and act accordingly (not move in that direction, turn its head to look at them, etc.). We ended up installing a wide field camera in a hidden location to allow the operator to see the environment directly around the robot. This gave Sparky the "body awareness" it needed.

Discussion and Conclusions

Users enjoyed interacting with Sparky and treated it as a living thing, usually a pet or young child. Kids were more engaged than adults and had responses that varied with gender and age. No one seemed to find the robot disturbing or inappropriate. We seemed, therefore, to have avoided any uncanny valleys (with the possible exception of vocalizations). It also got attention easily, fulfilling another of our design goals.

A friendly robot usually prompted subjects (older boys excepted) to touch the robot, mimic its motions and speak out loud to it. With the exception of older boys, a sad, nervous or afraid robot generally provoked a compassionate response. Older boys did not generally demonstrate this compassion but seemed to respect the robot only if it fought back when abused.

Our interactions with users also showed a potential need for future (autonomous) social robots to have a somewhat different sensory suite than current devices. For instance, we found it very helpful in creating a rich interaction to "sense" the location of bodies, faces and even individual eyes on users. We also found it helpful to read basic facial expressions, such as smiles and frowns. This argues for a more sophisticated vision system, one focused on dealing with people. Additionally, it was essential to be aware of participants' immediate proximity and to know where the robot was being touched. This may mean the development of a better artificial skin for robots. If possessed by an autonomous robot, the types of sensing listed above would support many of the behaviors (i.e. touching, mimicry, etc.) that users found so compelling when interacting with a teleoperated Sparky.

Fortunately, there are some traditional robotic skills that Sparky, if it were autonomous, might *not* need. For instance, there was no particular need for advanced mapping or navigation. Sparky's goal was to be social and that required paying primary attention to people rather than efficient navigation to a target. In fact, people usually came right to it; there was no need to seek them out. A robot that could pay attention to people in its field of view and had enough navigation to avoid bumping into objects would probably do quite well.

In addition to using somewhat different "senses" than traditional robots, Sparky demonstrated a different kind of intelligence than is traditionally associated with robots. Most standard research robots today use artificial intelligence to accomplish a clearly delineated set of manipulations in the world. The robot has a task and there is one correct answer. We see current and future social robots as using a form of artificial intelligence often referred to as "perceived intelligence" (Isbister 1995) which holds human experience as the metric for success. The robot need only do something "reasonable" to be a success. In short, social robots can succeed with a variety of actions and there are many correct answers for a given situation. Our experience bears this out; Sparky was often perceived as acting sensibly even when suffering a serious control malfunction that left him behaving erratically. Though it will be challenging to build these new social capabilities into mobile robots, humans are perhaps a more forgiving environment than roboticists are accustomed to.

Thoughts on Physical Interfaces to Computing

Our project began as an exploration of new computing interfaces and made a quick transition to building a small "cute" robot. Our results are necessarily rooted in that particular exercise. Nevertheless, watching and summarizing our experiences has left us thinking about how these lessons might be applied to broad interface questions.

As an "anti-workstation", Sparky ended up being an extremely physical system. Most of its communications took place with a dynamic, gesturing body (though sounds played a role also). Users responded with their own bodies, rather than with mouse clicks or key presses as in traditional interfaces to computation. This dialog, of human bodies relating to artificial ones, is a powerful one precisely because we already know body language. Designers could do well to use this "body-centric" type of communication in their work.

Sparky used its body to communicate such "strong" emotions as happy, sad and angry. These emotions would almost certainly look odd and forced in other, more typical, interface situations. However, there are more muted emotions that map readily to common states in machines. For instance, attentiveness was easily discerned just by watching the head of the robot as it followed a person's motion. This emotional state maps readily to a device's success at paying attention to a user. Similarly, Sparky could intuitively demonstrate a certain energy level using posture and the pace of its motion. This emotional state could have been related to the state of its batteries.

In short, by choosing emotional states reasonable to a given device and portraying that state in simple body language, interface designers may very well be able to open a relatively new channel for communicating with users. This "body-centric" technique, borrowed from robotics but potentially applicable broadly, should be a subject of further research.

Acknowledgements

We would like to thank several members of Interval's staff: Jesse Dorogusker, John Ananny, Paul Korff, Gerald Rogerson, Pat Engeman, Brad Niven, and Dan Psomas. Thanks also go to Debby Hindus for help with various drafts.

References

Amaya, K. and Bruderlin, A. 1996 Emotion from Motion. In *Graphics Interface*, 222-229. Toronto, Canada

Badler, N. Barsky, B. and Zeltzer, D. 1991. *Making Them Move: Mechanics, Control, and Animation of Articulated Figures*. San Mateo, CA: Morgan Kaufmann Publishers.

Blumberg, B. and Galyean, T. 1995 Multi-Level Direction of Autonomous Creatures for Real-Time Virtual Environments. In *Computer Graphics* 30(3), 47-54.

Breazeal, C. and Scassellati, B. 2000. Infant-like Social Interactions Between a Robot and a Human Caretaker. To appear in *Special issue of Adaptive Behavior on Simulation Models of Social Agents*.

Bruderlin A.& Williams, L. 1995. Motion Signal Processing. In *Computer Graphics*: 30(3).

Darwin, C. 1872. *The Expression of the Emotions in Man and Animals*. Oxford University Press.

Hamm, J. 1982 *Cartooning the Head and Figure.* Perigree.

Isbister, K. 1995. Perceived Intelligence and the Design of Computer Characters. M.A. thesis, Dept. of Communication, Stanford Univ.

Perlin, K. and Goldberg, A. 1996. Improv: A System for Scripting Interactive Actors in Virtual Worlds. In *Computer Graphics*: 29(3).

Perlin, K. 1985. An image synthesizer. In *Computer Graphics*: 19(3), 287-293.

Picard, R. 1997. *Affective Computing.* MIT Press.

Reichard, J. 1978. *Robots: Fact, Fiction and Prediction.* Penguin Books, 26-27.

Schiano, D.J., Ehrlich, S., Rahardja, K. & Sheridan, K. 2000. Face to InterFace:Facial affect in (hu)man and machine. To be published in *Proceedings of CHI2000,* NewYork: ACM.

Sims, K. 1994. Evolving Virtual Creatures. *In Proceedings of SIGGRAPH '94.* NewYork: ACM.

Snibbe, S., Scheeff, M. and Rahardja, K. 1999 A Layered Architecture for Lifelike Robotic Motion. *Proceedings of the 9th International Conference on Advanced Robotics.* Tokyo, Japan.

Thomas, F. and Johnston, O. 1981 *The Illusion of Life: Disney Animation*. Hyperion.