

Expressing Thought: Improving Robot Readability with Animation Principles

Leila Takayama
Willow Garage
68 Willow Road
Menlo Park, CA, USA, 94043
takayama@willowgarage.com

Doug Dooley
Pixar Animation Studios
1200 Park Avenue
Emeryville, CA, USA 94608
dooley@pixar.com

Wendy Ju
Willow Garage
68 Willow Road
Menlo Park, CA, USA, 94043
wendyju@willowgarage.com

ABSTRACT

The animation techniques of anticipation and reaction can help create robot behaviors that are human readable such that people can figure out what the robot is doing, reasonably predict what the robot will do next, and ultimately interact with the robot in an effective way. By showing forethought before action and expressing a reaction to the task outcome (success or failure), we prototyped a set of human-robot interaction behaviors. In a 2 (forethought vs. none: between) x 2 (reaction to outcome vs. none: between) x 2 (success vs. failure task outcome: within) experiment, we tested the influences of forethought and reaction upon people's perceptions of the robot and the robot's readability. In this online video prototype experiment ($N=273$), we have found support for the hypothesis that perceptions of robots are influenced by robots showing forethought, the task outcome (success or failure), and showing goal-oriented reactions to those task outcomes. Implications for theory and design are discussed.

Categories and Subject Descriptors

H.1.2. [Information Systems Applications]: Models and Principles—*User/Machine Systems*; H.5.2 [Information Systems Applications]: Information Interfaces and Presentation (e.g., HCI)—*User Interfaces*

General Terms

Human Factors, Design

Keywords

Human-robot interaction, Animation

1. INTRODUCTION

Our ability to interact smoothly with one another in everyday life depends in part on our ability to interpret each other's actions. Proper "performance" of actions can allow

bicycle riders to avoid pedestrians (as noted by [19]), let us clarify who we are speaking to [9], or help us signal our discomfort, for instance, when others come too close [1]. Robots that operate in public settings (e.g., offices, homes, airports, marketplaces) can be safer and more effective at performing work if they are designed with similarly human-readable behaviors [4]. By making the intended actions of robots more readily apparent to both interactants and bystanders, we can improve people's abilities to coordinate their actions with that of robots, much in the way that we coordinate our actions with one another every day.

One of the biggest challenges is to show when robots are busy "thinking" or planning to act. It makes pragmatic sense for robots to remain still while computing trajectories or planning movement, but the tendency for robots to jump suddenly from dead stop to action once the course of action is calculated can startle people, or catch them standing too close. Even robotics researchers occasionally mistake cogitating robots for idle robots and have narrowly missed being hit by darting arms or moving treads. The challenge, though, is that human expression of thought often relies on very subtle changes in facial and bodily expression. To make robot internal "thought processes" more readily observable, we turned to the practices of animators, who have a deep expertise in making inanimate objects come to life with readable actions. We focused specifically on pre- and post- action expressions of forethought and reaction as ways of helping people to understand when the robot is "thinking of acting." To test our hypotheses that these forethought and reaction cues would make robot actions more readable, we conducted a controlled experiment where we screened animated clips of a robot trying to accomplish a variety of tasks, with and without forethought or reaction, asking viewers to interpret and rate the clips.

1.1 Animation Principles

While both the believability of characters and the readability of character actions are important goals in animation, we primarily focus here on improving the readability of actions. Nearly 30 years ago, Frank Thomas and Ollie Johnston presented a set of animation principles from the Disney Studio [21]. Among these key principles were the principles of anticipation and follow-through, which prepare the audience for an action and help the animated action reach a realistic conclusion. Six years later, John Lasseter (Pixar Animation Studios) wrote an influential ACM paper about applying those animation techniques to 3-D computer animation, focusing on anticipation as a key way of making

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HRI'11, March 6–9, 2011, Lausanne, Switzerland.

Copyright 2011 ACM 978-1-4503-0561-7/11/03 ...\$10.00.

actions readable [15]. Lasseter addresses the importance of animating thought: “If a character were to move about in a series of unrelated actions, it would seem obvious that the animator was moving it, not the character itself. All the movements and actions of a character are the result of its thought process. In creating a ‘thinking character,’ the animator gives life to the character by connecting its actions with a thought process.” He describes how leading actions with the eyes or head helps to create anticipation, by establishing that the actions are thought through before being carried out. The pacing of the “thought” need to match that of the upcoming action: “Consider a character wanting to snatch some cheese from a mouse trap; the eyes will lead the snatch by quite a bit because this is a big decision. The character needs time to think, ‘...Hmm...This looks tricky, is this cheese really worth it or is it just processed American cheese food?...Oh what the heck...,’ he decides, and snatches the cheese.” Conversely, the body should move first when the thinking character is caught by surprise.

Animation principles have informed human-robot interaction research and development in numerous ways. Much of this work has focused on how animator’s principles for conveying character can be used to make robots more socially engaging and expressive. The keynote presentation for HRI 2009, for instance, highlighted how animation and character development could inform human-robot interaction design and research, discussing previous work in audio animatronic figures at Disney and two-way interactions with people via teleoperated robots [16]. More recently, a subset of animation principles [21] have been proposed for socially assistive robots such as Sparky, which might use caricatured behaviors to provide social skills training for children with autism spectrum disorder [17] and in robotic agents such as iCat [22], Tofu and Miso, which were designed to explore personality and expressiveness through principles such as secondary action and the illusion of thinking [24].

HRI designers have also adopted the tools and techniques used by animators to create coherent and believable actions. For instance, the designers of socially interactive robots such as iCat and Public Anemone [3] used animator techniques of dynamic weighting and blending to author the robot’s behaviors and responses to interactants. Animations themselves can provide a useful prototyping medium for testing how designed robot behaviors “read” without the investment involved with building and programming a physical robot; this lower cost enables designers to create interaction “sketches” that inform the design before engineering decisions are made that might constrain later design options [6]. Similar work [25] using video simulations of robots show high levels of agreement in how study participants behave in live and videotaped human robot interaction studies.

1.2 Human-readable actions

Anthropologists [9], sociologists [8], and psychologists [1, 7] have long studied how people use nonverbal communications in their interactions with one another. Nonverbal communications are different from other behaviors in that they have a specific communicative message for an implied audience; they might indicate to another that the “speaker” is engaged, or angry, or interested in interacting. For this reason, some have applied the term “emotive actuation” [10] to the use of such behaviors, in robotics.

One of the key features of nonverbal communications is

that they are non-focal; people are able to interact non-verbally with numerous others simultaneously. This makes such actions useful in coordinating joint action, for instance, with a musical partner. Shimon, the marimba-playing robot, uses the preparation and follow through to exhibit anticipatory gestures that enable musical ensembles to play together, which also enable the robot to play improvised jazz music with humans in synchrony [12]. Thus, the robot holding up a mallet, making visually obvious movements toward the next note, allows the human musician to synchronize his performance with the robot’s performance. This work is consistent with an embodied approach to human-robot interaction that perception and action are integrated, not separate; fluid joint action between people and robots requires that robot behave according to internal drivers and clearly communicate those drivers to people [11].

Other nonverbal communication techniques, such as gaze and orientation, can be used to signal an interactant’s willingness to engage without explicitly accosting people. Studies run with a RWI B21 robot with leg-tracking capabilities found that people were significantly more likely to stop and interact with the robot if it turned its screen towards people passing by [5]. A robot’s gaze behaviors can be used to influence the conversational roles that people take in an interaction, e.g., addressee, bystander, and overhearer; this influences how involved people feel in the conversation and how they feel toward the robot (a Robovie R2) [18].

Some social roboticists have been particularly interested in physical communication of emotion as a way to improve the interpretability of its actions. When emotionally expressive motions are generated by drives [2] (or goals), then those emotional expressions are playing a key functional role in demonstrating the robot’s progress toward its goals. In these cases, the displays of emotions are used as dynamic states, distinct from uses of emotion as indications of intrinsic traits such as personality or character.

It is critical in every case that the actions in question be human-readable; otherwise they would be wholly ineffective at coordinating joint action, engaging public interaction, or demonstrating robot intent. Thus, tests of robot readability can be useful to those who are designing non-verbal communication behaviors to support human-robot interaction.

2. STUDY DESIGN

The current study was designed to address the overall question of: Can (and if so, how should) human-robot interaction use principles from animation to make robot behaviors more readable and to influence people’s perceptions of robots?

More specifically, the current study explored these hypotheses:

(H1) Showing forethought (using anticipation [15, 21]) before performing an action will improve a robot’s readability as demonstrated through (a) written interpretations of the robot actions and sureness about those interpretations and (b) people’s subjective ratings of the robot.

(H2) Showing a goal-oriented response (reaction) to the task outcome (success or failure) will positively influence people’s subjective ratings of the robot regardless of whether the robot succeeded or failed on the task.

With these hypotheses in mind, we designed a controlled experiment that manipulated showing forethought (FT) or not (NFT) and showing an a reaction (R) to the task outcome (success or failure) or not (NR). We balanced successes and failures such that each person saw the robot succeed in half of the tasks and fail in the other half. Both forethought and reaction were between-respondents variables.

The dependent variables we measured were how people described the robot’s intentions (before the core action), how sure they were of their description, and how the respondents perceived the robot in terms of the following adjectives: appealing, approachable, competent, confident, intelligent, safe, superior.

2.1 Stimuli

The animations of the PR2 were created by our co-author, who is a professional character animator.

The mesh and kinematic model for this robot came from the PR2 (personal robot 2) platform unified robot description format (URDF), which was converted into a Collada file (an XML schema). The Collada model was imported into an animation software suite. For the purposes of this study, not all of the kinematic constraints were maintained, particularly the speed of the torso lift. Four tasks were selected in order to cover a variety of activities: opening a door, delivering a drink to a customer, ushering a person into a room, and requesting help from a person to plug into an outlet. These tasks were chosen from a set of tasks that an actual PR2 could perform and because they represented a range of tasks along two dimensions: (1) Doing a task for a person vs. doing a task for itself and (2) Doing an action-oriented task vs. doing a communication-oriented task. See Figure 1 to see a screen shot from each of the four task scenarios.




Figure 1: Four scenarios used in the current study: Delivering a drink, ushering a person, opening a door, and getting power

The behavior or performance of the robot was designed in two parts. The first part was the functional task (e.g., open the door). The second part was the set of expressions that communicated what the robot was planning to do (i.e., showing forethought) or the robot’s goal-oriented response to the task success or failure (i.e., showing reaction). As much as was possible, we modeled the control conditions (no forethought, no reaction) after the behavior exhibited by

the current PR2 Beta robot. See Figure 2 for more detailed depictions of two of those scenarios in terms of how they differed across experiment conditions.

2.1.1 Forethought and Reaction

In the no forethought condition, the animations of the robot depicted only the functional task part of the performance and nothing else, e.g., handing the drink to the customer. This was just the functional movement required by the robot to perform the functional task. The animations that showed forethought included the functional task part of the performance as well as the expressive part of the performance. The expressive part of the performance was also split into layers that showed engagement and confidence; it also used different timing to show the robot’s plans more clearly.

For the no reaction condition, the animated robot showed no reaction to task performance outcomes (success or failure), only the functional movements, e.g., releasing the drink. The animations that showed reaction to task performance outcomes also displayed goal-oriented expressive motions. By driving the expressive motions out of an internal motivation to succeed on each functional task, the PR2 was animated to show disappointment when it failed on the task and to show happiness when it succeeded on the task.

2.1.2 Engagement

To illustrate engagement, our animations had the robot just translate forward a very small distance, or translate back a very small distance. Engagement can also be shown by having a robot rotate its whole body forward or back from very near its bottom base. PR2 did not have this type of kinematic joint, so a translate motion was selected. The animations that showed forethought and reactions showed engagement with the person or the object of interest; the animations that did not show forethought or reactions did not use engagement.

2.1.3 Confidence

Confidence was demonstrated by having the torso translate up and down a very small amount. Upward torso lift showed confidence and positive thoughts. Downward torso movement showed a lack of confidence and negative thoughts. Together with timing, this element was critical to showing happiness about task success and disappointment at task failure.

2.1.4 Timing

The timing for the body language was crucial. To give the robot the appearance of thought, it had to be designed to show expressive movement just an instant before it performed the functional task, similar to the way that humans do. It also required the robot to look in the direction it was about to turn. This adds time to the performance, but allows the robot to appear as though it is thinking about the task before it performs it. This was used when showing happiness upon successful task completion (e.g., rolling away at a lively pace) vs. showing disappointment upon task failure (e.g., slumping off at a slow pace).

2.2 Functional vs. Expressive Motions

Although we have simplified the formulation of these animations into functional task motions (e.g., grabbing the door

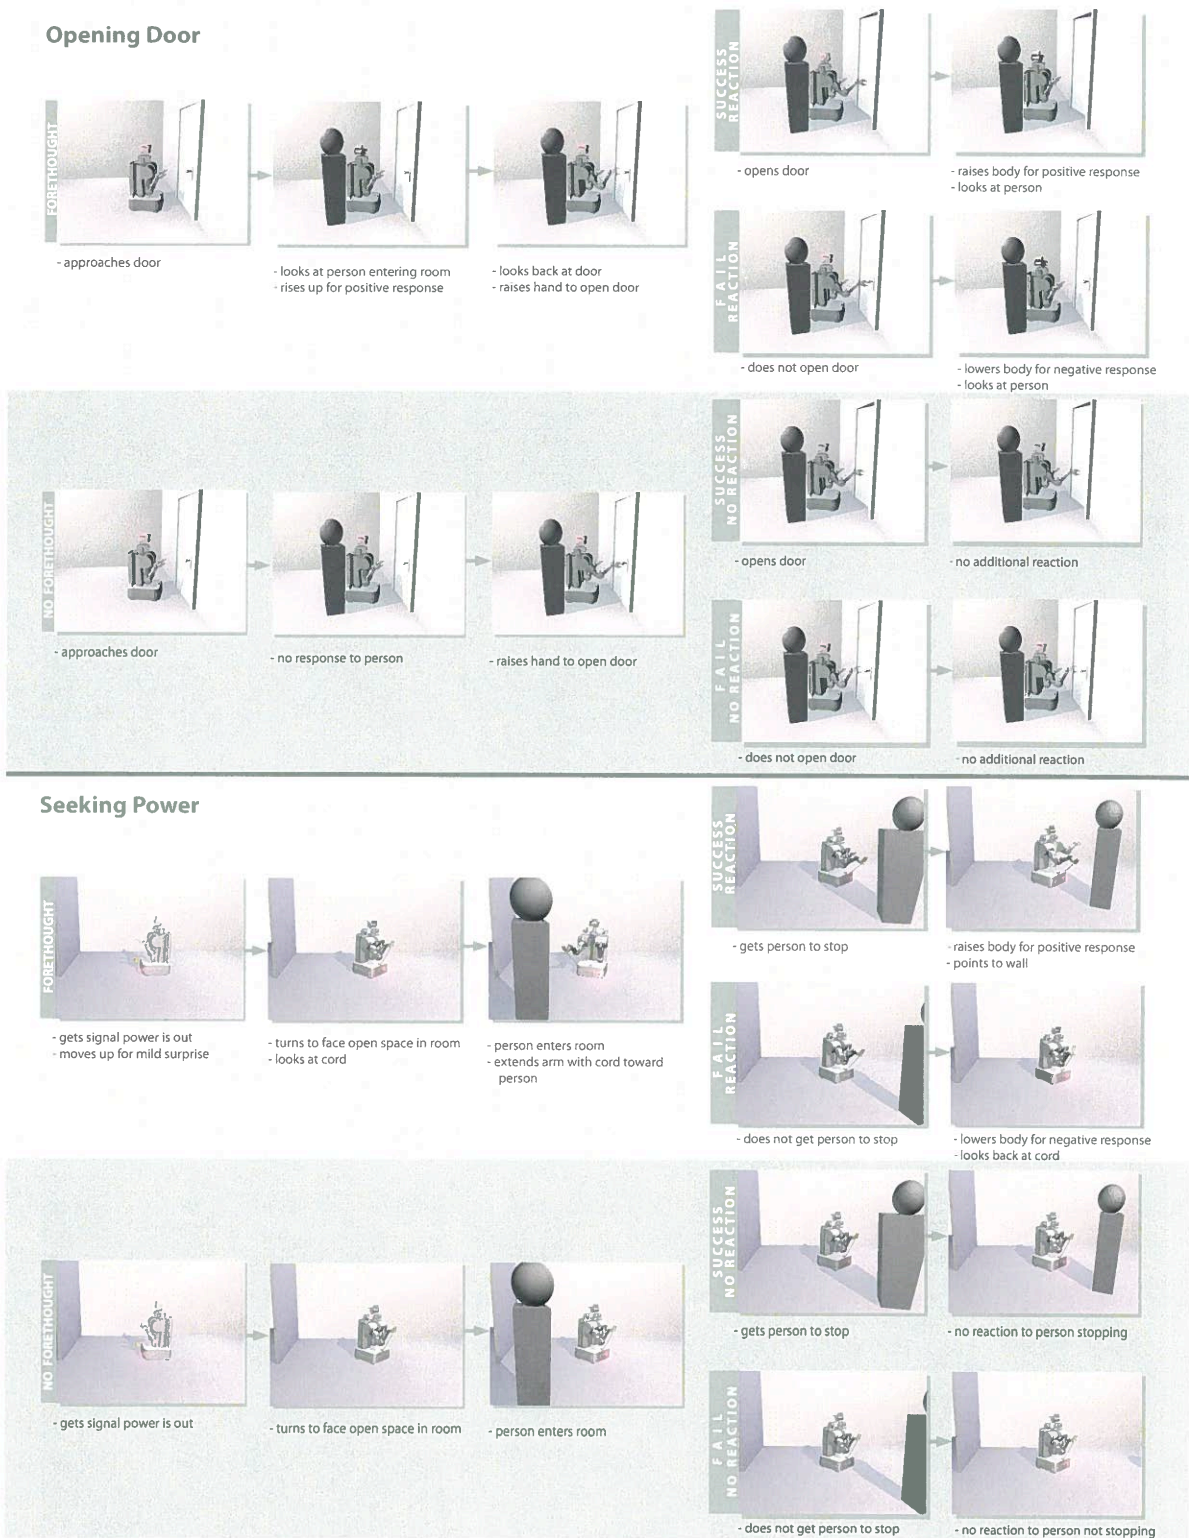


Figure 2: Frames from two of the animations used in this experiment: For each scenario, there were two possible pre-action animations (showing forethought or no forethought), two possible functional task outcomes (success or failure), and two possible post-action animations (showing a reaction to the task outcome or not).

knob) vs. expressive motions (e.g., looking around the door handle and scratching its head), we do not subscribe to the idea that these are completely separate concepts. There are many situations in which being expressive is a functional part of a task (e.g., when a robot requests help from a person to put its plug into an outlet socket). That is why two of the four scenarios used in this study involve communicative actions as part of the primary task and two of them are more solitary tasks.

2.3 Methods

In this 2 (forethought vs. none: between) x 2 (reaction to outcome vs. none: between) x 2 (success vs. failure outcome: within) study design, each respondent was exposed to only one of four possible experiment conditions—forethought + reaction, forethought + no reaction, no forethought + reaction, no forethought + no reaction. The four different tasks were randomized for order and the two successes and two failures were also randomized for order. An individual respondent saw either forethought in the animations or no forethought in the animations. That respondent saw either reactions to task outcomes or no reactions.

There is a precedent for video prototyping studies in human-robot interaction (e.g., [20, 25]) and human-computer interaction (e.g., [13]), although these precedents usually feature live video. Performing a video prototyping study with animations allowed us to provide consistent experiences for each of the study participants, to test a variety of task domains and to engage a geographically diverse set of study participants. In terms of design research, using animations allowed us to test the behaviors we would like to build before locking-in the design a robot would need to physically perform the behaviors. Finally, the animated video prototype allowed us to test embodied behaviors without subjecting study participants from any risk of bodily harm. Of course, there are differences between interacting with a robot and seeing someone else interact with a robot; there are also differences between interacting with a physical robot and merely seeing it on a screen [14]. However, we felt that a video prototype study was a reasonable approximation for live interaction for the reasons stated above.

2.3.1 Respondents

Two-hundred and seventy-three adult volunteer respondents were recruited through local mailing lists and the Mechanical Turk service provided by Amazon.com.

2.3.2 Study Protocol

Each respondent saw four pre-action animations (i.e., behaviors of the robot and person before the robot actually opened the door, delivered the drink, asked for help, or directed the person). On each web page, there was one animation and one set of these questions:

- 1) Please describe what you see happening in this clip.
- 2a) Please describe what you think the robot is trying to do in the video.
- 2b) How confident do you feel about your answer to question 2a?
- 3) If you were the person depicted with the robot in the clip, what would you do immediately after seeing the robot do what happened in the clip?
- 4) Please rate the robot in the video based on the

following parameters: Unappealing – Very appealing; Unintelligent – Very intelligent; Incompetent – Very competent; Subordinate to you – Superior to you; Very unsafe – Very safe; Not approachable – Very approachable; Not confident – Very confident.

See Figure 3 for a screen shot of an actual web page from the online questionnaire.

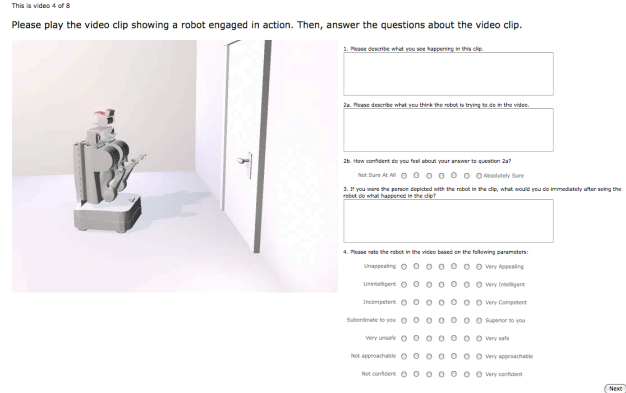


Figure 3: Sample page from the online questionnaire

After the four pre-action animations, the respondent saw the four post-action animations, which showed the task outcome (i.e., success or failure of the action) and the reaction or lack of reaction to the task outcome. The same set of questions was asked for each of the four post-action animations.

Upon completing the eight pages of animation clips and their accompanying questionnaires, the respondent was shown a debriefing paragraph about the purpose and goals of the study. The respondent was provided with contact information for the experimenters in case it was necessary to discuss questions or concerns about the study.

3. DATA ANALYSIS

Readability was measured in two ways. First, the open-ended responses about what the respondent thought the robot was about to do (questionnaire item 2a) were coded in terms of the presence of the top two keywords—verbs chosen as being most descriptive of what the robot was trying to do. For the door animation, the keywords were “open” and “enter.” For the drink animation, the keywords were “serve” and “deliver.” For the ushering animation, the keywords were “direct” and “greet.” For the power animation, the keywords were “plug” and “charge.” If the open-ended description of the robot’s intention included at least one of the appropriate keywords, then it was marked as being correct (1). Otherwise, it was marked as being incorrect (0). This is a Boolean classification based on the presence of keywords, which is a strategy used in automatic survey classifiers [23]. Second, respondents rated how sure they felt about their readings of the robot’s behavior (1=not sure at all, 7=absolutely sure).

The adjective rating data (questionnaire item 4) were analyzed using a repeated measures analysis of variance (ANOVA) with a Bonferroni correction for the cut-off p -value of .0071

(= .05 / 7) because there were seven adjectives in the list. The ratings for the adjectives were treated as dependent variables that were scored on a 1-7 scale.

First, the pre-action animation responses were analyzed with regard to the influence of showing forethought as a between-respondents variable (ANOVA with an error that accounted for repeated measures of each respondent). Second, the post-action animation responses were analyzed with regard to the influence of goal-oriented response as a between-respondents variable and the task outcome (success vs. failure) as a within-respondents variable (repeated measures ANOVA).

4. RESULTS

4.1 Readability of Robot Forethought

Forethought was not found to be a significant predictor of keyword match in the participant’s descriptions of robot activity, $F(1,1083) = 0.82, p = .37$ (n.s.). Variation between these particular animations is a potential cause for this non-significant result. The activity types were found to be significant predictors of keyword match: activity for person vs. robot, $F(1,1083) = 40.51, p < .001$; communicative vs. action oriented activity, $F(1,1083) = 37.90, p < .001$; and their interaction, $F(1,1083) = 25.18, p < .001$. Upon reviewing their descriptives (e.g., means and standard errors), we found large differences in keyword matching rates of each of the animations (1 = correct and 0 = incorrect)—opening door ($M = 0.84, SE = 0.02$), requesting power ($M = 0.78, SE = 0.03$), delivering drink ($M = 0.64, SE = 0.03$), and ushering ($M = 0.29, SE = 0.03$). See Figure 4.

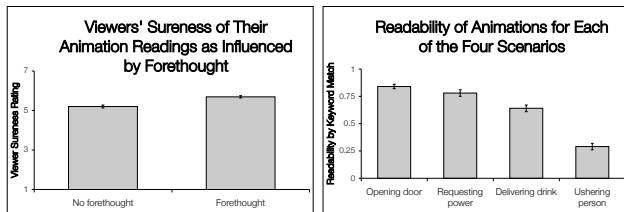


Figure 4: Means and SEs for how sure people were about their readings of the robot behaviors and keyword matching rates across the four scenarios

Participants in the forethought condition were significantly more sure of their descriptions ($M = 5.69, SE = 0.06$) than in the condition where the robot did not show forethought ($M = 5.20, SE = 0.07$), $F(1,1067) = 26.43, p < .0001$. This supports part of Hypothesis 1a in that forethought improved the confidence with which people interpreted the robot behaviors.

4.2 Perceptions of Robot Forethought

Showing forethought made the robot seem more appealing ($M = 4.83, SE = 0.06$) than when the robot did not show forethought ($M = 4.27, SE = 0.07$), $F(1,1078) = 34.18, p < .001$. Showing forethought also made the robot seem to be more approachable ($M = 5.05, SE = 0.06$) than when it did not show forethought ($M = 4.54, SE = 0.07$), $F(1,1072) = 27.05, p < .001$. Perceptions of robot appeal and approach-

ability were moderately correlated, $r = .61, p < .001$. Their scales ranged from 1 to 7. See Figure 5.

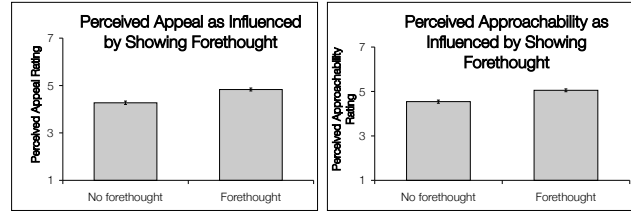


Figure 5: Means and SEs for perceived appeal and approachability of the robot as influenced by showing or not showing forethought

Perceptions of appeal (Pearson $r = .42, p < .001$) and approachability (Pearson $r = .31, p < .001$) were weakly correlated with the respondents’ sureness of the descriptions. Showing forethought was not found to have significant influence upon other perceptions of the robot. These results present support for Hypothesis 1b that showing forethought would improve people’s subjective ratings of the robot.

4.3 Perceptions of Robot Reaction

Because the competence and intelligence ratings were strongly correlated, $r = .83, p < .001$, we constructed an unweighted average of the two items to form a single factor that we could generalize as seeming to be “smart” (competent and intelligent). As expected, succeeding on a task made the robot seem to be smarter ($M = 4.74, SE = 0.07$) than when it failed ($M = 3.86, SE = 0.07$), $F(1,1063) = 28.75, p < .001$. The robot also seemed to be more confident when it succeeded on the task ($M = 4.68, SE = 0.07$) than when it failed ($M = 3.96, SE = 0.08$), $F(1,1055) = 32.26, p < .001$. These were main effects in the repeated measures ANOVAs described above. These same main effects were statistically significant at the .0001 level when analyzing the “competent” and “intelligent” items separately. See Figure 6.

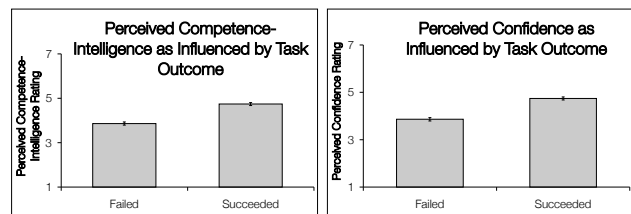


Figure 6: Means and SEs for competence-intelligence factor and confidence as influenced by task outcome

Showing a goal-oriented reaction to a task outcome (i.e., disappointment in response to failure and happiness in response to success) also made the robot seem to be smarter ($M = 4.72, SE = 0.07$) than when it did not react ($M = 3.86, SE = 0.07$), accounting for whether or not the robot succeeded or failed on the task, $F(1,1063) = 18.40, p < .001$. See Figure 7. These results present support for Hypothesis 2 that showing a goal-oriented response to the task outcome

positively influenced people’s subjective perceptions of the robot, regardless of whether it completed the task successfully.

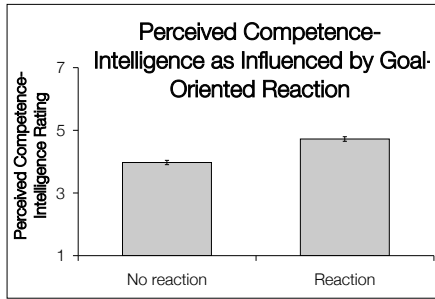


Figure 7: Means and SEs for competence-intelligence as influenced by showing a reaction to the task outcome

Succeeding and reacting were not found to have significant influences upon other perceptions of the robot.

5. DISCUSSION

We found mixed support for Hypothesis 1a (that showing forethought would improve a robot’s readability); while we did not find support for H1a in terms of the keyword match, we did find support for H1a in terms of how sure people felt about their readings of the robot’s intentions.

We also found support Hypothesis 1b (that showing forethought would improve people’s perceptions of the robot). Hypothesis 1b was supported in terms of appeal and approachability.

Furthermore, Hypothesis 2 (that showing reaction to the task outcome would improve perceptions of the robot) was supported in terms of the perceived “smartness” (a combination of “intelligence” and “competence”) and confidence of the robot.

In terms of keyword matches, the readability of the robot’s behaviors varied more between tasks than they did between showing forethought vs. not showing forethought (H1a). Because each animation was staged very clearly, it is possible that viewers were better able to focus on the specifics of the environment to tell what the robot was trying to do. See Figure 4. For example, the door opening animation only included a robot, a door, and a person. The presence of the door, which is not present in the other animations, provides the viewer with a strong cue that something is going to happen with the door. To address this issue, future work could use actual robots in a more realistic physical environment that is held constant across tasks.

The improvement of appeal and approachability of the robot that showed forethought (H1b) and improvement in perceived intelligence/competence of the robot that showed reaction to its success/failure (H2) indicate that these two lessons learned from animation might be used to improve the perceptions of robots in human-robot interactions.

5.1 Implications for Theory and Design

At a high level, this study demonstrates the potential benefits of using animation principles to influence people’s per-

ceptions of robots in human-robot interactions. In practice, the executive level processing of an autonomous robotic system can identify upcoming tasks and trigger a task-oriented expressive behavior that shows forethought immediately before performing a functional task to help increase people’s anticipation for the task that is to come. If the robot shows forethought before performing a functional action, people will be more likely to see the robot as being more appealing, approachable, and sure of its subsequent actions. This can be important for human-robot interactions, particularly ones that require interactions in close proximity.

The executive system can also identify the status of each of a robot’s goals to trigger reaction behaviors consistent with the success or failure to complete a task. Showing a goal-oriented reaction to a task outcome can positively influence people’s perceptions of the autonomous robot. This can hold true even in the face of situations where the robot fails to perform functional tasks. While it is clearly better if the robot succeeds in performing functional tasks, it is possible to mitigate the negative perceptions of the robot by having it acknowledge its failures. When the robot successfully completes a functional task, it can also benefit from expressing happiness in relation to the successful completion of the task.

5.2 Limitations and Future Work

There are several limitations of the current work that inform the next steps in this research. First, as with any video prototype study, there are limitations to the methodology, particularly that people do not necessarily respond in the same way to videos of robots as they do to physical robots [14]. We would like to check whether people’s reactions to these robot behaviors will be correlated but stronger in the live condition as suggested by [25] and [14]. Although we are currently able to base physical robot behavior on animations from Maya on the PR2, the setup is as yet somewhat risky to study participants because it does not yet incorporate collision avoidance. It is important for future work to test these techniques in the animation of physical robots.

Second, the current study is that it only investigates the influences of forethought and reaction upon human-robot interactions across four contexts—opening a door, delivering a drink, requesting help with recharging, and ushering a person in a particular direction. We would like to try animation techniques to illustrate more variations upon anticipation so as to suggest intent, cogitation, hesitation, enthusiasm (for example), and also to explore these techniques in a wider array of scenarios and task domains where service-oriented robots might be found.

Third, there is extensive work to be done in teasing out the exact dimensions these animation principles that most improve the readability of robot actions; in the current study, we used a combination of dimensions rather than isolated ones because we wanted to test the overall approach before isolating dimensions.

Fourth, we have learned that it would be better to allow viewers to watch each animation only once (not repeatedly) and to see how quickly and with what certainty viewers read the behaviors. This is more consistent with current practices in animation for measuring readability of motion pictures and will be used in future work. Future work could use existing methods of screening animations for testing the readability of robot behaviors.

6. CONCLUSIONS

This study makes several contributions to the realm of human-robot interaction. First, we have demonstrated the use of robot animations as a way of testing rapid prototypes of interactive robot behavior prior to building and programming physical robots. Secondly, we have demonstrated several techniques for improving a robot's readability for the often inscrutable modes where the robots are planning or evaluating goal-oriented action.

Drawing from principles of animation, we were able to illustrate forethought and reaction. By employing techniques of engagement, confidence and timing, we are able to help people read robot behaviors with more certainty. In this on-line video prototype experiment ($N=273$), we found support for the hypothesis that perceptions of robots are positively influenced by robots showing forethought, the task outcome (success or failure), and showing goal-oriented reactions to those task outcomes.

We found that showing forethought makes people more sure of their interpretations of robot behavior, and make the robot seem more appealing and approachable. We also discovered showing a reaction to the task outcome can make the robot seem to be more intelligent/capable, even if it fails to achieve the functional task at hand.

While animation principles for showing motion have been written down [15][21], there are other unwritten animation principles and techniques, including individual styles, that also deal with showing *thought*. Because robots already “show” motion, there may be more to gain from exploring ways to show robot thought. The success of such animated behaviors on robots could be evaluated with regard to how quickly people can “read” the robot, accurately predicting what the robot will do next, and thereby influencing the fluency of human-robot interactions.

Moving forward, we believe that the results of this study are a testament to the utility of applying animation techniques to the craft of robot behavior design. The ability to improve people's interpretations of what the robots are doing and whether or not the task is successfully completed increases people's confidence and willingness to engage in interaction with these robots, and might help to make these robots more safe to interact with as well. We aim to expand the repertoire of animated robot behaviors and application domains to better take advantage of these capabilities.

7. ACKNOWLEDGMENTS

Thanks go to all of the volunteer participants in this study. Thanks also go to John Hsu and Rosen Diankov for generating the PR2 model that we used in this study.

8. REFERENCES

- [1] M. Argyle. *Bodily communication*. Taylor & Francis, 1988.
- [2] C. Breazeal. A motivation system for regulation human-robot interaction. In *AAAI*, pages 54–61, 1998.
- [3] C. Breazeal, A. Brooks, J. Gray, M. Hancher, J. McBean, D. Stiehl, and J. Strickon. Interactive robot theater. *Comm. of the ACM*, 46:76–85, 2003.
- [4] C. Breazeal, C. D. Kidd, A. L. Thomaz, G. Hoffman, and M. Berlin. Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *IROS*, pages 708–713, 2005.
- [5] A. Bruce, I. Nourbakhsh, and Y. R. Simmons. The role of expressiveness and attention in human-robot interaction. In *ICRA*, pages 4138–4142, 2002.
- [6] B. Buxton. *Sketching user experiences: Getting the design right and the right design*. Morgan Kaufmann, 2007.
- [7] H. Clark. Using language. *Computational Linguistics*, 23(4), 1996.
- [8] E. Goffman. *The presentation of self in everyday life*. Harmondsworth, 1978.
- [9] E. Hall. *The hidden dimension*. Doubleday, Garden City, NY, 1966.
- [10] J. Harris and E. Sharlin. Exploring emotive actuation and its role in human-robot interaction. In *HRI*, pages 95–96, 2010.
- [11] G. Hoffman and C. Breazeal. Robotic partners' bodies and minds: An embodied approach to fluid human-robot collaboration. In *Cognitive Robotics*, 2006.
- [12] G. Hoffman and G. Weinberg. Shimon: An interactive improvisational robotic marimba player. In *CHI*, pages 3097–3102, 2010.
- [13] W. Ju and L. Takayama. Approachability: How people interpret automatic door movement as gesture. *International Journal of Design*, 3(2), 2009.
- [14] S. Kiesler, A. Powers, S. R. Fussell, and C. Torrey. Anthropomorphic interactions with a robot and robot-like agent. *Social Cognition*, 2:169–181, 26.
- [15] J. Lasseter. Principles of traditional animation applied to 3d computer animation. In *SIGGRAPH*, pages 35–44, 1987.
- [16] A. J. Madhani. Bringing physical characters to life. In *HRI*, pages 1–2, 2009.
- [17] R. Mead and M. J. Mataric. Automated caricature of robot expressions in socially assisted human-robot interaction. In *HRI workshop*, 2010.
- [18] B. Mutlu, T. Shiwa, T. Kanda, H. Ishiguro, and N. Hagita. Footing in human-robot conversations: How robots might shape participant roles using gaze cues. In *HRI*, pages 61–68, 2010.
- [19] D. Norman. *The design of future things*. Basic Books, 2007.
- [20] D. S. Syrdal, N. Otero, and K. Dautenhahn. Video prototyping in human-robot interaction: Results from a qualitative study. In *HRI*, pages 1–8, 2008.
- [21] F. Thomas and O. Johnston. *The Illusion of Life: Disney Animation*. Hyperion, 1981.
- [22] A. Van Breemen, P. Res, and N. Eindhoven. Animation engine for believable interactive user-interface robots. In *IROS*, pages 2873–2878, 2004.
- [23] P. Viechnicki. *A performance evaluation of automatic survey classifiers*, pages 244–256. Springer-Verlag, 1998.
- [24] R. Wistort. Only robots on the inside. *Interactions*, 17(2):72–74, 2010.
- [25] S. Woods, M. L. Walters, K. L. Koay, and K. Dautenhahn. Comparing human robot interaction scenarios using live and video based methods. In *Workshop on Advanced Motion Control*, pages 27–29, 2006.