The 24th IEEE International Symposium on Robot and
Human Interactive Communication
August 21-September 4, 2015. Kobe, Japan.

Taking Candy from a Robot: Speed Features and Candy Accessibility Predict Human Response

Heather Knight, Manuela Veloso, Reid Simmons, *Members, IEEE*

***Abstract*— In our experiment, two autonomously moving costumed robots visit 256 offices during a 'reverse' trick-or-treating task close to Halloween. Our behavioral data supports the idea that people interpret a robot's non-verbal cues, as the robots' costuming and baskets of candy seem to have communicated an implicit offer of candy. In fact, one third of our detection instances occurred during robot transit, i.e., while the robots were making no verbal offer. We find that candy accessibility dominates any social influence of robot orientation and that robot speed influences both whether people will interrupt a robot in transit (slow more interruptible) and whether they will respond to its verbal offer (fast more salient).**

I. Introduction

Our research interest is how robot motion and physical cues, such as orientation to a goal, influence human response to a mobile robot. We are also curious how these cues impact people's attributions about robots, and whether they can clarify or confound human interpretation of a robot's goals. As a naturalistic exploration of this topic, we implement a candy-delivery behavior on two autonomous 1.3-meter-tall CoBot robots [1] with four-wheeled omnidirectional bases and baskets for transporting objects (Fig. 1). We also added down-facing sensors that we use to detect when people are taking candy.

This is the third year that the CoBots have performed a reverse trick-or-treating behavior close to Halloween, but the first time that we varied its motion characteristics experimentally. During the task, the robots travel to all the offices on each floor, verbally offering people to take candy upon its arrival to each. Deployed since 2009 to perform a variety of tasks, they came pre-installed with baskets (either in the front or the back of the robot) that we filled with candy for the experiment. The robots are capable of rotation and translation, thus they are ideal for exploring the impact of simple motions on human perceptions and behavior.

Our variables included robot orientation to the office (*social* or *asocial*) and robot navigation speed (*fast* or *slow).* Our hypotheses were: 1) that people at the offices would take candy more often from robots with natural social facing (i.e., looking directly into the offices, instead of away), and 2) that the robots moving at *fast* speed would seem more dedicated to their tasks. Surprisingly, we find that a robot's social facing has no impact on whether or how fast people take candy. Instead, people are most influenced by the proximity of the candy-basket. Consistent with our second hypothesis,

*Research supported by the National Science Foundation.

H. Knight, M. Veloso and R. Simmons are with the Robotics Institute at Carnegie Mellon University, Pittsburgh, PA 15203 (phone: 412-268-3818; fax: 412-268-6436; e-mails: hknight, mmv, reids@cs.cmu.edu).

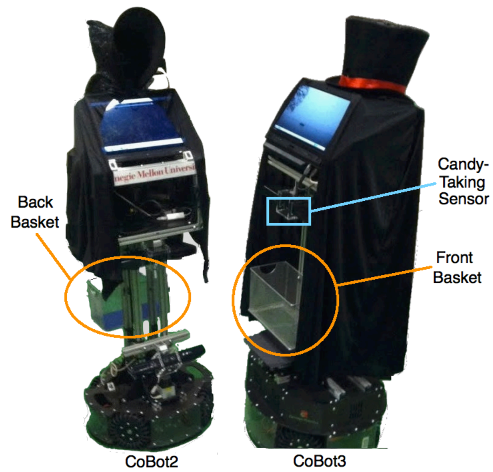

Figure 1. *The robots: CoBot2 costumed as a witch with a back-installed candy basket, CoBot3 costumed as a magician with front-installed candy basket. Candy-taking detectors installed above the baskets facing down.*

people took candy more often when a robot approached an office at higher speed.

We were also surprised to discover that one-third of the candy-taking instances occurred while the robots were in transit, in other words, when they were not verbally offering candy. While the orientation condition and basket location do not impact the transit path, we did find that hallway candy-taking behavior did change in regard to the robot's *fast* or *slow* speed condition. In fact, people take candy over twice as often from a robot traveling at *slow* speed, even when accounting for relative time spent traveling.

II. Related Work

Autonomous robots are increasingly entering human environments, from shopping center guides [2] to café companions [3], from delivery robots in our hospitals [4] to university resources [5][6], and cleaners of our homes [7]. Whether sharing common spaces or engaging residents and bystanders directly, these robots often benefit from a shared understanding of social rules [3][6], and when these rules *do* or *do not* apply to particular robots in particular task contexts.

As the examples above highlight, more and more researchers are bringing their research studies out of the laboratory. While the challenge of such investigations is that natural environments are complex and unexpected confounds may enter the experiment (e.g. how many people took candy from the CoBot robots in the hallway during transit), they also present realistic versions of the target environments for which we design social robots. The unique contributions of our work include the real-world autonomous setting in which we conducted our experiment, and our specific insights about expressing robot body language via motion.

While previous research has investigated aspects of robot motion communications, such as proxemics [8], social navigation [9], and human attributions about combinations of robot motion characteristics [10], our work explores simple but generalizable principles of how robot motion and orientation impacts human behavior during a naturalistic robot task. For example, in the author's previous work applying Laban Effort principles to robot motion [11], orientation was a clear feature used by study participants to interpret the agent's expression. In this experiment, however, we found that though people made note of a robot's social orientation, that orientation did not impact human behavioral response. We will explore these diverging findings further in our results and discussion sections.

III. Background: The CoBots

Since 2009, the CoBot robots have performed a variety of tasks throughout our robotics and computer science buildings, including meeting and guiding someone to a location, delivering objects and/or messages [12], and serving as semi-autonomous tele-presence agents [13]. Just over 1.3 meters tall, the CoBot robots move via their four-wheeled omni-directional bases. Although researchers can monitor robot operations remotely via localization and local video data, the robots seldom require human intervention.

The CoBots autonomously complete their tasks, using depth cameras and laser rangefinders for localization and navigation [14]. When necessary, they can seek human assistance via symbiotic autonomy [6], e.g., asking bystanders for help pressing the elevator button. An on-board tablet provides the computational platform to run the robot's task behaviors, as well as the algorithms for sensing and control. It also provides a graphical and speech-based interface for interaction.

IV. Autonomous Candy Delivery

In this experiment, we used two of the four existing CoBot robots: CoBot2 and CoBot3. While the hardware for the two robots is mostly identical, because of previous work, their basket locations differ (Fig. 1), namely Cobot2 has a back-installed basket, while Cobot3 has a basket installed in the front. We also added customized Halloween costumes, both consisting of a cape and a hat, but differing in theme and cape length so as to not interfere with candy access.

In this section, we overview the reverse trick-or-treating behavior that the two robots autonomously performed, then describe our setup to detect when people were taking candy from the robots.

A. Reverse Trick-Or-Treating Program

For the third year in a row, researchers sent CoBot robots to deliver candy to people's offices in costume. This year was the first year that more than one CoBot was in operation at the same time, and it was the first year that we varied robot motion conditions such that we could analyze the impact of our experimental features on people's candy-taking behavior.

The reverse trick-or-treating task can be called with any floor currently included in the CoBot's navigation system, automatically loading a desired sequence of office locations, using ascending numerical order. At each step, the robot sets its next target office location, loads its experimental conditions (randomly choosing between *fast* vs. *slow* speed, and *social* vs. *asocial* orientation), then attempts to navigate to that location at the appropriate speed.

Upon arrival to the office, the robot centers itself at the door, sets its orientation angle to the current setting, and says "Knock, knock, is anyone there?" and/or "Happy Halloween! Please take some candy, then press the done button." The robot's verbalization also prints out on its tablet touch screen. If there is no touchscreen response, the robot will wait for up to 30 seconds then leave the office, otherwise it leaves when the user presses the "Done" button. Upon departure, it loads its next office destination and experimental settings, the process continuing until it has visited the full set of offices for that floor, also returning to offices with no response.

During operation, the CoBot data logging [1] relevant to this experiment included: sequences of offices visited, robot transit times, robot task durations, experimental condition settings, and candy-taking detection distance data (including raw signal, estimated distance and timestamps). While the robot operations were autonomous, there was often a researcher on site down the hallway who could observe and/or overhear people's reactions to or comments about the robots. Experimenters only intervened if the robot was stuck, needed more candy or, on some floors, required a door to be opened to enter a corridor.

B. Detection Methodology

Each CoBot was equipped with Sharp IR distance sensors (range: 10cm-80cm), which we installed above each candy basket. Alternative sensors might include a weight sensor (would be better for running average, but less reliable for candy-taking) or a vision system tracking hands (high algorithmic load, less privacy). Though the IR sensors are simple, they provide complex insights into human behavior, while allowing users to remain anonymous.

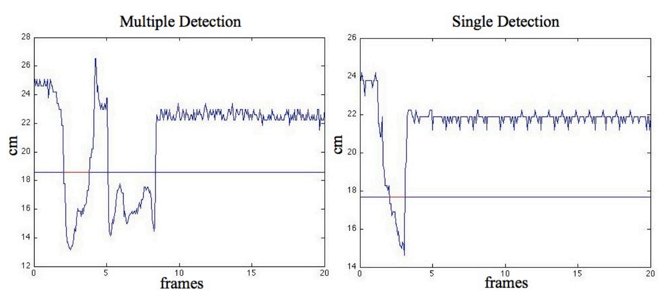

Figure 2. *Sample detections of people taking candy, (left) two detections, (right) one detection. Straight line presents threshold for leftmost detection.*

The candy-taking sensor publishes ~every 50ms (a rate of 20 frames/sec). Using this data, we 1) estimate candy levels (running average over one minute of data), and 2) detect when people put their hand in the basket to take candy. To accomplish the second, our system looks for signals lasting at lest 150ms that exhibit a minimum of 7cm deviation from running average (Fig. 2). We parameterized these numbers to eliminate false positives (as determined by human labels). In other words, we have a conservative detector, in which some candy-taking instances may be missed but there are few false-positives (final rate was 2% or 3/151 detections, by human annotation).

As part of the detection process, we group detections into a single *instance* if there is a gap of <500ms between them (threshold determined from human behavior samples). If the detections are separated by more than that, they are considered multiple detections, either from one person taking candy several times, or more than one person taking candy. We also note the duration of each detection *instance* in *frames* (20 frames/sec), as longer detection durations may indicate several people taking candy in overlapping intervals, and/or a more relaxed manner of a single person taking candy. We display sample graphs of our detections in Fig 2.

V. EXPERIMENT

In our analysis, we intended to analyze the impact of our experimental conditions on people's candy-taking behavior during two task contexts: *at office* or *in transit*. When the robots were at an office, they would verbally offer the occupants candy, but when they were in transit, they made no verbal offer. In the following subsections, we review the experimental conditions, then overview our data collection.

A. Experimental Conditions

Our experimental conditions include orientation (*social* indicates orienting toward office or *asocial* indicates orienting away) and speed (*slow* or *fast)*. The experimental conditions were selected at random each time an office location was added to a robot's task planner.

- **Orientation:** When a robot arrived at an office, it would either face directly into the office at 0 degrees (*social,* i.e., toward interaction partner) or directly away from the office at 180 degrees (*asocial*), then ask the occupants if they would like some candy. The *social* and *asocial* labels are meant to be category abstractions rather than dictionary definitions.

- **Speed:** The robots either traveled at 0.3 m/s (the *slow* condition) or 0.75 m/s (the *fast* condition), this speed setting was relevant to both its transit between offices speed and its approach speed to an office. Motor noise increases with speed.

We also track candy location. Each robot had a permanent costume and candy basket. In our analysis, we will refer to candy location as *front basket* and *back basket* (Fig 3). There is a possibility that the costuming impacted human response, but we tried to make them fairly similar. Cobot2 wore a short cape and witch hat, with a back-installed candy basket, while Cobot3 wore a long cape and magician hat with a front-installed basket (Fig. 1).

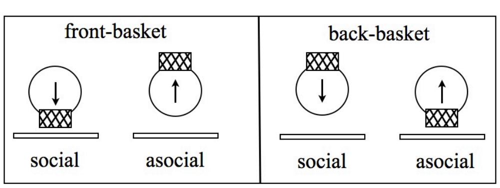

Figure 3. *Full set of orientations (*social *is toward or* asocial *is away from office) and candy-basket permutations.*

B. Data Collection Overview

The two robots successfully visited 256 offices (302 paths attempted, extra paths due to duplicated attempts at reaching the same office after localization corrections or lack of response) across four floors of our computer science and robotics buildings (Table 1). The robots traveled a combined distance of ~2.3km and time of 4 hours, 34 minutes and 24 seconds (Fig. 3). In total, they detected 148 instances of people taking candy, dispersing over 27 lbs. of candy.

TABLE I. REVERSE TRICK-OR-TREATINGS STATS

Floor	Duration*	Distance**	#Offices	# Detects
GHC6	5871	967	99	30
GHC9	5114	584	100	30
NSH4	3042	427	55	33
GHC7	2437	310	48	55
	16464	2288	302	148

*Duration in seconds including startup time, **Distance in meters including travel from other floors

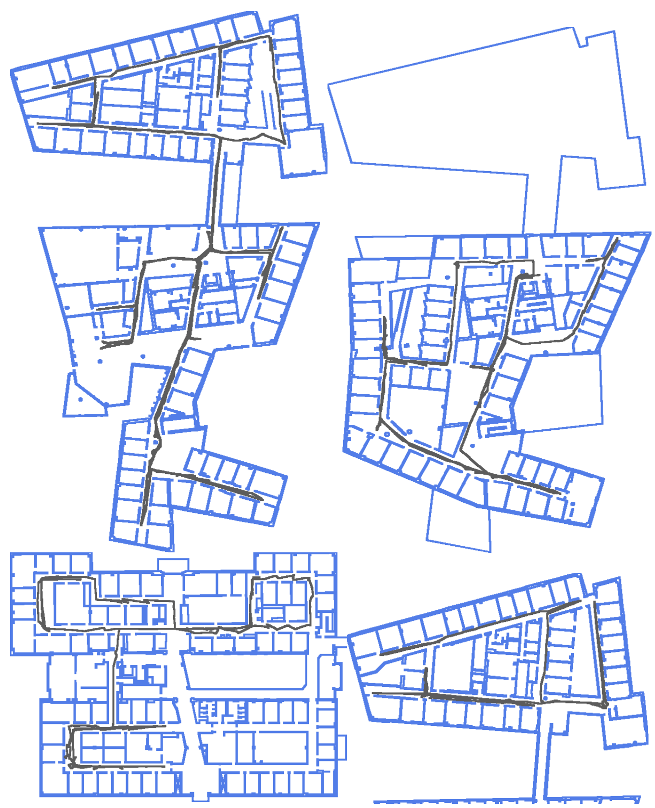

Figure 4. *Birds-eye view of the four floors in which Cobot2 and Cobot3 visited offices to offer people candy, overlayed with their transit paths. Top row: GHC6 and GHC 9 (Cobot2), Bottom row: NSH4 and GHC7 (Cobot3)*

A further breakdown of the detection contexts by floor in Table I help suggest the social composition of each, e.g., how many people took candy at the offices may reflect how many people were in their offices to begin with, and the number of transit detections may relate to the number of people were passing through the hallways. Logistically, it made sense to assign particular robots to particular floors, but population differences in the floors visited and differing times of days of those visits may have also impacted the total number of candy-detection instances we collected. As annotated in Fig. 4, Cobot2 visited GHC6 and GHC9, while Cobot3 visited NSH4 and GHC7.

C. Representative Detection Signal

In Fig. 5, we present an example candy-taking-detection signal from the entire set of GHC9 data. The candy level varies as candy gets taken, shifts or is added. The other three files have similar detection distributions. There is slightly more ambient noise in the Cobot3 (+/-3cm) than Cobot2 signal (+/-2cm), perhaps due to vibration, though both are well under the detection threshold.

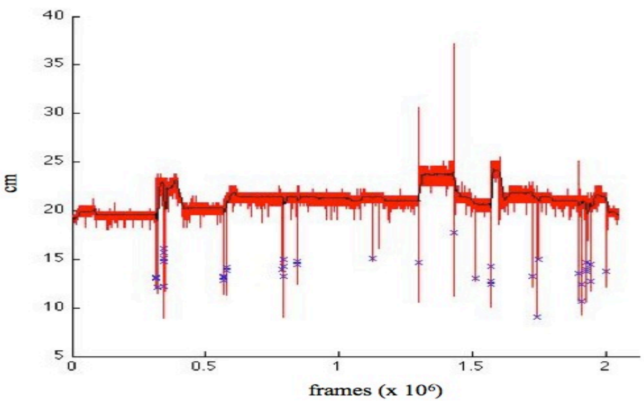

Figure 5. *Detection Results for GHC9: Signal in red, running average in black and blue X's marking the 30 candy-taking detection instances.*

VI. RESULTS

In this section we review user comments and quantitative results. Our statistical analyses include Pearson Chi Square tests of correlations between categories of data (e.g., X conditions strongly relate to the presence of Y), and ANOVA analyses seeking to establish relationships between mean data samples (X conditions predict the likelihood of Y level of outcome).

We will use *instance* to refer to candy detection events and *frames* as a duration unit of how long people spend taking candy (20 frames/sec, i.e. 50ms each). At the offices, candy accessibility was the best predictor of whether people would take candy at the office and people responded most often to *fast* robots addressing them directly. During transit, people more often interrupted *slow* robots.

A. Qualitative Results

Study conductors were often in earshot during robot operation. When the robots were facing the wrong direction many people commented anecdotally that a robot had "made a mistake," "miscalculated," or "had a bug." This indicates that they know the robot is facing the incorrect direction in the *asocial* condition. In addition, faster robot motions may have intimidated people. This is supported by anecdotal comments that the *fast* robot was "kind of scary," and "I wasn't sure it was going to stop." Finally, study conductors often overheard people greeting or thanking the robots after taking candy from them in the hallway, despite the lack of verbal offer. They are probably inferring an implicit offer of candy via non-verbal indicators and prior knowledge, feeling no compunctions about stopping, greeting, and/or inviting their friends to also take candy from the robot.

B. Impact of Task Context

During their candy delivery activities (Table II), the robots spent slightly more time (41.7%) *in transit* than *at offices* (31.3%), 1.3 times as much. They also spent just over a quarter of their time doing something else, like beginning/ending programs, pausing for human assisted re-localization (if stuck), or refilling the candy basket.

TABLE II. TOTAL TIME DISTRIBUTION BY TASK CONTEXT

	In Transit	At Office	Other
Total Time (s)	6859	5160	4445
% Time	41.7%	31.3%	27.0%

Overall, people took candy from the robots almost twice as often when they were visiting offices than while they were in transit (Table III). It makes sense that more detections occurred at the offices because that is the only location where the robots make a verbal offer for people to take candy. Thus, the surprising result is that so much of the candy-taking (more than one-third) occurred while the robot was in transit.

TABLE III. INSTANCES (N=148) AND DURATION OF CANDY-TAKING

	Instances	# Frames
Office	93 (62.8%)	2440 (60.8%)
Transit	55 (37.1%)	1574 (39.2%)

TABLE IV. REVERSE TRICK-OR-TREATINGS STATS

Location	Time	Transit Detects	Office Detects	Total
GHC6	10am	18	12	30
GHC9	1pm	4	26	30
NSH4	1pm	11	22	33
GHC7	3pm	22	23	55

C. Impact of Robot Orientation

Our first experimental variable was robot orientation upon arrival to the office: *social* or *asocial.* We also track candy-basket location: *front-basket* or *back-basket.* We diagram the four possible permutations of orientation and basket location in Fig. 6. We only look at office candy-taking detections, as robot orientation angle only presents itself upon robot arrival to each office.

We find that people's main drive in responding to the robot is candy access not orientation toward the office door. Our evidence is that people respond most quickly (Table VII) to a robot presenting the candy basket in closest proximity, regardless of overall body facing. Duration of candy-taking was strongly predicted by which robot was making the offer (Table VI), indicating either a strong influence of basket location or highlighting population differences among the differing floors at the times of day they were visited. We also present statistical analyses that provide further insight into this data.

Contradicting our first hypothesis, there are slightly more (1.2 times as many) candy detections in response to the robot in the *asocial* condition as the *social* condition (Table V). If we look at the numbers more closely, most of this effect is due to the heightened likelihood of people taking candy from the asocial orientation robot with the back-facing basket.

TABLE V. NORMALIZED DETECTIONS BY ORIENT AT OFFICE (N=93): RATIO OF DETECTIONS TO CANDY OFFERS WITHIN EACH CONDITION

	Back-basket	**Front-basket**	**All baskets**
Social Orient	16.9% (14 detects / 83)	65.0% (26 detects / 40)	33% (40/123)
Asocial Orient	29.6% (24 detects / 81)	55.8% (29 detects / 52)	40% (53/133)
All orientations	23% (38/164)	60% (55/92)	

To evaluate our detection findings statistically, we consider two measures: how well robot orientation predicts *whether* people will take candy (Pearson Chi Square), and how well orientation predicts *number* of detections, i.e., how many times people will take candy (ANOVA).

We find two trends supporting the information value of robot orientation toward predicting people's likelihood of taking candy from a robot with a back-facing basket (comparing the 16.9% ratio to 29.6% ratio in Table V). Specifically, a Pearson test correlating orientation condition to presence of candy-taking detection(s) across all office visits with back-facing baskets has a p-value of 0.0648 (N=164), thus there is a trend relationship between orientation and Boolean detection of candy-taking instances.

Our second test is to run an ANOVA analysis to see if orientation condition can predict *number* of detections, again within the office-detections with back-facing baskets, finding a p-value of 0.0929, another statistical trend. Both tests support the relationship between orientation and candy taking in the back-facing office setting. Namely, people are more likely to take candy from a robot with a back-facing basket in the *asocial* orientation rather than *social* one, because that means the candy is more accessible and visible.

Using a similar approach, we also find that robots facing the office with a *social* orientation condition are more likely to entice people to take candy if they have front-facing baskets (comparing the 16.9% ratio with the 65.0% ratio in Table V). Limiting our analysis to whether basket-location can predict detections within each of the two orientation conditions, we find that basket location has a strong relationship with *whether* people take candy from robot in *social* orientation (N=123). In fact, the Pearson Test shows a p-value of 0.0003* (very significant). An ANOVA analysis looking at the effect of basket location on predicting candy-taking detections for robots in social facing (also N=123) also finds significance with a p-value of 0.0019*. Again, candy accessibility dominates response, this time with two statistically significant results.

All of the above results support that people were least likely to take candy from a robot in the *social* orientation condition with a back-facing basket (perhaps because it is confusing and less accessible). Statistically, the remaining results about likelihood of taking candy are fairly similar. For example, basket location does not correlate with whether or how often people will take candy from a robot in the *asocial* condition (comparing the 29.6% with the 55.8% in Table V), with respective p-values of 0.2266 (Pearson Test) and 0.2035 (ANOVA). There is a similar lack of significance when evaluating whether orientation condition predicts whether people will take candy from a robot with a front-facing basket, although numerically, Table V does show higher relative detection numbers for more-accessible *social* orientation + front-basket setup (65%) as compared to *asocial* + front-basket (55.8%).

Our next major finding is that people were more likely to take candy from the robot with a front-facing basket (60%), than the robot with the back-facing basket (23%), see Table V. Our theory is that it is more intuitive for the robot to be offering candy in the front, and it also makes sense that the earlier in the day visits (GHC6) and less social floors (GHC9) would find fewer people in their offices or open to interaction, as they are further from classrooms and transit corridors (Table IV). Our overall Table I numbers also show a lower ratio of candy-taking detections to offices visited for GHC6 and GHC9 (Cobot2 floors) than NSH4 and GHC7 (Cobot3 floors). We see additional evidence that Cobot3 was more in demand xin Table VI, where mean detection durations more than double for the front-facing basket (more candy-taking may result in more overlapping detections). In fact, an ANOVA analysis of whether the basket-orientation alone can predict detection durations succeeds with a p-value of 0.0006*, i.e., high statistical significance. In contrast, orientation predicting detection duration has a p-value of 0.7872, i.e., irrelevant.

TABLE VI. MEAN DETECTION DURATION BY ORIENT AT OFFICE*

	Back-basket	**Front-basket**	**All baskets**
Social Orient	12.6 (11.9std)	30.0 (29.8std)	24.0 (26.2)
Asocial Orient	14.5 (15.1std)	38.0 (40.0std)	27.3 (33.5)
All orientations	13.8 (14.0)	34.2 (35.7)	

*Duration in frames (every 50ms)

Finally, we evaluate the mean time that the robots spent at the offices with detections across the orientation conditions in Table VII. *Social* orientation + front-basket and *asocial* orientation + back-basket result in the shortest office visits (i.e., people take the candy quickly and dismiss the robot when the candy is in closest proximity). As indicated by the qualitative results, the explanation for this data is not that people did not find the *asocial* condition normal, but rather that they cared more about the fact that there was candy present.

TABLE VII. MEAN TIME AT OFFICES WITH DETECTIONS BY ORIENT & CANDY BASKET LOCATION*

	Back-basket	**Front-basket**	**All baskets**
Social	29.0	23.5	26.2
Asocial	22.2	27.5	24.9
All orientations	25.6	25.5	25.9

*Time in seconds

D. *Impact of Robot Speed*

Our second experimental variable was robot speed condition: *fast* or *slow*. We find diverging results by context (*at office* versus *in transit)*. Namely, people more frequently respond to a *fast*-moving robot when it approaches their *office,* but are more likely to take candy from a *slow*-moving robot *in transit*.

We summarize our overall detection frequency results for how often people take candy from the robots by speed condition in Table VIII. By far, the lowest detection ratio occurred during *fast* transit (at a rate of 7.2%). The next lowest was for *slow* transit (27.6%), a large improvement over *fast*. The highest ratios occurred at offices, 41.1% for *fast* and 31.8% for *slow*. Overall, people took candy at the offices most often, marginally preferring the *fast* robot at the office but generally avoiding the *fast* robot in transit.

TABLE VIII. NORMALIZED DETECTIONS BY SPEED & CONTEXT (N=148)
RATIO OF DETECTIONS TO SAMPLES WITHIN EACH CONDITION

	Office	**Transit**	**All contexts**
Slow	31.8% (42 detects / 132)	27.6% (45 detects / 163)	29.5% (87/295)
Fast	41.1% (51 detects / 124)	7.2% (10 detects / 139)	23% 61/263
All speeds	36.3% (93/256)	18.2% (55/302)	

Mean detection durations were congruent with our detection instances findings (Table IX). Namely, the average candy-taking duration for a *slow*-moving robot lasts longer during transit (33 frames), and the average detection duration for the *fast*-moving robot lasts longer at the office (29 frames). In other words, in the speed conditions where people are more likely to take candy, they also spend more time taking that candy. This may occur because longer durations indicate more candy taking by a single or multiple individuals. For example, if three people take candy at overlapping intervals, there may be 4 seconds of continuous detection, while if just one did; it might be closer to 1.5 seconds of detection. Longer duration might also reflect a relaxed attitude with which people take candy (which may be particularly relevant for the *slow* moving transit condition).

TABLE IX. MEAN DETECTION DURATION BY SPEED AND CONTEXT*

	Office	**Transit**	**All contexts**
Slow	22.1 (3.2 std)	32.6 (5.6 std)	27.6 (30.9)
Fast	28.9 (36.5std)	14.0 (9.3 std)	26.5 (34.0)
All speeds	25.9 (30.4)	29.2 (34.9)	

*Duration in frames (every 50ms)

Impact of Robot Speed at Office

As found in the previous section, people took candy slightly more often if the robot approached at high speed than when the robot approached the office more slowly (Table VIII). Note that the robot speed condition set the overall transit velocity, which people in the offices would only visually or auditorily experience during the robot's final approach to position outside their office doors.

We next evaluate whether the higher frequency of taking candy from the robot approaching an office at high speed is statistically significant. A Pearson Chi Square Test of *whether* speed condition correlates with office candy-taking has a p-value of 0.141, not quite statistically significant, but close to a trend. We also perform an ANOVA analysis tying speed condition to the *number* of office detections, finding a p-value of 0.224 – not significant. After combining multiple detections (the 93 office detection instances occur at 41 offices), we reduce our sample to N=41, out of a total of 256 offices that robots visited successfully. We may just not have a large enough sample size.

Accounting for opportunity cost (see Table X), people take candy 1.16 times as often from the *fast* robots approaching their office than the *slow* ones.

TABLE X. OFFICE TIME TOTALS BY SPEED CONDITION

Slow: 2607 frames (50.5%)	**Fast:** 2553 frames (49.5%)

Similar to our findings about how candy proximity at the office led to a faster response rate, we also find that a robot approaching an office in the *fast* speed condition leaves the office sooner (in 22 vs. 29 seconds), indicating that people may respond more quickly to the high-speed robot at the office (Table XI). Though not yet substantiated statistically, the idea that the robot's higher speed results in higher and/or faster human response rates to robot requests is fascinating, as it would indicate that human accordance to a robot's verbal communications could be influenced by that robot's motion patterns, including the sounds that those motions produce.

TABLE XI. MEAN TIME AT OFFICES WITH DETECTIONS BY SPEED

Slow	29.0 (23.6 std)
Fast	22.2 (27.5 std)

*Time in seconds

Impact of Robot Speed in Transit

During the robot transit task, people take candy more often and for longer mean durations while the robot is moving slowly (Table VIII & IX). We also note that, statistically, neither orientation nor basket location had a statistical relationship with transit detections.

The speed condition findings do indeed show a trend toward predicting how many times people will take candy during the robot's transit. Because of multiple detections along the paths our 55 transit detection instances reduced to N=19 paths with detections (out of 302 attempted paths to an office). In combining detections, we probably lost too much information to establish correlations between speed condition and *whether* there would be transit detections (Pearson test showed little support for relationship between speed condition to Boolean transit detections with p-value 0.7225). Because the ANOVA analysis retains information about the number of detections on each path, it had better results. Using speed condition to predict the *number* of transit detections results in a p-value of 0.077, a clear trend toward statistical significance. Given the small sample size, it is surprising and promising that we were able to detect this result.

In this case, the detection instances for each speed condition have widely varying opportunity cost. The *slow* robot spent twice as much time in transit as the *fast* one (Table XII). Even accounting for this, however, we see a strong influence of speed on people's frequency of taking candy; people took candy 2.19 times as often from a slow-moving robot as a fast-moving robot.

TABLE XII. TRANSIT TIME TOTALS BY SPEED CONDITION

Slow: 4632 frames (67.5%)	**Fast:** 2227 frames (32.5%)

Given the large numerical disparity between transit candy-taking and the statistical trend showing that speed can predict mean values of candy-detections, there is clearly value in further exploring the relationship between speed and people's attributions about robot interruptibility and approachability during transit. To be safe, we recommend the following transit speeds for future robot behavioral designs: use slower speeds if you want people to feel welcome to interrupt the robot, and faster speeds if it is more important for the robot to complete its current task. These settings could be dynamically altered depending on how over or under-scheduled your robot happens to be over the course of a day.

VII. DISCUSSION

As our reverse trick-or-treating program was only designed to offer candy at office doors, it was striking that over one third of the candy-taking detections occurred during transit. Possible explanations for this apparent discrepancy include:

- *non-verbal offer*: the costuming and presence of a basket of candy implicitly communicated an offer for people to take candy to which people responded,
- *prior knowledge:* people knew that the CoBots' task that day was to share candy, as it has been reverse trick-or-treating for several years, and responded appropriately
- *stealing:* baskets of candy are very tempting and thus people take candy without considering/caring whether the robot would want them to take the candy.

Of the 55 candy-taking detections that occurred while the robot was in transit, 49 consisted of people blocking the robots' path in order to take candy, with only 6 instances of people taking candy from a robot's basket while it remained in motion. Because of how people were interrupting the robot, we believe that *non-verbal* offer and *prior knowledge* were most powerful.

Our orientation results demonstrated that robots do not need to behave like people if they are enabling activities that people are intrinsically motivated to complete. While basket location alone had no impact on how long the robot spent at each office, if the basket location matched the robot-facing angle, people could collect candy more efficiently. Thus, we believe that the dominant influence on how long people detained the robot at their offices was congruence between robot orientation and basket location.

It was also noteworthy that so many people interpreted the robots' *asocial* orientation condition as a miscalculation or bug in their verbal descriptions to the study conductors, rather than being intentional. This assumption is consistent with previous findings in which a cheating robot's verbal misreporting of the correct winner [15] or who had won a rock-paper-scissor match [16] were interpreted as computation errors. Perhaps if we had made our *asocial* positioning more extreme, e.g., continuing to orient such that the robot was facing away from the person, or if the repositioning made it harder for them to get to the candy, we would see more people interpreting the robot's behavior as intentional. They might even think it was teasing them.

Numerically, people took candy more often at the offices after a higher speed approach. The fast speed may have made the robot seem adamant about its task (we frequently anthropomorphize machines [17]). Another explanation was that it was easier to hear a robot approaching in the high-speed condition. In fact, the robots' *fast* condition may be more salient to people within the offices for both reasons; sound is an important component of motion. In contrast, in addition to being less intimidating and easier to catch, the slow robots may have also seemed to place less importance on their current task.

Thus, if researchers want a robot to seem more interruptible to people, perhaps they should include nonverbal cues, such as moving at a slower speed toward its current target, that make it seem more approachable.

VIII. CONCLUSION

Interactive robots are increasingly providing services and value to people in real world settings. In this work, we deployed two autonomous robots to deliver candy to four floors of offices, located in two separate buildings. These robots navigated, localized and completed their reverse trick-or-treating tasks autonomously, varying speed and orientation-to-office conditions for each office on their list in random fashion.

We were surprised to discover that natural robot facing at the office had no impact on overall candy-taking behavior and that one-third of all candy-detection instances occurred in the hallway. We believe that the costuming and previous deliveries communicated an implicit non-verbal offer of candy to bystanders in the hallway. The high volume of people taking candy during transit and the ubiquitous blocking of the robot's path supports this deduction.

Our most important statistical findings were that a robot's social positioning is less important than candy accessibility, and that people prefer to interrupt a robot moving slowly, rather than one moving quickly. We believe that people's desire to take candy overwhelms any concern they might have about a robot's correct social facing, and that slow robots seem more interruptible than fast robots. In addition, people take candy more often and quickly from the *fast* robot at the offices. Thus, higher speed approaches may be more salient to people for a robot making a direct request.

In future work, it would be useful to compare these results to a study in which the robot were to ask people to complete a less desirable task, such as filling out a survey. In such cases, we might find a stronger effect from natural social positioning (why bother responding to a robot facing the wrong direction), and lower interruption frequency during transit, even with clear brochures and heavy advertisement. We would also be curious to see how these results would generalize to other task applications.

In summary, this work provides several key insights:

- Robots do not need to behave like people if they are enabling activities that people are intrinsically motivated to complete.
- People can interpret non-verbal indications that a robot is offering candy without an explicit verbal offer.
- People are more likely to interrupt a robot in a low speed condition, thus, robot transit speeds could be used to influence operational and interaction goals.
- Human accordance to a robot's request may be influenced by that robot's motion patterns.

ACKNOWLEDGMENT

Thanks to the CoBot team, especially Joydeep Biwas and Eleanor Avrunin. Thanks also to Aaron Steinfeld and the National Science Foundation for supporting this work.

REFERENCES

[1] J. Biwas, *Vector Map-Based, Non-Markov Localization for Long-Term Deployment of Autonomous Mobile Robots*. Thesis, Robotics Institute, Carnegie Mellon University, December 2014.

[2] T. Kanda, et al, "An affective guide robot in a shopping mall," in *Proc. of Int'l Conference on Human-Robot Interaction*, 2009.

[3] A. von der Pütten, N. Krämer, C. Becker-Asano, K. Ogawa, S. Nishio and H. Ishiguro, "The Uncanny in the Wild. Analysis of Unscripted Human-Android Interaction in the Field," in *International Journal of Social Robotics*, vol. 6, no. 1, pp. 67-83, January, 2014.

[4] B. Mutlu and J. Forlizzi. "Robots in organizations: the role of workflow, social, and environmental factors in human-robot interaction," in *Proc. Int'l Conf. on Human-Robot Interaction*, 2008.

[5] R. Kirby, et al, "Designing Robots for Long-Term Social Interaction." in *Proc. Int'l Conf on Intelligent Robots and Systems*, 2005.

[6] S. Rosenthal, J. Biswas, and M. Veloso, "An Effective Personal Mobile Robot Agent Through Symbiotic Human-Robot Interaction," in *Proc. Conf. on Autonomous Agents and Multi-Agent Systems*, 2010.

[7] J. Forlizzi and C. DiSalvo, "Service robots in the domestic environment: a study of the roomba vacuum in the home," in *Proc. of Int'l Conference on Human-Robot Interaction*. 2006.

[8] L. Takayama, and C. Pantofaru. "Influences on proxemic behaviors in human-robot interaction," in *Proc. IROS '09 the Int'l Conference on Intelligent Robots and Systems*, 2009.

[9] R. Kirby, J. Forlizzi and R. Simmons, "Natural Person-Following Behavior for Social Robots," in *Proc. HRI '07 Int'l Conf on Human-Robot Interaction*, March, 2007.

[10] M. Karg, et al, "Body Movements for Affective Expression: A Survey of Automatic Recognition and Generation," in *IEEE Transactions on Affective Computing,* 4.4 (2013): 341-359.

[11] H. Knight and R. Simmons, "Expressive Motion with X, Y and Theta: Laban Effort Features for Mobile Robots," in *Proc. IEEE Int'l Symposium on Robot and Human Interactive Communication*, 2014.

[12] B. Coltin and M. Veloso, "Online Pickup and Delivery Planning with Transfers for Mobile Robots," in *Proceedings of ICRA'14, the IEEE/RSJ Int'l Conference on Robotics and Automation*, 2014.

[13] B. Coltin, J. Biswas, D. Pomerleau, and M. Veloso, "Effective Semi-autonomous Telepresence," in *Proc. of RoboCup Symposium*, 2011.

[14] J. Biswas and M. Veloso, "Localization and Navigation of the CoBots Over Long-term Deployments," *International Journal of Robotics Research*, December 2013, vol. 32, no. 14, pages 1679-1694.

[15] M. Vazquez, A. May, A. Steinfeld, and W. Chen, "A Deceptive Robot Referee in a Multiplayer Gaming Environment," in *Proc. Int'l Conf. on Collaboration Technologies and Systems*, 2011.

[16] E. Short, J. Hart, M. Vu, and B. Scassellati, "No fair!!: An interaction with a cheating robot." in *Int'l Conf. Human-Robot Interaction*, 2010.

[17] B. Reeves and C. Nass, *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places.* Center for the Study of Language and Information, Jan 29, 2003.