

Towards More Efficient Navigation for Robots and Humans

David V. Lu and William D. Smart

Abstract—Effective robot navigation in the presence of humans is hard. Not only do human obstacles move, they react to the movements of the robot according to instinct and social rules. In order to efficiently navigate around each other, both the robot and the human must move in a way that takes the other into account. Failure to do so can lead to a lowering of the perceived quality of the interaction and, more importantly, it can also delay one or both parties, causing them to be less efficient in whatever task they are trying to achieve.

In this paper, we present a system capable of creating more efficient corridor navigation behaviors by manipulating existing navigation algorithms and introducing social cues from the robot to the human. We give the results of a user study, demonstrating the effectiveness of our system, and discuss how it can be applied more generally to a wide variety of situations.

I. INTRODUCTION

Subtle differences in robot behavior can lead to significant differences in how humans react to that behavior, and to their perception of the robot. These subtleties can be the difference between behaviors that are socially acceptable and those that are not, between clear and meaningful gestures and ones that are confusing. In many cases, *how* a robot performs an action is just as important as *what* it does.

Our overall goal is to improve the interactions between a human and a robot, by making subtle changes to the robot's behavior. Before we begin, however, we must define what we mean by *improve*. For the purposes of the work reported here, there are three ways to improve an interaction: (1) to increase the efficiency with which the robot performs the task in which the interaction happens; (2) causing the human's subjective perceptions of the robot to become more positive; and (3) to increase the efficiency with which the *human* performs whatever task they are engaged in. This third aspect is one that is often neglected, but is one of the main reasons that we want people to interact with robots; to make the humans more efficient at what they do.

Take, for example, the seemingly simple task of navigating in an indoor office environment, as seen in figure 1. Successful algorithms to ensure collision-free obstacle-avoiding paths have existed for decades. However, with the addition of humans to the environment, the problem becomes more difficult. Not only do obstacles around the robot move, they adjust their movements in relation to how the robot moves. As a result, the robot must take into account the fact that some obstacles are intelligent, independent agents. We claim that the robot needs to use some form of predictive model for




Fig. 1. An Example of Standard vs. Social Navigation - These diagrams show results from an experiment as a robot (blue square) and a human participant (yellow circle) pass each other in a hallway. In the top example, using the standard navigation, the person is forced to slow down drastically while the robot passes (c). With the social navigation in the bottom example, the person is able to pass the robot passes the robot with much greater ease and ends up completing the task quicker as a result.

Lu is a PhD candidate in the Dept. of Computer Science at Washington University, St. Louis, Missouri, 63130 davidlu@wustl.edu
 Smart is faculty in the Department of Mechanical Engineering, Oregon State University bill.smart@oregonstate.edu

how such agents are likely to behave around it, and navigate accordingly.

Furthermore, a human interacting with a robot is likely going to develop their own predictive model, often known in the psychology and philosophy literatures as a theory of mind, for the robot. Prior studies [1], [2] have shown that people subconsciously treat technology socially. They attribute to it a rational mind similar to their own, capable of making decisions and interacting in traditional human ways. If the robot does not conform to this theory, if it does not fit their predictive model, then the person will initially be confused by the robot's behavior, and will be forced to learn how it is likely to behave. This is a slow process and will impede the progress of the interaction. However, if the robot fits their expectations, then they are able to use all of their (often implicit) prior knowledge of social interactions to predict how the robot will behave, leading to a more efficient interaction.

In this paper, we describe the use of two different social cues on a mobile robot as it passes a human in a corridor: changing the path that the robot takes, and whether or not it makes eye contact with the human. We measure the effectiveness of these two cues by their effect on the time taken for both the robot and the human to traverse the corridor, and verify their effectiveness with a user study.

II. RELATED WORK

Several systems have addressed the problems of navigating [3]–[5] and localizing [6] robots in the presence of people. However, most systems effectively ignore the problem, treating all obstacles, human or not, equally (see, for example [7]).

Algorithms that *do* adjust their navigation in people-specific ways take a number of forms. The patrol robot studied by Hayashi et. al [8] is an excellent example of a robot that uses different paths and gaze behaviors around people in order to affect the interaction. However, that effect was only studied using subjective survey-based measures, and did not look at the resulting efficiency of the interaction.

Kirby [9], [10] also addressed problems of social navigation by altering the robot's paths with respect to people's personal space and other social conventions. The resulting interactions were evaluated using both qualitative measures (people's opinion) and analysis of where the robot passes the humans. There is also similar work which considers proxemic behavior as a way to find optimal positioning for interactions [11]. None of these works consider the human's behavior and efficiency as measures of the quality of the interaction.

The work of Sisbot et. al [12] integrates the relative positioning of the human and robot, the human's field of view and other concerns into the robot's internal representation of the world, and uses that information to plan more human-friendly paths. In order to justify how such a system would be socially beneficial, the authors use human evaluation studies that occurred prior to the design of the algorithm rather than evaluations of the algorithm itself.

There is also work that views pedestrians as dynamic obstacles, which the robot then uses to create a spatio-temporally optimal plan [13]. However, this work assumes fixed human trajectories, and thus does not take into account how the humans might react to the robot's behaviors as it approaches them.

In addition to manipulating the path, many robots also use gaze to communicate information about the task [14]. Gaze communicates "joint attention" and helps establish mutual understanding of the scope of the interactions [15]. Such gazes can also help to resolve ambiguities that are not otherwise communicated [16], and can help the human decide on an appropriate action to take. Gaze has also been shown to affect a person's proxemic preferences around a robot [17], resulting in people varying their approach distance to a robot based on the robot's gaze.

III. SOCIAL NAVIGATION BEHAVIORS

As a robot and a person approach each other in a hallway, it is often unclear what the robot should do in order to most efficiently pass the person. When two people pass each other, they have a shared body of implicit knowledge about social situations, and swap a multitude of subtle social cues in order to manage the interaction. Both of these are typically missing in the human-robot setting. One source of problems is the path taken by the robot; most path-planning algorithms used for navigation will direct the robot to drive straight down the center of the hallway until it gets close to the person. Only when it cannot move any further without risking a collision will it move to the side or steer around the person. While this approach is valid for inanimate objects, or even other robots, it introduces two problems when dealing with humans. First, the robot does not make clear to the human which side of the hallway they should pass on, forcing them to guess, or to rely on the prevailing social norms (which are not implemented in traditional path-planners). Worse yet, the robot will likely not make any indication that it has even detected the human. This makes it unclear whether the robot is driving straight forward because it has not seen the human, or because it *has* seen them, but has chosen not to react. Both of these problems can lead to an uncomfortable interaction for the human, because of an inability to predict what the robot will do and, often, the physical proximity of the robot as it treats the human as just another obstacle.

In this paper, we address these two problems, inappropriate paths and poor signaling of intent, with two techniques. Similarly to the work discussed in section II, we modify the robot's costmaps to reflect the social behaviors we want to show in the planned paths. To better communicate the robot's intent to the human, we use a gaze behavior that directs the robot's head at either the human or at the hallway ahead.

Our goal in implementing these behaviors is to make the interaction more natural, and hence efficient, for the human. In a pilot study with two human actors in a motion-capture environment [18], we explicitly studied the scenario of two people passing each other in a hallway. We observed that

both gaze and the relative position of each person played a role in the passing behavior.

Our general hypothesis is that by employing some combination of these two behaviors, we can improve the interactions in the three ways discussed in the introduction: improving robot performance, improving the person's impression of the robot, and affecting the person's behavior. We hypothesize that these changes will lead to more predictable robot behavior, ultimately resulting in a more efficient interaction.

A. The Robot

The work reported in this paper was done on a PR2 robot, a mobile manipulation platform developed by Willow Garage, running the open-source Robot Operating System (ROS) software [19]. The robot has a quasi-holonomic base, two arms, and a pan-tilt head. In this work, the arms are not used and were kept in the folded position. The PR2 is equipped with two laser range-finders, one at ankle height, and one below the head that tilts up and down, and multiple cameras. The cameras in the head give the appearance of eyes, which is important for a gaze behavior, but the data from the cameras are not used in the experiment reported here.

B. Lasers for People Detection

Before causing the robot to act appropriately in the presence of people, we must first determine where the people are. In the work reported here, we use the laser range-finders to detect people, primarily because of their accuracy in our experimental environment. We note, however, in a more general setting, another more powerful people-detector might be more appropriate.

We reduce the problem of detection people to that of detecting legs. Our leg detection technique is based on the algorithm of Arras et al [20] and extends an implementation developed at Willow Garage by Caroline Pantofaru.¹ Like Arras and Pantofaru, we use a group of low-level classifiers to determine the probability that a sequence of laser readings is a leg or not. These leg probabilities are then passed to an algorithm that pairs the individual legs, based on distance constraints, and tracks the resulting leg-pairs, which correspond to a person under our assumptions. While there may be more advanced information about the person beyond their location through time that may help adjust the robot's behavior, this approach uses less processing than more advanced state recognition and provides an adequate advance over the previous implementation.

C. Gaze Behavior

We use the detected locations of humans to control the gaze behavior. The default behavior is to have the robot looking straight ahead. When a person is detected, the position of their head is estimated approximately, using the center of their detected legs and an average human height. The robot's pan/tilt head is then pointed at this position, simulating gaze. Initial prototyping revealed that having the

robot look at the person constantly gave an impression of being "creepy", which is consistent with psychological studies with two human subjects [21]. Furthermore, this persistent gaze lowered the person's confidence that the robot knew where it was going, since it was constantly not looking in the direction it was traveling. Hence, for the actual experiment, we chose the robot's gaze to be on a cycle: 5 seconds of looking at the person, followed by 5 seconds of looking ahead. The intent was to give the person acknowledgement that the robot saw them, while avoiding the downsides of a constant leer. We acknowledge that more sophisticated models are possible, but chose instead to focus on improving the robot's path planning, which turned out to be much more problematic.

D. Socially Aware Costmaps

We choose to implement our changes to the paths the robot drive along within the existing ROS navigation and costmap architecture. This allowed for flexibility in choosing which path planning algorithm to use and retains the space-oriented nature of our navigation constraints. Another possible approach is to create waypoints that reflect the desired paths we would like to create. However, this approach runs into trouble in confined environments such as the corridor where the points along the desired path may result in colliding with the wall.

1) *Standard Navigation Algorithm:* The navigation capabilities in ROS [7]² are used on many different robot platforms. The navigation system is flexible in terms of the variety of sensor information it can use. This information is fed into the costmap and then different path planning algorithms (e.g. Dijkstra's, A*) can be used to find a path that minimizes the total cost of the cells the robot passes through. Dynamic/moving obstacles are dealt with by replanning often enough to avoid the obstacles. Implicitly, the robot assumes that people will get out of the way on their own.

However, in the ways discussed earlier, it does not work optimally in the presence of people. First, it treats all obstacles read in by its sensors equally; steering around people as though they were any other obstacle. The resulting paths often move uncomfortably close to people. Furthermore, the path taken is designed to be as short as possible, due to the nature of the default costmaps, not to communicate any information about where it is going to go.

Unfortunately, ROS does not have features that allow us to solve this problem. By manipulating the costmap, we can indirectly affect the paths that the robot chooses to take, while retaining compatibility with all of the existing robots using ROS. The implementation of the navigation system only allows users to manipulate the costmap in certain ways, treating the input data in one of two ways: cells in the costmap are either marked as occupied, or are cleared when a sensor ray passes through them. Free space is set to a low default cost (0), whereas obstacles are marked with a

¹http://ros.org/wiki/leg_detector

²<http://ros.org/wiki/navigation>




Fig. 2. Three different paths with three different costmaps. The robot is marked in blue, the detected person in red and the path in yellow. The white areas show the locations of obstacles (either the walls or the person’s legs). The tints of green show the other costs, ranging from bright green lethal points to darker tints for lower values.

high lethal cost (255). All nearby points that would put the robot in collision with the obstacles are assigned a slightly lower lethal cost (254). Points immediately surrounding the lethal points are assigned a cost that exponentially decays relative to the distance to the lethal points, ensuring that the robot does not drive too close to the wall unless it must. A visualization of this costmap schema can be seen in the top portion of figure 2.

The default implementation of the navigation system often results in the robot often driving within 50cm of obstacles, which can be quite unsettling when that obstacle is a person. People are generally not comfortable with a robot entering what Hall defines as the person’s intimate proxemic space [22], despite the fact that the robot does not collide with them.

The problem is that the current implementation only deals with hard constraints on the costmap space. If we were to try to increase the space around each person with these constraints, many paths that are valid (non-colliding) would be excluded. If the environment is cluttered, then there may be no possible path with hard constraints applied. A more appropriate approach would use softer constraints on the planned paths, where we can specify that we would rather a robot did not enter certain spaces, but would allow it if needed.

2) *Flexible Costmap Changes*: We have implemented an alternate version of the navigation system which allows for arbitrary changes to be integrated into the costmap. After the sensor information is added and the costs are updated around each obstacle, our version allows for user-specified, independently-compiled plugins to make changes

to the costmap. This architecture enables users of the new costmap to be able to insert any manner of changes to the costmap without requiring recompilation of the core navigation system. This architecture also runs the plugin code in the same process as the rest of the system, avoiding the large memory transfers that would be required if the costmap changes were specified in another process, and transmitted over the normal ROS communication channels.

We experimented with the use of Gaussian costmap adjustments, as suggested by Kirby [10] and others, as seen in the middle portion of figure 2.

In this approach, values are added to the costmap around the human’s detected location, according to a 2d Gaussian, possibly taking into account the person’s direction of travel. This approach works well in the general case, as it causes the robot to take smooth paths that are farther away from people. It also works as an analogy to the personal space/proxemic concerns discussed earlier. However, it requires fine tuning of the parameters to get the desired behavior, and we found that sometimes there are no (readily-found) values of the parameters that realize the desired behavior. For instance, in our corridor, we wanted the robot to move to the side of the hallway it was going to pass on as soon as possible. In theory, increasing the amplitude or variance on the Gaussian function should result in the paths moving further away from the center of the Gaussian. However, the path planning algorithms need to find balance between the lowest cost paths and the shortest paths. Once the costs become sufficiently high (either by increasing the amplitude or variance), the decrease in penalty for moving further from the obstacle is outweighed by the cost of taking a longer path. This results in the robot reverting to its earlier behavior of driving straight toward the person along the shortest path.

Instead of the Gaussian approach, we implemented a cost function that was more specific to the environmental context. This “linear” cost function varied depending on what side of the hallway the person was detected on. If the person was detected on the left, the costs would decline linearly from the left side to the right side (and similarly if they were on the right), with values ranging from 140 on the undesirable side of the hall to 20 on the other. Similar to the Gaussian function, this would cause the lowest cost next to the person to be all the way against the wall. However, this would persist through the entire hallway, resulting in the robot moving to the side of the hallway earlier than in the Gaussian, as can be seen in figure 2. This sends a social signal to the human sooner and more clearly, giving them time to interpret and react to it. Furthermore, this costmap modification requires little tuning to get the designed behavior. A further discussion of the irregularities and behavior of Gaussians in costmaps can be found in our upcoming paper [23].

While this solution is less general than the Gaussian function, we argue that the environment informs human-human social behavior. We instinctively follow the architectural cues in buildings. This context can be, we claim, captured in special-purpose costmap modifications that depend on the current location of the robot.

Our new version of the navigation system is available at our open-source repository: <http://code.google.com/p/wu-robotics>. We encourage interested readers to download it, try it out, and contribute their own context-specific costmap-modification plugins.

IV. EXPERIMENTAL DESIGN

A. Procedure

The goal of our experimental design was to test how the two social navigation behaviors we implemented would affect interactions between people and the robot. We devised a fetch and carry task, in which participants were told that they would be working with the robot to deliver boxes to different rooms on a hallway. This scenario was selected for three reasons. First, it is typical of the kind of interaction a person unfamiliar with robots is likely to encounter in the near future. Robots are being used in uncontrolled environments, such as hospitals, where people may encounter a robot with no prior explanation as to what it is likely to do. Secondly, delivering packages or mail is a task that robots are already capable of doing. Finally, giving the person a specific task gives them reason to complete the task in an efficient manner. The boxes were shoe-box sized and empty, so that any changes in speed due to the weight or impedance of carrying a box is negligible. The robot is 67cm wide and the hallway is approximately 150cm wide, meaning that people will be able to pass the robot in most cases, regardless of whether the robot is all the way to one side or not.

Participants were recruited from the Washington University campus community. All participants were 18 years or older, with English language skills, and no prior experience working with robots. After the consent process, the scenario was explained to them. They were instructed to fetch or deliver a box three times, which necessitated them walking from room A (at one end of the hallway) to room C (on the far end) and back. On the first trip, the robot would be driving toward the participant as they went to room C and stationary at the door to room B (midway down the hall) on the way back. On the second and third trips, the participant would encounter a stationary robot on the way out, and the robot driving toward them (and room C) on the way back. In the cases where the robot was stationary, the robot's position required the participant to navigate around the motionless obstacle, allowing us to get a base time for how long they would take to walk down the hall irrespective of the navigation algorithm in use.

We use our two navigation behaviors to form a 2x2 study design. The robot's interaction in the hallway was dictated by one of these four sets of behaviors.

- C1 Standard navigation, no gaze behavior
- C2 Standard navigation, gaze behavior
- C3 Social navigation, no gaze behavior
- C4 Social navigation, gaze behavior

At the conclusion of the three deliveries, the participant was given a questionnaire to fill out. The participants were compensated for their time.

Thirty people participated in the study (N=30), with ages ranging from 18 to 70 (mean=27.7, SD=12.4, median=23). 60% (N=18) of the participants were female, and roughly half (53%, N=16) were native to the United States. Conditions C2 and C3 were experienced by seven participants each, while eight participants each experienced conditions C1 and C4. The data for each participant are divided up into six trials, defined as going from one end of the hallway to the other. Out of a total of 180 possible trials down the hallway, 177 are eligible for analysis, with the remaining three removed due to sensor malfunctions or external interference.

For measuring the location of both the robot and the participant, in addition to the robot's laser scanner, a Microsoft Kinect device was placed at each end of the hallway. These devices were stationary, and the only elements moving within the field of view were the robot and participant, allowing for very accurate measurements of the the locations of the targets through time. Signalling distance information was derived from these measurements as well.

B. Hypotheses

We chose to study the following hypotheses with our study.

H1: People encountering the robot in the No Gaze condition will take longer to walk past the robot than those encountering the Gaze condition.

H2: People encountering the robot navigating with the social algorithm will take less time to walk past the it than when it uses the standard navigation.

V. RESULTS

The metrics used here to measure the effects of the different conditions are similar to the "Passage Behavior Parameters" used by Pacchierotti et al [5]. (Pacchierotti et al did not use the parameters as metrics; rather, they used them to manipulate the conditions of their experiment.) The two main metrics are *speed* and *signaling distance*, for both the human and the robot. We use speed to represent the human and robot's ability to get the task completed. Faster speeds are indicative of more efficient interactions. There are several different ways that we measure the speed. First, we have overall average speed, measured over the time that the human participant is in the hallway. We notate this with HS for the human's speed, and RS for the robot's speed. To further examine the vital period before the human and robot have passed each other in the hallway, we use metrics for the speed before they pass, HS_0 and RS_0 . Signaling distance is the distance between the person and the robot when the person moves towards the side of the hallway they will pass on. This is based on the assumption that people walk down the center of the hallway until they have seen their obstacle *and* decided how they are going to pass the obstacle.³ The person's signaling distance is HD and the robot's is RD .

An illustrative example of the differences in navigation styles can be seen in figure 1. Consider the top example from

³This assumption was, observationally, borne out in our experiments.

TABLE I
METRICS FOR EXAMPLES SHOWN IN FIGURE 1.

$[m/s]$	HS	RS	HS_0	RS_0	HD	RD
Trial #26 (C1)	0.88	0.26	0.74	0.37	2.92	1.44
Trial #175 (C4)	1.18	0.41	1.13	0.42	6.98	3.14

a trial under condition C1. Both robot and human start by moving straight down the middle of the hallway (1a). As they get nearer, the human moves to the side of the hallway (1b); however, the robot continues on the same path as before. This causes the person to slow down almost to a stop as they wait for the robot to pass by (1c). Although the robot does in fact yield slightly, it only does so once the person has already reacted to its path by stopping. Both parties continue on their way after they pass (1d). Contrast this to the bottom example under condition C4. This trial starts off similarly with both moving down the center of the hallway (1e). However, they both move to the sides of the hallway early on (1f) allowing for a nice quick pass (1g). This allows both the human and robot to reach their destinations sooner than the previous example (1h). To consider these two examples in concrete terms, see table I to see the metrics.

First, let us examine the robot signaling distance (RD) under the two navigation conditions. Comparing all the trials where the robot was moving, we find that the social navigation condition had a significantly larger ($p < 0.002$) average signaling distance (1.7m) than the standard navigation condition (3.0) on average. With the signaling distance almost doubling, our hope was that this would also have an effect on the speed of the person as well.

It turns out that, under certain conditions, the human speed before passing, HS_0 , is significantly greater when using the social navigation. A full listing of the speeds over all trials can be seen in table II. While not all of the differences listed in the table are statistically significant, there are some notable trends.⁴ In support of H2, if we compare C2 and C4 (looking at all the trials with the gaze condition), we find that the human participants are significantly faster ($p < 0.04$) while using the social navigation ($HS_0(C4) = 1.196ms^{-1}$) compared to the standard navigation ($HS_0(C2) = 1.117ms^{-1}$). The trend is similar in C1 and C3, with $HS_0(C3) = 1.242ms^{-1}$ and $HS_0(C1) = 1.217ms^{-1}$, though not it is not statistically significant ($p = 0.29$).

One surprising result was the relative speeds, comparing across the gaze conditions. Contrary to what we predicted in H1, the immobile head (NoGaze) actually led to participants walking faster than in the Gaze behavior ($HS(NoGaze) = 1.243ms^{-1}$, $HS(Gaze) = 1.186ms^{-1}$, $p < 0.02$). We had hypothesized that the gaze behavior would let the person know that the robot saw them and, hence, the person would walk more confidently, and more quickly, down the corridor. This proved not to be the case. We believe that the additional behavior was distracting for two reasons. First, diverting the gaze often meant the robot

⁴Statistical significance is determined by using a 1-tailed t-test with the speed for each trial considered as a separate measurement.

TABLE II
SPEED AND PASSING DISTANCE OF ROBOT AND HUMAN IN ALL EXPERIMENTS. DARKER BACKGROUNDS INDICATE HIGHER VALUES.

	HS	RS	HS_0	RS_0	HD	RD
C1	1.24	0.33	1.22	0.38	5.23	1.52
C2	1.16	0.36	1.12	0.42	5.40	1.89
C3	1.25	0.34	1.24	0.38	4.80	2.68
C4	1.21	0.32	1.20	0.36	4.68	3.31
Standard (C1+C2)	1.20	0.34	1.17	0.40	5.31	1.70
Social (C3+C4)	1.23	0.33	1.22	0.37	4.74	3.01
NoGaze (C1+C3)	1.24	0.34	1.23	0.38	5.02	2.08
Gaze (C2+C4)	1.19	0.34	1.16	0.39	5.03	2.63
Overall	1.21	0.34	1.19	0.39	5.02	2.35

would be traveling in a separate direction than the direction it was looking in. This breaks from Kirby’s definition of desirable robot navigation qualities by not facing the direction of travel. The second reason is that glancing at the other entity in the hallway could be construed as looking to start an interaction, which could slow the person down. We find these two reasons to be adequate for explaining why we found the opposite of our original hypothesis.

The survey data did not produce any relevant statistically significant results. We discuss this further in section VI.

There were a number of other interesting results we were able to gather from our data. As expected, the stationary robot leads to significantly higher average human speed ($p < 0.003$, $1.179ms^{-1}$ vs. $1.251ms^{-1}$) than when the robot is moving. This comes as no surprise since it is easier to predict where a stationary object will be. This is also shown by the fact that the signaling distance is much greater when the robot is stationary ($p < 0.0001$, $HD(robot_stationary) = 6.178m$ vs. $HD(robot_moving) = 3.870m$). When the robot is not moving, the person is able to figure out what side of the hallway the robot will be on much more quickly, and thus can make their decision about which side of the hall to walk on themselves.

We also found that people were more likely to pass the robot on their right side in the social navigation condition (87% in C4) than in non-social condition (67% in C2) ($p = 0.06$). This fits with the prevailing social norms in the United States, where the experiment was conducted.

Additionally, in the trials where the robot was not moving (i.e. was not displaying any particular navigation behavior) we still noticed a change in the person’s behavior. Based on which navigation algorithm they saw in the rest of their trials, the person would decrease their signaling distance (HD) when interacting with an immobile social navigating robot ($p < 0.03$).

VI. DISCUSSION

From the results of this user study, we conclude that there are multiple different effective ways for robots to navigate around people, depending on the desired priorities. If we place sole priority on getting *people* to where they need to go, the clear solution is to park the robot. Keep the robot out of the way, and people will get to their destinations promptly.

If the priority is solely on the *robot's* navigation effectiveness, the best strategy might be the (untested) strategy of ignoring the obstacles associated with detected people, on the assumption that they will move out of the way as needed. However, this strategy requires that the human endow the robot with higher status and thus give way to its priorities, which may not be the case. Nevertheless, most situations call for a mixing of priorities. The system cannot be optimized for just one party or another. The best solution for the combined robot/human system likely involves elements of social navigation and additional secondary actions that help the robot convey its intent to the humans with whom it interacts.

One unintentional result of social navigation is that the robot moved more slowly once the person was detected. The local planning algorithm, which sets the speed, causes the robot to move at speeds inversely proportional to the sum of the costmap cells it is going to drive into. Future work will look to eliminate such a constraint to enable both the human the robot to be more efficient with social navigation, as well as incorporating a more sophisticated gaze behavior.

Another conclusion that can be drawn from our work is that people will interact with the different algorithms with different levels of effectiveness despite the fact that there is no statistically significant difference in their opinion of the robot. This suggests that people are not consciously aware of many of these subtle social cues, regardless of the effect they have on the interaction. This fact lends support to part of Reeves' and Nass' findings [2] that social interaction with machines is often subconscious and inherently social.

The social navigation behaviors we implemented are obviously not an exhaustive set. The linear costmap adjustment we used is only one of many context sensitive adjustments the robot may need. It is not designed to fit the general case of person-aware navigation, but rather to fit into a general solution. The work described in section II describes a number of other costmap changes that could be integrated into this system. It is for this reason that we maintain that a modifiable architecture like the one implemented here is essential, in that it allows people to implement any sort of costmap adjustment and integrate it with the state-of-the-art navigation system that ROS already provides. Among the many benefits of having a flexible open-source extensible solution is that it can be widely adopted on many systems, and can be used to integrate even more contextual information into the costmap. Having an open solution to the problem of social navigation paves the way for a future in which robots can integrate as many factors as humans do when planning paths around other people.

VII. ACKNOWLEDGEMENTS

Many thanks to the user study participants, Daniel A. Lazewatsky, Robert Pless, Elise Lu and Daniel B. Allan.

REFERENCES

- [1] C. Nass, Y. Moon, and P. Carney, "Are People Polite to Computers? Responses to Computer-Based Interviewing Systems," *Journal of Applied Social Psychology*, vol. 29, no. 5, pp. 1093–1109, 1999.
- [2] B. Reeves and C. Nass, *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. New York, NY: Cambridge University Press, 1996.
- [3] M. Yoda and Y. Shiota, "The Mobile Robot Which Passes A Man," in *Proc. RO-MAN'97*, Japan, 1997, pp. 112–117.
- [4] P. Althaus, H. Ishiguro, T. Kanda, T. Miyashita, and H. Christensen, "Navigation for Human-Robot Interaction Tasks," in *Proc. ICRA'04*, vol. 2, New Orleans, LA, 2004, pp. 1894–1900.
- [5] E. Pacchierotti, H. Christensen, and P. Jensfelt, "Human-Robot Embodied Interaction in Hallway Settings: a Pilot User Study," in *Proc. RO-MAN 2005*, Nashville, Tennessee, 2005, pp. 164–171.
- [6] S. Thrun, M. Bennewitz, W. Burgard, A. Cremers, F. Dellaert, D. Fox, D. Hahnel, C. Rosenberg, N. Roy, J. Schulte, *et al.*, "MINERVA: A Second-Generation Museum Tour-Guide Robot," in *Proc. ICRA 1999*, vol. 3, Detroit, Michigan, 1999.
- [7] E. Marder-Eppstein, E. Berger, T. Foote, B. P. Gerkey, and K. Konolige, "The Office Marathon: Robust Navigation in an Indoor Office Environment," in *Proc. ICRA 2010*, Anchorage, Alaska, 2010.
- [8] K. Hayashi, M. Shiomi, T. Kanda, and N. Hagita, "Friendly Patrolling: A Model of Natural Encounters," in *Proc. RSS 2012*, Sydney, Australia, 2012, p. 121.
- [9] R. Kirby, R. Simmons, and J. Forlizzi, "COMPANION: A Constraint-Optimizing Method for Person-Acceptable Navigation," in *Proc. RO-MAN 2009*, Toyama, Japan, 2009, pp. 607–612.
- [10] R. Kirby, "Social Robot Navigation," Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, May 2010.
- [11] S. Tranberg Hansen, M. Svenstrup, H. Andersen, and T. Bak, "Adaptive Human Aware Navigation Based on Motion Pattern Analysis," in *Proc. RO-MAN 2009*, Toyama, Japan, 2009, pp. 927–932.
- [12] E. Sisbot, L. Marin-Urias, R. Alami, and T. Simeon, "A Human Aware Mobile Robot Motion Planner," *IEEE Transactions on Robotics*, vol. 23, no. 5, pp. 874–883, 2007.
- [13] M. Phillips and M. Likhachev, "SIPP: Safe Interval Path Planning for Dynamic Environments," in *Proc. ICRA 2011*, Shanghai, China, 2011, pp. 5628–5635.
- [14] S. Fussell, L. Setlock, and E. Parker, "Where do Helpers Look?: Gaze Targets During Collaborative Physical Tasks," in *Proc. CHI 2003*, Fort Lauderdale, Florida, 2003, pp. 768–769.
- [15] M. Staudte and M. W. Crocker, "Visual Attention in Spoken Human-Robot Interaction," in *Proc. HRI 2009*, San Diego, California, 2009, pp. 77–84.
- [16] R. Ros, S. Lemaignan, E. A. Sisbot, R. Alami, J. Steinwender, K. Hamann, and F. Warneken, "Which One? Grounding the Referent Based on Efficient Human-Robot Interaction," in *Proc. RO-MAN 2010*, Viareggio, Italy, 2010, pp. 570–575.
- [17] L. Takayama and C. Pantofaru, "Influences on Proxemic Behaviors in Human-Robot Interaction," in *Proc. IROS 2009*, St. Louis, Missouri, 2009, pp. 5495–5502.
- [18] D. V. Lu, A. Pileggi, and W. D. Smart, "Multi-person Motion Capture Dataset for Analyzing Human Interaction," in *RSS Workshop on Human-Robot Interaction*, Los Angeles, California, July 2011.
- [19] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng, "ROS: an open-source Robot Operating System," in *ICRA Workshop on Open Source Software*, Kobe, Japan, 2009.
- [20] K. Arras, O. Mozos, and W. Burgard, "Using Boosted Features for the Detection of People in 2D Range Data," in *Proc. ICRA 2007*, San Diego, California, 2007, pp. 3402–3407.
- [21] M. Argyle and M. Cook, *Gaze and Mutual Gaze*. Cambridge University Press, 1976.
- [22] E. Hall, *The Hidden Dimension*. Doubleday, 1966.
- [23] D. V. Lu, D. B. Allan, and W. D. Smart, "Tuning Cost Functions for Social Navigation," *Under review*, 2013.