



Evaluating directional cost models in navigation

Thibault Kruse, Alexandra Kirsch, Harmish Khambhaita, Rachid Alami

► To cite this version:

Thibault Kruse, Alexandra Kirsch, Harmish Khambhaita, Rachid Alami. Evaluating directional cost models in navigation. Proc. ACM/IEEE International Conference on Human-Robot Interaction (HRI), Mar 2014, Bielefeld, Germany. pp.350 - 357, 10.1145/2559636.2559662 . hal-01568842

HAL Id: hal-01568842

<https://laas.hal.science/hal-01568842>

Submitted on 25 Jul 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Evaluating Directional Cost Models in Navigation

Thibault Kruse
IAS, TU München
Boltzmann Str. 3
85748 Garching, Germany
kruset@in.tum.de

Alexandra Kirsch
Eberhard Karls Universität
Tübingen, Department of
Computer Science
Sand 14, 72076 Tübingen,
Germany
alexandra.kirsch@uni-
tuebingen.de

Harmish Khambhaita,
Rachid Alami
LAAS/CNRS
7 Avenue du Colonel Roche
31077 Toulouse Cedex 4,
France
{harmish.khambhaita,
rachid.alami}@laas.fr

ABSTRACT

A common approach to social distancing in robot navigation are spatial cost functions around humans that cause the robot to prefer paths that do not come too close to humans. However, in unpredictably dynamic scenarios, following such paths may produce robot behavior that appears confused. The concept of directional costs in cost functions [9] is supposed to alleviate this problem without incurring the problem of combinatorial explosions using temporal planning. With directional cost functions, a robot attempts to solve spatial conflicts by adjusting the velocity instead of the path, where possible. To complement results from simulations, in this paper we describe a user study we conducted with a PR2 robot and human participants to evaluate the new cost function type. The study shows that the real robot behavior is similar to the observations in simulation, and that participants rate the robot behavior less confusing with the adapted cost model. The study also shows other important behavior cues that can influence motion legibility.

1 Introduction

Navigation among humans must not just be safe, but also agreeable to humans. The concept of “legibility” describes one of the qualities beside safety that make robot behavior agreeable. Legibility expresses how easily an uninstructed person can estimate internal states of the robot from observation. This is important for interactions as it reduces distraction. Legibility is threatened by all robot behaviors that lead to wrong beliefs and expectations about the robot. The state-of-the-art in human-aware navigation planning does not take into account the concepts of legibility and cues, and can therefore produce confusing robot behavior.

Respecting a social distance around *standing* humans just means preferring to keep a minimum distance from the location of persons. But in dynamic situations the future paths of persons have to be taken into account for social distancing in path planning. To some degree, the future paths of humans can be predicted with various methods [16]. However the reliability of any prediction method quickly degrades in time, and there are necessarily many occasions of unexpected human motions, even more so when a robot moves nearby. Even with rather reliable prediction, replanning during execution is inevitable, requiring a planning process that scales well.

Our previous publication [9] presents an adapted cost model *ContextCost* that addresses those challenges. The

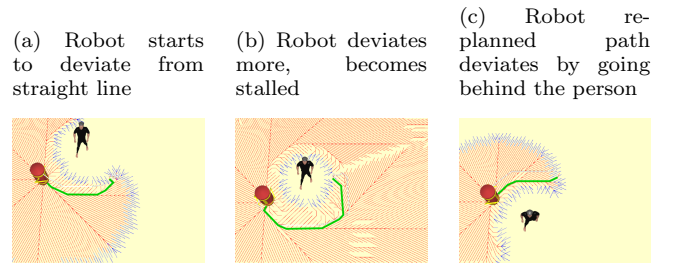


Figure 1: Effects of replanning with cost model *Static* in simulation. The robot is planning paths going from left to right, a simulated person walks from top to bottom without stopping for the robot. Taken from [9].

cost model *ContextCost* defines social costs for a path search in a grid depending on the motion direction of present humans and of the potential movement of the robot. In particular it may discard social costs for approaching humans when calculating a path, relying on the local planner to adjust robot velocity instead to solve a spatial conflict.

This cost model shifts the burden of respecting social distancing around moving humans from global planning to local planning and thus makes the robot follow more goal-directed paths in the presence of moving humans. We previously assumed this increases the legibility of the robot’s navigation intention, but we wanted to support this assumption by a user study.

Legibility in general is increased by explicit signaling and also any aspect of behavior that can be used as cue to the actual robot intention. Legibility is different from predictability, which requires an observer to imagine the most likely future state, determined by many other factors than the robot internal state [4]. As an easy reminder of this notion of legibility: For the task of reading text, legibility may relate to how easy it is to understand a given piece of text, whereas predictability could refer to how easy it is to predict the following text.

In robot motion, the robot motion direction and path are natural cues for the robot’s navigational intentions, as the natural way to move to a goal is moving forward on the most efficient path towards the goal, which is of course the straight line in the absence of obstacles. So legibility in navigation is naturally driven by the cues given by the robot’s direction of movement and path.




Figure 2: Robot (in blue) behavior with static cost models.

State-of-the art social navigation frameworks consider proxemics for social distancing, producing paths that prefer to avoid a circular region around any person in the environment. A path planner then finds a path that avoids this region. When the human is moving this region moves as well, so consecutive searches for a path can find different solutions.

Figure 1 shows an inevitable problem when navigating based on replanning with such a model of proxemic regions. The path planner finds a path to the goal of the robot while avoiding to approach the human moving from top to the bottom in the figure. As the situation evolves, the human position changes, and the robot has to replan to adapt. So, the robot plans very different paths during the approach of its goal, as the situation changes. The resulting directions and path of the robot are not helpful as cues to the robot intention. Figure 2 illustrates the resulting robot path in such situations.

Note that temporal planning alone does not generally fix this problem in realistic HRI. This is partly due to human motion being difficult to predict, in particular when being influenced by the robot motion. Also most temporal planning approaches can only be applied within boundaries. Such as the maximum look-ahead distance of a local planner like the Dynamic Window Approach (DWA) [6], where any conflicts at the edge of the boundary generate undesirable robot behavior.

Also in temporal planning, the planner has additional degrees of freedom to solve a conflict by using different velocities at different times. So it does not scale well, and requires an additional non-trivial cost function to balance optimization based on path length, path duration, and social factors like path legibility.

The context dependent cost model *ContextCost* has been compared in simulation to human motion strategies in 2-agent orthogonal crossing scenarios [1]. Encouraged by the results, this paper further evaluates the cost function in a real world setting when presenting experiment participants to a real world autonomously moving robot.

In particular we investigate whether the phenomenon can be reproduced on the PR2 robot when navigating with a present person, and whether our cost function improves the perceived quality of the motion of the PR2.

2 Related Work

The theory of proxemics describes social distancing behaviors of humans during interactions with other persons [7]. In recent years research in human-aware navigation has been based on theses proxemics rules for realizing more human like behavior during robot motion and non motion tasks. For detailed survey of human-aware navigation research we refer the reader to [10].

Sisbot et. al. [19] describes the human-aware navigation planner (HANP) that provides human-aware navigation plans considering proxemic distances around static humans in the environment. In extension to that work, we use a cost-map that adapts costs for moving humans based on the compatibility of motion directions, and we use an adapted local-planner that reduces velocity on a given path up to zero while waiting for moving humans to pass.

The general idea of considering moving humans not as obstacles to avoid, but instead considering the space around them as “soon to be free” is also exploited by Müller et al. [13], who make a robot follow persons going in the same direction rather than trying to find a way around them.

In [11] Kuderer et al. present an approach to predict future human trajectories based on a learned set of features that capture relevant characteristics of human trajectories in populated environment. Predictions of human trajectories are used as robot trajectories for social motion. Other researchers investigate temporal planning to solve spatial conflicts in very confined areas [15, 18]. However, these approaches do not pay special attention to legibility.

In [12] the authors have carried out a video-based user study comparing human-aware and non-human-aware navigation algorithms in situations where a human crosses the path of a robot. Users were shown first-person perspective videos of crossing situations with a robot using different navigation algorithms and asked to predict the robot motion and rate their own emotions during the interaction shown in the videos. Legibility was measured as a correlate of predictability.

Similar user-studies investigating optimal parameters for robots approaching standing persons have been conducted [2, 3, 22]. Sardar et al. [17] performed a user study where a robot approached a standing person concentrating on another task, and found significant differences in socially normative behavior for robots and humans. Within a holistic perspective on technology acceptance such as established by Heerink et al. [8], our work focuses on functional acceptance. The study by Pacchierotti et al. [14] also looks at a mutually dynamic situation of passing each other. Foka et al. [5] also did experiments using their POMDP-based navigation framework, but did not evaluate legibility.

3 Approach

The purpose of the user study was to evaluate the usefulness of the adapted cost model *ContextCost* in dynamic realworld situations over standard cost models with proxemics (model *Static*). Both cost models use a proxemic model for social distancing. They guide search for a path on a 2D grid by assigning costs to grid cells, such that a search algorithm can find the minimum connecting path between a start and a goal cell, which is assumed to be socially optimal.

In both cost models, for each human H the costs can be based on several different specialized functions f_i , such as for safety (as in figure 1) or visibility. For cost model *Static*, the costs for each grid cell can be computed using the function ζ'_{Static} :

$$\zeta'_{Static}(H, w_i) = \max(f_1(H, w_1), f_2(H, w_i), \dots, f_k(H, w_i))$$

Figure 3 illustrates the cost model *ContextCost* described in more detail in [9], which is based on the concept of spatial compatibility. When two paths are spatially compatible, the robot can stay on its intended (straight) path. Only in

H position of human
 \hat{hd} heading of human
 \hat{hd}' translated \hat{hd}
 \hat{hd}^{-1} inverse \hat{hd}'
 w_i i-th waypoint on path
 P closest point on human path to w_i
 d_P distance of P and w_i
 \hat{rd} potential heading of robot at w_i
 α angle between \hat{rd} and \hat{hd}^{-1}




Figure 3: Context dependent costs *ContextCost*, from [9].

incompatible situations, specific avoidance action has to be taken. Spatial compatibility is measured by the distance d_P between a waypoint and the estimated path of the human. The second measure is the difference in heading angles α , which we call directional compatibility. Using the angle α and distance d_P , we can define the incompatibility function ϕ that calculates a number $\in [0, 1]$ with which we can multiply the costs of the model *Static*. ϕ is defined such that when a path segment should not conflict with a predicted human motion, social costs do not apply.

$$\phi(H, w_i, w_{i-1}) = \begin{cases} 1 & , \text{if } \hat{hd} \text{ undefined} \\ 1 & , \text{if } d_P \leq d_{low} \\ 0 & , \text{if } d_P \geq d_{high} \\ 0 & , \text{if } \alpha \geq \alpha_{max} \\ \frac{d_P - d_{low}}{d_{high} - d_{low}} \cdot \frac{\alpha}{\alpha_{max}} & , \text{otherwise} \end{cases}$$

The arguments w_i, w_{i-1} to ϕ as opposed to only w_i underline why this cost function is called directional, it takes into account from which direction (which preceding cell w_{i-1}) the relevant cell is to be approached.

ϕ has three parameters d_{low} , d_{high} and α_{max} to tweak what motions are considered compatible. We use this function ϕ to modify the cost model *Static* by replacing ζ'_{Static} with $\zeta'_{ContextCost}$ as shown in equation below.

$$\zeta'_{ContextCost}(H, w_i, w_{i-1}) = \zeta'_{Static}(H, w_i) \cdot \phi(H, w_i, w_{i-1})$$

Thus robot path direction and human direction are considered incompatible if the human and robot could frontally run into each other. In all other cases, the motions are compatible, meaning it is likely the situation can work out well if the robot just reduces its velocity on the path while the human is in the way. So in the situation shown in Figure 1, a robot using cost model *ContextCost* would not deviate from the straight line, but instead reduce its velocity.

For the experiment, the “move_base” navigation framework of the PR2 robot was modified. As a global planner for “move_base” we used HANP [19] to enable human-aware path planning, and invoke the global planner at a regular frequency of 10Hz. HANP was configured with one of two cost models $\zeta'_{ContextCost}$ and ζ'_{Static} as a parameter of the experiment trials.

As a local planner, for this experiment we used a waypoint-following algorithm without general obstacle avoidance that adapts the robot velocity to prevent collisions with perceived humans. The waypoint follower causes the robot to visit each waypoint of the global plan, by turning towards each waypoint in turn and then moving forward to reach it (as opposed to allowing sideways or backward

motions). The strategy of velocity adaptation was varied according to the global planning strategy. When using the simple proxemics based cost model *Static*, the local planner tried to move at maximum velocity on the path, except when a human was in front of it at 1m distance (center to center), in which case no forward motion was allowed. The latter is a measure to grant a minimum feeling of safety for human participants. Using the cost model *ContextCost*, the velocity was adapted by projecting the motions of the robot and the human into the future, and selecting a velocity that would in those projections not predict their mutual distance to decrease below 1.3m. This value was established in preliminary trial as yielding a distance that felt safe while keeping the robot close enough to establish a crossing conflict situation.




Figure 4: Idealized velocity profile scheme for 90° crossing situations with cost model *ContextCost*. v_{pref} indicates the preferred travel velocity, v_{stop} a velocity so low a robot is considered not traveling.

In preliminary trials performed with a different set of participants we faced two main challenges when controlling the robot velocity: “stopping late” and “restarting early”. Figure 4 shows a velocity profile we imagine to be optimal for crossing the path of a human when giving way. It describes the velocity of moving forward, not rotations. There are three distinct phases for braking, waiting, and restarting. Ideally, there is no acceleration during the “Brake” phase, and no significant acceleration during the “Wait” phase. So the ideal robot behavior cannot merely be described in terms of distance and velocity, but acceleration (and possibly jerk) can have significant effect on humans nearby.

The phases should be clearly observable by the human whose path is being crossed, and the timing of the phases can be crucial to human comfort. In preliminary trials, the robot would sometimes initiate the braking phase too late, sometimes re-accelerate slightly during braking and waiting phases, or initiate the “Restart” phase too early. Each case caused some participants to feel uncomfortable, and the discomfort was strong enough to significantly influence the participants’ answers.

Figure 5a shows the transition to the “Brake” phase. When the robot decelerated late or had the slightest acceleration motion, participants felt the situation was uncomfortable. We believe this is because participants fear the robot to accelerate to full speed at the smallest symptom of acceleration. Since the short-time robot acceleration was caused by noisy data and prediction, we suppressed it by adding uncertainty to the perception data, i.e. by predicting low, nominal and high human velocities, and selecting a velocity that avoids collisions with all three prediction models.

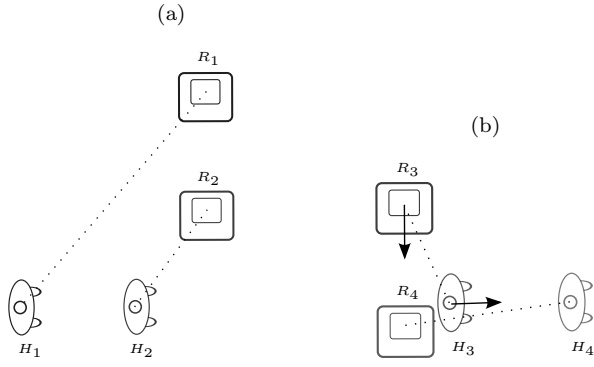


Figure 5: Situations relevant for comfort while crossing with adapted cost model *ContextCost*. R_t and H_t are position of the robot and human at time-step t respectively. Stopping late in (a) or starting early in (b) cause discomfort.

Cost model:	<i>Static</i>	<i>ContextCost</i>
Path:	Changing paths	Straight line
Velocity:	Maximum possible speed without collision	Reduce speed to keep distance

Table 1: Comparison of experiment conditions

This causes the robot to select low velocities more often, and makes participants less nervous. Figure 5b shows the “Restart” phase. Sometimes the participants felt uncomfortable expressing the robot moved “too early”. We believe again that the participants fear the robot quickly accelerating to full speed, so we corrected this behavior by projecting the robot position not at the desired speed, but at full speed, which caused the robot to accelerate later, and did not discomfort the participants anymore. Using both increased uncertainty about the future human positions and projecting the robot at full speed, the robot behavior in the crossing situation became rather similar to the scheme in Figure 4.

The differences of the two strategies used during the experiments are summarized in table 1. In the cost model *Static*, the robot attempts to move always at the maximum allowed speed and adapting the path to not cross the current human position nor an area in front of the human. The path frequently changes as the human moves, causing the robot to swerve. Using the cost model *ContextCost*, the robot attempts to follow a straight path to the goal, but reduces its velocity before reaching the crossing point, and accelerates again once the human has advanced enough.

The robot head was moving during the experiments to avoid the impression that the robot is ignoring the human. The control scheme turned the robot head to point at the participant head whenever the robot head could point towards the participant without deviating more than 90° from the frontal position, and vice versa for the participant. Otherwise the robot head pointed straight forward. When moving, the head moved slowly at 0.3 rad/s . Based on participant remarks this appeared appropriate except in rare cases when the robot was very close and moved in front of the participant, in which case the gaze seemed a bit provocative, which is similar to other results in robot gaze studies [20]. The participants were informed before the experiment that

the head will move, to prevent the discovery of this from distracting. The robot arms remained folded during all the experiment in the default PR2 “tucked” position. Only two participants mentioned this afterward as a bit unnatural, as opposed to making prompting gestures.

4 Evaluation

The user study was set up to determine whether the robot behavior with simple static cost models in crossing situations can be reproduced in the real world, and whether the context-dependent cost model visibly improves the robot motion quality.

Legibility is difficult to quantify, as a robot may behave very legible with respect to one aspect of internal state, but at the same time obfuscate another aspect of internal state. Therefore the experiment instead establishes a good knowledge of the robot internal state for the participants, showing how the participants can become confused by the robot behavior under assumptions about the robot internal state.

17 participants were given the task to act as interfering walkers with a robot moving from one spot to another in an area without static obstacles. The human participants were tracked using an infrared motion capture system, based on four passive infrared markers placed on a lightweight helmet the participants wore and 10 wall-mounted cameras, as shown in Figure 7. The room has an experimental area with offices and desks surrounding it, during the experiment people were quietly working around the experiment area.

The participants where all given the same set of written instructions to read. The instructions stated that multiple different robot strategies would be used, but not how many or what variations to expect.

The procedure was the following: In the preparation stage, the robot and the participant moved independently to their starting positions. The starting position of the participant was marked on the floor. This situation can be seen in Figure 7. When both were ready, the experiment instructor pressed a button to start the execution phase. This triggered the robot motion planning and also made the robot say “go”. The participant and the robot would start moving towards their respective goal positions at roughly the same time. The map in Figure 6 shows the coordinates. The participant goal was a small shelf on which lay a questionnaire. Going towards a shelf for a purpose rather than going onto a marked spot seemed to benefit a natural walking behavior. The goal area of the participant was slightly elevated (10 cm), and there was a low barricade with a gate spanning 1.5 m, seen in white in Figure 7. This was an existing feature of the room. The participants started their movement at coordinates (0.5, -3.5) on the map, but would not be visible to the motion capture system before reaching roughly (2.5, 3.5). Again we decided to maximize the walking distance because we were mostly interested in ratings about the robot, not human gait.

When the participants reached the shelf, they could then immediately rate the robot behavior on the questionnaire placed on the shelf. Once the participant and the robot had both reached their goal positions, the execution phase ended. Then a new cycle would begin with the preparation phase. An instructor was present and visible during the experiment, sitting at the desk shown in Figure 6, to react in cases of emergencies or disturbances, but also to reassure the participant and verify they were following the procedure. Also




Figure 6: Top-Down view on the room layout, showing the experiment path, the instructor’s desk, and the shelf where the participants marked their answer after each trial




Figure 7: Experiment setup showing wall-mounted cameras (2 out of 10), participant at start position with helmet and passive marker (1 out of 4), robot at start position, the intended paths crossing at 90°

the instructor modified the robot maximum speed between 0.45 m/s and 0.55 m/s according to the participant’s walking speed during the first 3 trials to create a spatial conflict without imposing a walking speed on the participants. The robot would approach alternately from the left or the right hand side.

Each participant was asked to perform the same task 13 times, and was told that the first 3 cycles would not be taken into account, but serve to get used to the situation. The robot strategy was randomized but never the same three times in a row. All participants were presented with both strategies to allow them a comparative opinion.

As opposed to studies on human gait, the experiment was interactive, meaning the robot behavior is a function of its algorithms, the person’s behavior (in particular the walking speed), the adapted robot maximum velocity and of noisy technical side effects like varying WiFi network latencies. Because of that, the captured data of the human trajectories is only of limited usefulness to interpret exact motion patterns.

- Number of Participants: 2 female, 15 male
- Age range: 22 - 34

- Mother language: 11 French, 6 Other
- Education level: 11 Master, 5 PhD
- Robot experience: 11 None, 6 with less than 6 years

After each run, the participants were asked to rate the robot’s performance according to predefined questions:

- Please rate the robot behavior (clear vs. confusing)
- Please rate the crossing situation (comfort):

The participants were given a Semantic Differential Scale from 1 to 5, the extremes labeled “clear”, “confusing”, “comfortable”, “uncomfortable” respectively. A preliminary trial had asked the participants to rate their “surprise”, but this word also had positive connotations and was also rated for the novelty of the whole situation. “clear” vs. “confusing” seemed better to express legibility issues.

Based on the design of the cost model *ContextCost*, the hypotheses for the outcome of the experiment were the following:

- H1: The robot behavior observed in simulation for *Static* also occurs in the experiment
- H2: Participants rate path adaptation as more confusing than velocity adaptation
- H3: Participants rate path adaptation as more uncomfortable than velocity adaptation

The alternatives are trivially negations of the given hypothesis. H2 and H3 would be confirmed if significant difference between the reported ratings were found.

5 Results

We collected data from 170 valid trial runs as shown in Table 2. Figure 8 shows summary plots of a representative run using cost model *Static*, and Figure 9 shows the same plots for a representative run using *ContextCost*. Figures 8a and 9a show the positions of robot and human over time for two representative samples we selected. Since we could not measure the human position from the start, we added the starting position as a single circle in the plot. Figures 8a and 8f shows how the robot at a distance of 1.2m slightly deviates from the straight line towards its left and rotates. The human passes during that time. Once the human has passed, the robot plans paths behind the human as shown in figure Figure 8b, thus turning to its right before eventually regaining the straight path. Compare that to Figures 9a and 9f, where the robot also decelerated to a minimal distance of 1.2m, but never rotated, and never planned anything but a straight path as shown in Figure 9b. This difference was already explained in [9], so here we just point out that the same occurs in a real world experiment.

In Figure 9d at time 3.8 there is a small increase of velocity. we could not pin down the cause of this, which could be with the prediction algorithm or robot wheel control. We point it out only to show that such a small and short acceleration already counted as a cue to some participants that the robot was maybe re-accelerating.

Figure 10 and 11 summarizes participant’s answers to the questionnaire. Each bar in the figure represents the number of participants who answered with the value on the horizon-

	<i>Static</i>	<i>ContextCost</i>
robot start at (6, -0.5)	39	45
robot start at (6, -6.5)	41	45

Table 2: Conditions for the 170 valid trials.

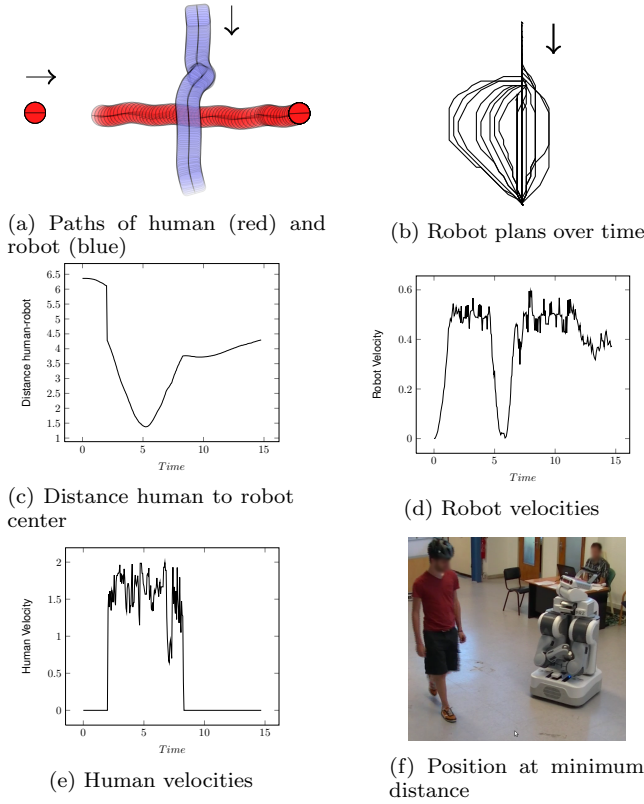


Figure 8: Sample 1 with cost model *Static*. In (a) the round positions are the human moving from left to right. (b) shows what the diverse plans the robot generate while moving.

tal axes for respective question. Figure 10 shows the differences in reported discomfort. The visible difference as given by a one-way ANOVA indicates trials with cost model *ContextCost* were rated more comfortable ($M=1.5$, $SD=0.723$) than for *Static* ($M=2.873$, $SD=1.265$), $p < 0.001$. The same is shown for behavior clarity in Figure 11. The trials with cost model *ContextCost* were rated more clear ($M=1.522$, $SD=0.707$) than for *Static* ($M=2.684$, $SD=1.215$), $p < 0.001$. The ratings for clarity and discomfort were also correlated as revealed by Pearson's R test ($R=0.771$, $p<0.001$). In their remarks, some participants usually argued that both values were independent. They explained that they rated strangeness whenever the robot orientation changed (as it only did with cost model *Static*), while they rated discomfort high when the robot passed first (only happened with *Static*) or had jerky acceleration motions (only happened with *ContextCost*). However, given the strong correlation possibly most participants did not clearly separate both concepts.

Hypothesis H1 is validated by the robot path behavior we saw, as displayed in Figures 8 and 12, in comparison to Figure 9a. The distinct responses for clarity and comfort validate Hypotheses H2 and H3.

While the figures for strategy *ContextCost* are very similar over all 90 runs, there was a lot of variation over the runs with *Static*. This was due to the different exact circumstances of crossing, which depended on the motion of the participant relative to the robot. To show some of the variety we observed, Figure 12 shows the paths for additional

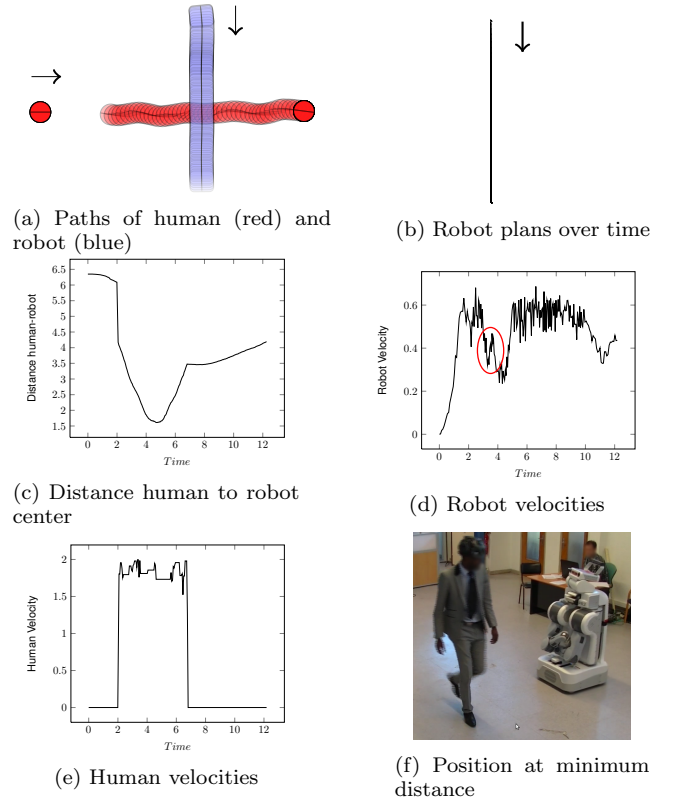


Figure 9: Sample 2 with cost model *ContextCost*. Same plots as in Figure 8.




Figure 10: Participant ratings of situation: (1=Comfortable, 5=Uncomfortable).




Figure 11: Participant ratings of robot behavior: (1=Clear, 5=Strange).

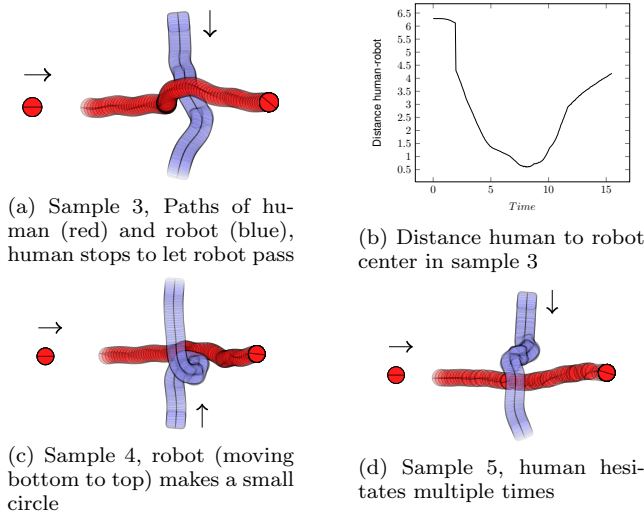


Figure 12: Additional Examples with cost model *Static*, showing variants of robot behavior reacting to different human behaviors

samples. In Figure 12a the participant let the robot move first. This happened 20 out of 80 times with cost model *Static*, but never with *ContextCost*. The minimal distance of 0.5m was a result of the participant passing closely behind the robot, and such low distances were consistently observed when participants let the robot pass first. Figure 12c shows a more extreme robot motion deviation as the robot retraced its step following a path that led away from the human. Finally Figure 12d shows a sample where the participant hesitated several times before passing the robot.

No statistically significant differences were found comparing the robot approaching from the right or left-hand side of the human. The observed human average walking velocities mainly varied between 1.5 m/s and 1.7 m/s with no clear trend detected over trials. While Figures 9e and 8e appear different, no pattern seemed evidence over all 170 runs.

6 Discussion

The charts in Figure 10 and 11 still show some degree of variation over the robot clarity of motion and discomfort. This is partially due to the fact that some participants preferred to use the full answer scale while others used only low numbers. However from the remarks we also know that even with the cost model *ContextCost*, there were symptoms the participants disliked. This is mainly due to remaining confusing symptoms in the robot acceleration behavior. As the robot velocity depends on a linear projection of the human motion and the robot, and the sensory data is noisy, the observable robot behavior for slowing down had tiny moments of acceleration, which were very visible to participants and caused mild confusion and discomfort.

As our first attempt at the experiment failed because robot acceleration was uncomfortable, we learned that three trial runs for warm-up may not generally be sufficient for this kind of study, for non-technical people.

Reported discomfort by participants seemed to always be related to uncertainty about future collisions. Participants felt generally disturbed when the robot base moved (or rotated) while the robot appeared to be “waiting for them to

pass”. Legibility of robot motion may be a quality that only becomes relevant once humans feel sufficiently reassured about their safety in the presence of a robot. Most participants chose a very straight line to their goal, as opposed to curving their path to accommodate the robot, very few participants however consistently walked on paths curved away from the robot approach direction. Two Participants reported that when the robot base rotated (in the case of cost model *Static*), they felt as if the robot tried to follow them or join them, three participants felt that the robot was trying to cut their way in such cases. Base rotations were reported as uncomfortable when it happened very close to the participants (due to the square shape of the robot, a rotation could also cause a collision with the person’s legs), but when it was at least 1.5 m away, no discomfort was reported.

The robot passing first (with cost model *Static* and the participant walking slowly) as in Figure 12a was accepted by some participants, others felt the robot should have stopped for them in any case. In such cases it is also notable that human participants pass very close behind the robot, almost touching the robot, so there was no visible notion of social self-distancing from the robot. This could mean that social distancing is generally not required for such situations because the duration where two agents are very close is also very short. However social distancing is likely to still be relevant for situations of following one another, or walking side by side, when the duration of proximity is longer. The observation that humans pass very close behind crossing robots, but do not like being passed close by the same robot in the inverse situation, hints at a difference in socially normative behaviors for humans and robots as also found in [17].

As additional remark on predictability, using the adapted cost model *ContextCost*, the robot behaved very similar for each run, even when participants walked at different speeds or hesitated. This consistency made the robot behavior also very predictable, as consistency is generally a factor increasing predictability. The predictability of the whole situation for third party observers may also be considered, in 20 of 80 runs with *Static*, the participant let the robot go first, but never did so in 90 runs with *ContextCost*, meaning the latter cost model leads to better predictability of the situation as a whole.

The study does not validate this approach for other conflict situations such as crossing at different angles or following at different speeds. Such studies would require a larger experiment area than we had available. Also the study does not compare the straight path behavior of a robot to paths curved “behind” the crossing human, which requires temporal planning to be found by a path planner.

It is also worth considering how the navigation approach scales to situations with more humans. While the robot is prone to the “freezing robot” problem [21], because in case of doubt it reduces velocity to zero, we believe for moderately crowded situations this behavior may still help to avoid making a complex situation even more complex, given the insecurity our participants reported to even slight changes in robot velocity. For densely crowded situations, a fully reactive navigation mode seems preferable.

The study has shown several cues that human observers use to guess the intentions of the robot. While the focus was on the path behavior of the robot, we also saw that minor accelerations were also used as a significant cue by humans, so those have to be avoided to prevent confusion. Possibly

the robot might require a social commitment behavior such that it commits to stopping or moving for at least a given time-span of one or two seconds. Cues like hectic motion of PR2 caster wheels during low velocity, or starting of robot cooling vents may not clearly indicate a specific robot intent, but are confusing as cues.

7 Conclusions

The user study was performed to validate the improvement of perceived robot motion quality when using an adapted cost model *ContextCost* for path planning. The study shows that using a real robot and uninstructed participants, a robot will show confusing behavior when using a simple proxemic path planner, behavior that humans will rate as strange or confusing. For crossing situations, the study shows that if a robot instead discards proxemic costs in the path planner and chooses a direct line to the goal, but also uses a local planner that adapts velocity to maintain a certain distance, the robot behavior becomes more stable, and humans rate the robot behavior as clearer. The study also shows that for such situations of conflict, robot acceleration is an important source of confusing robot cues, even if it happens for very short times and at small amounts. The study hints at the possibility that for short durations, a close encounter may not require social distancing at all, but safety distancing. Also the study has revealed that for legibility, it is not only crucial to provide useful cues to the robot internal state (such as its intention), but to also avoid accidentally creating misleading cues. This leads to the notion of “perspective taking” in path planning and local planning, such that a robot needs not only to reason about the effects of a *planned* trajectory, but also about the effects of the *imaginary trajectories* that observers predict for the robot, based on cues. So the implications of this user study on legibility extend to all other motion planning problems involving the presence of humans.

8 Acknowledgments

With the support of the Technische Universität München - Institute for Advanced Study and the CoTeSys (Cognition for Technical Systems) cluster of excellence and the Bavarian Academy of Sciences and Humanities. This work has partly been supported by the European Commission under contract number FP7-ICT-600877 (SPENCER).

9 References

- [1] P. Basili, M. Saglam, T. Kruse, M. Huber, A. Kirsch, and S. Glasauer. Strategies of locomotor collision avoidance. *Gait & Posture*, 37(3):385–390, 2013.
- [2] J. T. Butler and A. Agah. Psychological effects of behavior patterns of a mobile personal robot. *Autonomous Robots*, 10(2):185–202, 2001.
- [3] D. Carton, A. Turnwald, D. Wollherr, and M. Buss. Proactively approaching pedestrians with an autonomous mobile robot in urban environments. In *13th International Symposium on Experimental Robotics (ISER)*, 2012.
- [4] A. Dragan, K. Lee, and S. Srinivasa. Legibility and predictability of robot motion. In *Human-Robot Interaction*, March 2013.
- [5] A. F. Foka and P. E. Trahanias. Probabilistic autonomous robot navigation in dynamic environments with human motion prediction. *I. J. Social Robotics*, 2(1):79–94, 2010.
- [6] D. Fox, W. Burgard, and S. Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics and Automation Magazine*, 1997.
- [7] E. Hall. *The hidden dimension*. Anchor Books, 1966.
- [8] M. Heerink, B. J. A. Kröse, V. Evers, and B. J. Wielinga. Assessing acceptance of assistive social agent technology by older adults: the almere model. *I. J. Social Robotics*, 2(4):361–375, 2010.
- [9] T. Kruse, P. Basili, S. Glasauer, and A. Kirsch. Legible robot navigation in the proximity of moving humans. In *Advanced Robotics and its Social Impacts (ARSO), 2012 IEEE Workshop on*, pages 83–88, May 2012.
- [10] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch. Human-aware robot navigation: A survey. *Robotics and Autonomous Systems*, 61(12):1726 – 1743, 2013.
- [11] M. Kuderer, H. Kretzschmar, C. Sprunk, and W. Burgard. Feature-based prediction of trajectories for socially compliant navigation. In *Proceedings of Robotics: Science and Systems*, Sydney, Australia, July 2012.
- [12] C. Lichtenthäler, T. Lorenz, and A. Kirsch. Influence of legibility on perceived safety in a virtual human-robot path crossing task. In *RO-MAN, IEEE*, 2012.
- [13] J. Müller, C. Stachniss, K. O. Arras, and W. Burgard. Socially inspired motion planning for mobile robots in populated environments. In *International Conference on Cognitive Systems (CogSys’08)*, Karlsruhe, Germany, 2008.
- [14] E. Pacchierotti, H. Christensen, and P. Jensfelt. Evaluation of passing distance for social robots. In *RO-MAN, IEEE*, September 2006.
- [15] M. Phillips and M. Likhachev. Sipp: Safe interval path planning for dynamic environments. In *ICRA, IEEE*, 2011.
- [16] J. Rios-Martinez. *Socially-Aware Robot Navigation: combining Risk Assessment and Social Conventions*. PhD thesis, Université de Grenoble, 2012.
- [17] A. Sardar, M. Joosse, A. Weiss, and V. Evers. Don’t stand so close to me: users’ attitudinal and behavioral responses to personal space invasion by robots. In *HRI*, pages 229–230, 2012.
- [18] L. Scandolo and T. Fraichard. An anthropomorphic navigation scheme for dynamic scenarios. In *ICRA, IEEE*, February 2011.
- [19] E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Simeon. A human aware mobile robot motion planner. *IEEE Transactions on Robotics*, 23, 2007.
- [20] L. Takayama and C. Pantofaru. Influences on proxemic behaviors in human-robot interaction. In *IROS, IEEE/RSJ*, 2009.
- [21] P. Trautman and A. Krause. Unfreezing the robot: Navigation in dense, interacting crowds. In *IROS, IEEE/RSJ*, October 2010.
- [22] M. Walters, M. Oskoei, D. Syrdal, and K. Dautenhahn. A long-term human-robot proxemic study. In *RO-MAN, IEEE*, 2011.