Human Aware Navigation for Assistive Robotics

Dizan Vasquez, Procópio Stein, Jorge Rios-Martinez, Arturo Escobedo, Anne Spalanzani and Christian Laugier

Keywords: Human Aware Robotics, Assistive Robotics, Autonomous Navigation, Human Behavior Modeling

Abstract Ensuring proper living conditions for an ever growing number of elderly people is a significative challenge for many countries. The difficulty and cost of hiring and training specialized personnel has fostered research in assistive robotics as a viable alternative. In this context, an ideally suited and very relevant application is to transport people with reduced mobility. This may involve either autonomous or semi-autonomous transportation devices such as cars and wheelchairs.

For a working solution, a number of problems including safety, usability and economic feasibility have to be solved. This paper presents PAL's robotic wheelchair, an experimental platform to study and provide solutions to many of the aforementioned problems.

Dizan Vasquez
INRIA, Université Pierre Mendès France, LIG. e-mail: `dichodaemon@gmail.com`

Procópio Stein
INRIA, DEM - Universidade de Aveiro. e-mail: `procopiostein@gmail.com`

Jorge Rios-Martinez
INRIA, LIG. e-mail: `jorge.rios-martinez@inria.fr`

Arturo Escobedo
INRIA, LIG. e-mail: `jesus.escobedo-cabello@inria.fr`

Anne Spalanzani
Université Pierre Mendès France, INRIA, LIG. e-mail: `anne.spalanzani@inria.fr`

Christian Laugier
INRIA, LIG. e-mail: `christian.laugier@inria.fr`

Acknowledgements Procópio Stein and Arturo Escobedo are funded by the INRIA large initiative scale called PAL (Personal Assisted Living) `pal.inria.fr`. Procópio Stein is also funded by FCT. Jorge Rios-Martinez is also funded by CONACYT 250140/308006. Dizan Vasquez was funded by UPMF as invited Professor.

1 Motivation, problem statement and related work

Ensuring proper living conditions for an ever growing number of elderly people is a significative challenge for many countries. The difficulty and cost of hiring and training specialized personnel has fostered research in assistive robotics as a viable alternative. In this context, an ideally suited and very relevant application is to transport people with reduced mobility.

In particular, this paper studies the case of a robotic wheelchair. For such a system, it is crucial to take into account the actual needs and characteristics of both its users and the people around them. The platform discussed in this paper has been designed around the following requirements:

- *Safety*: The system should avoid collisions with both static and dynamic entities.
- *Usability*: People with motor disabilities often have problems using joysticks and other standard control devices. The system should account for this, for example by favoring the most "reasonable" actions when presented with an ambiguous command.
- *Comfort*: Strong accelerations can be untolerable and even dangerous for a wheelchair user, this imposes an additional constraint on how the robot may move.
- *Respect of social conventions*: When moving, a robot may considerably disturb people around it, especially when its behavior is perceived as unsocial. Even worse, the wheelchair's passenger may be held responsible for that behavior. It is thus important to produce socially acceptable motion.

From the technical standpoint these requirements imply that, in addition to the conventional robot tasks (e.g. localization, path execution) the following points should be specifically addressed:

- *Integrated motion-planning and long-term motion prediction*: Most human-populated environments are highly dynamic, requiring considerable look-ahead about how other objects will move in order to ensure collision-free robot motion under "comfortable" accelerations. This motivates the proposed integration of a long-term motion prediction algorithm based on the idea of typical behavior with a risk-based motion planning algorithm.
- *Interaction detection for socially acceptable robot-motion*: Our approach is based on the simple idea that, when people interact, they often adopt spatial formations implicitly forming "interaction zones". Thus, socially acceptable motion can be enforced by first detecting interaction zones and then computing the risk to invade them.

One of our main ambitions with this platform is to provide an open benchmark that could be used to compare and evaluate different approaches. This is an important task given the diversity of existing wheelchairs [1], including autonomous [2], semi-autonomous [3] and social aware systems [4, 5].

2 Technical Approach

Fig. 1 presents an overview of our system's architecture. It is divided into several subsystems:

1. *Tracking subsystem*: mobile objects are tracked both off-board and on-board the robotic wheelchair.
2. *Prediction subsystem*: the prediction subsystem processes data from the trackers and transforms it into probabilistic predictions about the configuration of the free space in the environment. It also features a "social filter", which detects present and future interactions and creates virtual obstacles corresponding to interaction zones.
3. *Navigation subsystem*: the navigation subsystem includes a laser-based localization module and a motion-planner which integrate predictions to compute safe trajectories that are fed to the execution module.

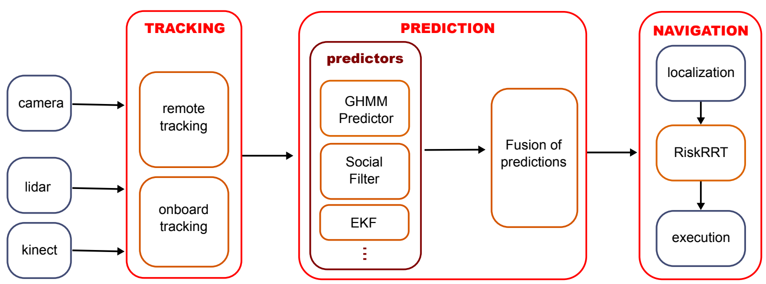

Fig. 1 Achitecture overview.

2.1 Tracking systems

The off-board tracker provides global information about moving obstacles and provides learning input for our motion prediction module.

At this point, we are still developing and testing our tracking systems. Meanwhile, we have performed several tests using augmented reality markers that people wear as hats. This has allowed us to validate the overall architecture, even if it is not a viable solution in the long run.

For the definitive version of the platform, we are working on a basic detect-then-track system, where moving objects are first detected using a Self-organizing network [6], after this, objects are encoded as a color histogram, and then detected in later frames using the mean-shift algorithm [7].

Finally, the different detections are used as input for a tracker based on the Joint Probabilistic data Association Filter.

On the other hand, the on-board system will provide detailed information about the objects that appear in the robot's perceptual field. Its main use is to identify interactions between people (e.g. two persons shaking hands). The on-board tracking performs leg detection using a LIDAR sensor and people detection using the kinect sensor, according to the technique described in [8].

2.2 Motion prediction

The motion prediction subsystem takes tracking data (i.e. position, orientation and velocity) and outputs K grids, representing the posterior probability of the space being occupied at times $\{t_1, \cdots, t_K\}$ in the future. Prediction itself is accomplished with a Growing Hidden Markov Model [9] and an Extended Kalman Filter but the grid representation makes it easy to experiment with other prediction algorithms. The prediction grids are then processed by a fusion module, which currently performs bayesian sensor fusion as described in [10].

In order to produce socially acceptable motion, we have proposed the "Social Filter", which integrates constraints inspired by social conventions in order to evaluate the risk of disturbance represented by navigation decisions. We focus on detecting and predicting conversations in the environment surrounding the wheelchair [5].

2.3 Navigation

Our navigation system is based on Risk-RRT [11], a partial motion planner which integrates motion predictions to provide safe trajectories. We have also extended the approach by including a mechanism to obtain socially acceptable behavior.

When the wheelchair is transporting a human, it will have to move in a populated environment where an "optimal" behavior may be perceived as unsocial. People will become uncomfortable if they are approached at a distance that is deemed to be too close, where the level of discomfort experienced by the person is related to the importance of his or her space. This simple idea was formalized introducing the concept of *personal space*, first proposed by Hall [12], which characterizes the space around a human being in terms of comfort to social activity.

Another interesting social situation arises when two or more of the persons in the environment are interacting. We model interactions using the concept of *o-space* which has been developed by sociologists [13]. This space can be observed in casual conversations among people where participants' position and orientation are used to establish boundaries of the space. This space is respected by other people and only participants are allowed to access to it, therefore the intrusion of a stranger causes discomfort. In our path planner, human friendly paths are generated by including a "Social Filter" which transforms those spaces into corresponding cost functions which lead the robot to avoid them. As a result, the choice of a best path done by RiskRRT is based on the "probability of success" calculated for every partial path considering the probability of not encountering a collision along the path and not entering in a personal space or an o-space [5].

2.3.1 Modeling Personal Space

We have modeled personal space as a mixture of two gaussians human centered, one for the front and one for the back of the space, the front is larger as people is more sensitive to this space. Fig. 2 shows an example of personal space as provided by the Social Filter.

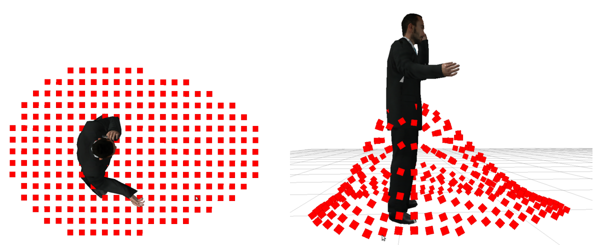

Fig. 2 Personal space calculated by Social Filter Module. Height of the gaussian means Risk of disturbance then maximum disturbance is located at human position.

2.3.2 Modeling o-Space

When more than two people are in conversation, they tend to make a formation with circular shape. The o-space could be taken as a circle whose center coincides with that of the inner space. For the specific case of two people, some formations, called F-formations, have been identified as being particularly frequent [13]. The social filter identifies individual F-formations (Vis-a-vis, L-Shape, C-Shape or V-Shape) and builds the corresponding o-space. in Fig. 3, the calculated o-space for a Vis-a-Vis interaction is shown.

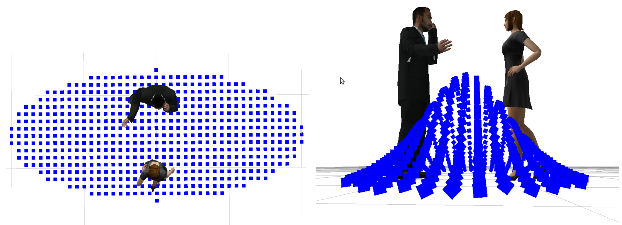

Fig. 3 O-space calculated by Social Filter Module for a Vis-a-Vis F-formation. Maximum risk of disturbance is located at o-space center, in the picture the disturbance is represented by height of Gaussian.

3 Experimental Results

It is important to highlight that the proposed experimental platform is an ongoing effort. Thus, the results described below should be considered preliminary. We have conducted experiments both in simulation and with the real platform as described in § 3.1 and 3.2, respectively.

Before going into the details of our results, it is convenient to discuss the graphical elements we will use in our figures. In our tests, humans are represented by a 3D model of a man or woman (4a), red points are used to represent the personal space that should be avoided by the robot. Finally, colored squares in front of the human represent a simple estimation of future positions, with each color representing a different moment in time.

The wheelchair (Fig. 4(b) is represented by a 3D model of a wheelchair surrounded by rounded points that represent explored RiskRRT nodes. As in the case of people, different colors are associated with different moments in time. The size of the points represents the computed risk of navigation to that position, where larger points mean bigger risks. Finally, a red solid line is used to represent the path to be followed, with a blue arrow indicating the robot's goal.

3.1 Simulation

3.1.1 Test scenarios

The tests focused on two main functions: predictive navigation and socially acceptable navigation. In the first case, people interfered with the robot's plans by either following the same path than the robot in the opposite direction or intersecting it at some point. In both cases the robot had to anticipate the human trajectories and generate an alternative collision-free plan.

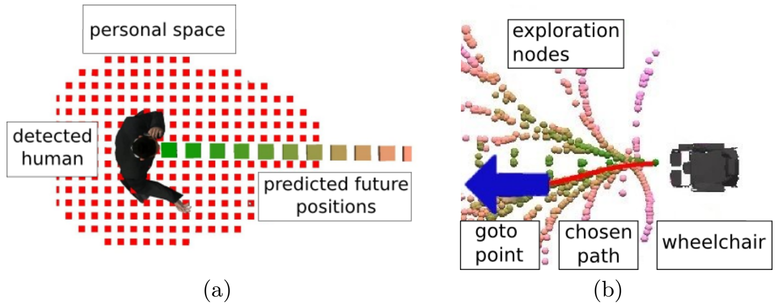

Fig. 4 (a) Human visualization and symbols meaning; (b) Robot visualization and symbols meaning

In the second case, we aimed to assess the capability of the robot to avoid disturbing or causing discomfort to persons that were not moving but were interacting with each other. People were arranged in a manner that the direct path to the robot's goal would be inside a social interaction zone, so a straight movement to the goal would violate the interaction zone and therefore, the robot had to find alternative paths.

3.1.2 Prediction and Navigation

We have conducted extensive tests of the RiskRRT algorithm in simulation. Fig. 5 shows one iteration of the navigation main loop. As it can be seen, the resulting trajectory differs from optimal trajectories obtained by traditional planning algorithms, the robot actually opts for a larger trajectory that avoids obstructing the moving pedestrians.

In all our simulations the speed of pedestrians has been fixed to one m/s and maximum speed of our wheelchair is also one m/s.

We have performed a number of tests to assess the effect of including prediction in our motion planning algorithm. Fig. 6 compares the paths that were obtained using predictions of pedestrian movements (left column) with those obtained without predictions (right column). The robot's initial position is on the left end of the corridor while the goal is at right end. Since the corridor is narrow, the only way to avoid colliding or disturbing the pedestrian is by moving aside in the corridor opening before continuing to the goal. In the figure, it is possible to see how, by using predictions, the wheelchair is capable to detect a possible collision in the middle of the corridor (6 a)) and to choose an alternative path to reach the goal. Without prediction it takes a straight path to reach the goal which, at first does not seem to pose any risk (6 b)) later, when the wheelchair detects the collision (6 d)) and tries to avoid the person, it is already too late.

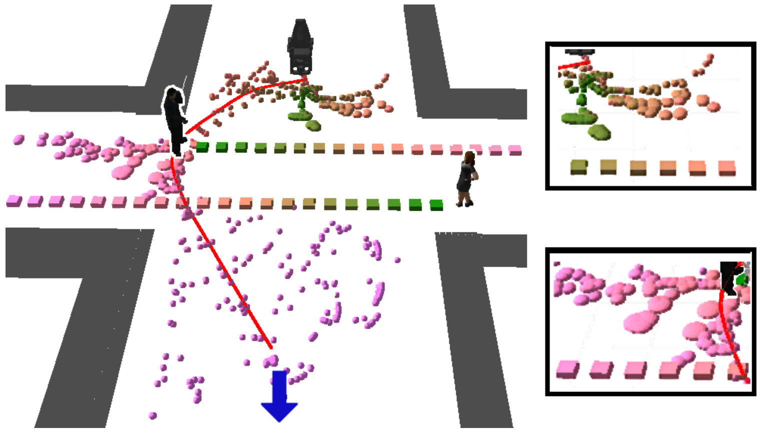

Fig. 5 Predictive navigation example. RiskRRT selected a plan (red line) to the goal (blue arrow). The chosen path leads the robot to pass by the back of the first person, and then reduces the speed to let the second person to pass as well. With this strategy, the robot minimizes the risk of collision and the discomfort caused for the two pedestrians. Once second person has passed, the algorithm choses a straighter path to the goal. Frames at the right of the figure show that estimated risk is bigger at future positions of the wheelchair (circles) which are close to predicted positions of pedestrians (squares).

3.1.3 Socially Acceptable Navigation

In order to test socially acceptable behavior, we conducted several simulation tests. Our first test scenario consisted of two interacting people, together with the wheelchair. We realized thirty executions of the planner in very similar conditions, as it can be seen in Fig. 7, when the social filter is off, the plans do avoid people but do not respect social space. When the social filter is turned on again, all the plans managed to respect interaction space without disturbing the involved people.

3.2 Real platform

3.2.1 Experimental platform

Our mobile platform (Fig. 8(b)) is a robotic wheelchair manufactured by BlueBotics for the European project MOVEMENT. It is built on a mobile base equipped with a SICK LMS-200 LIDAR, and a Microsoft Kinect RGBD camera. The wheelchair is also equipped with an on-board computer to take care of the low-level hardware control tasks, on top of that it also carries a notebook computer with the navigation, prediction and tracking software.

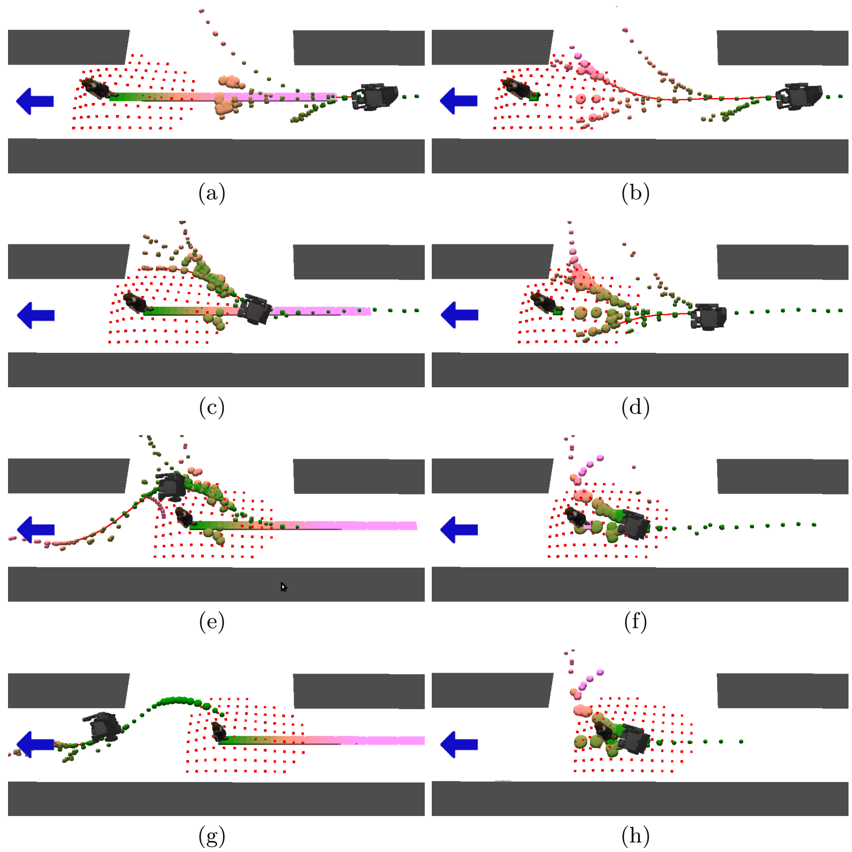

Fig. 6 Qualitative comparison of predictive navigation (first column) vs non predictive navigation (second column). Prediction helps to discover future high-risk states (a) and anticipate avoidance paths to finally reach the goal (g). Without prediction avoidance begins too late (f) and a collision is unavoidable (h).

In addition to the mobile platform, there is also an external camera (Fig. 8(a)) overlooking the test environment. It is connected to an external computer that communicates with the wheelchair via wireless network.

From the software perspective, the system has been implemented as a number of independent modules using the Robot Operating System (ROS) [14].

3.2.2 Motion Prediction

The proposed prediction algorithms has been extensively validated and compared about other state of the art techniques [9]. Our approach consistently

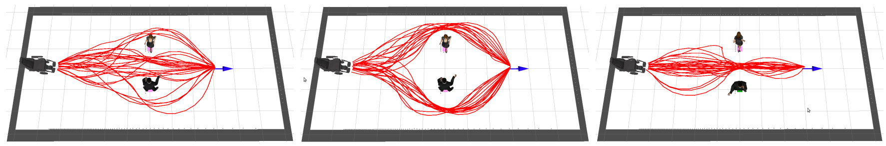

(a) (b) (c)

Fig. 7 Socially acceptable navigation. Each figure shows a sample of generated plans (in red) for thirthy executions of RiskRRT: a) without social filter social spaces are not respected, b) and c) with social filter, navigation is socially acceptable. In c) people are looking towards walls, therefore there is no social interacting zone, then navigation respects only their personal spaces.

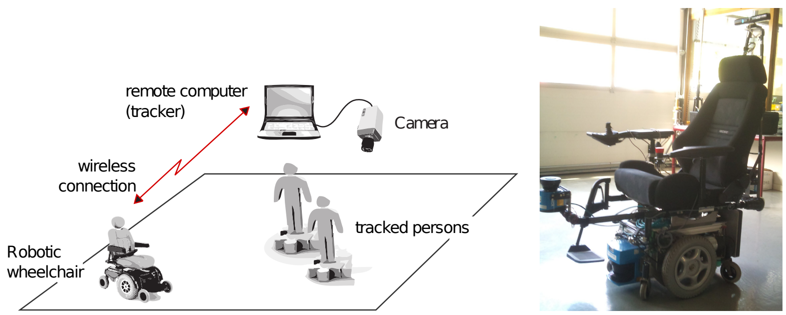

(a) Overview of the full experimental setting. (b) Robotic wheelchair.

Fig. 8 Experimental platform.

yields comparable predictions with much smaller models and is able to update its knowledge as new motion patterns are observed.

To validate the results obtained with our predictor, the scenario chosen to conduct the experiments is the main hall of INRIA Rhône Alpes (Fig. 9(a). It is an interesting choice as it has a large flow of people during different times of the day, entering and leaving the building during lunch hours and at the beginning and the end of a working day. These conditions provide a realistic and challenging place to conduct experiments on dynamic environments.

The GHMM has been trained using a set of 190 real trajectories. Volunteers were asked to move naturally among interest points in the environment, as the entrance of the hall, the two corridors and the two doors. Fig. 9(b) shows a sample of these trajectories, where the tree interest points located at the stairs illustrates the three separate paths that can be taken when climbing it.

A great advantage of the GHMM is it capability to automatically create, remove and merge redundant states while learning, which result in a more efficient training compared to classical HMM. Fig. 10 illustrates the learned states (represented by spheres) along the INRIA's hall.

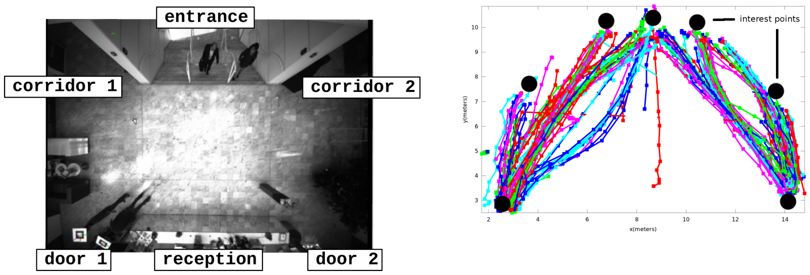

Fig. 9 (a) INRIA Rhône Alpes entrance hall. (b) Real trajectories used in the GHMM training.

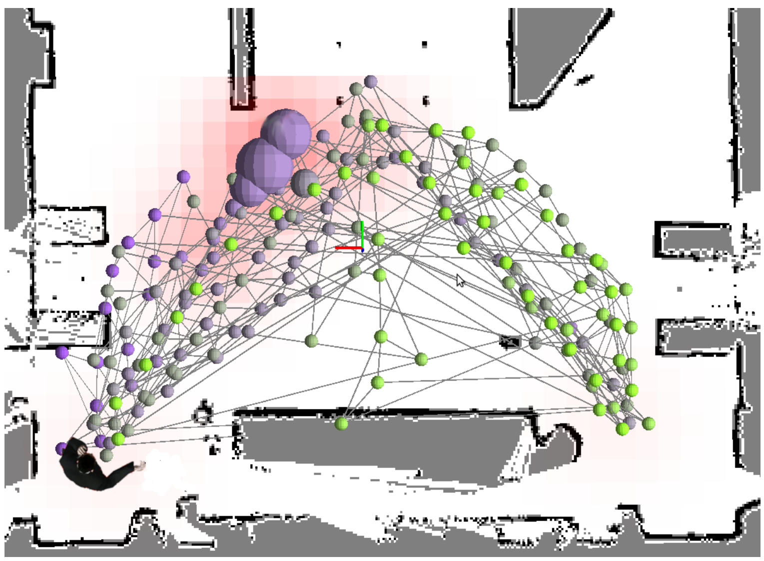

Fig. 10 GHMM learned states (represented by connected nodes) and the prediction of a goal for a person beginning to move from the left door (represented by larger nodes at the left portion of the stairs.

3.2.3 Socially acceptable navigation

Test were conducted in the INRIA hall, linking together the tracking, social filters and navigation modules, previously presented. The tracking module fed information to the social filter module which computed social interaction zones, according to the orientation and position of humans in the scene.

Fig. 11(a) shows one image of two persons interacting while the robot passes by, with a researcher closely following. Fig. 11(b) shows the same situation but taken from the overhanging camera linked to the tracker computer, where the robot position, its plan and intended destination can be seen.

(a) (b)

Fig. 11 (a) Experimental test with two interacting humans and a robotic chair navigating among them. (b) Overview camera image of the test scenario with the robot plan overlayed.

Several tests were conducted to evaluate the capability of the robot to avoid zones that would cause discomfort to the people interacting with each other. We also compared results with and without the social filter module, to demonstrate that not taking into account the zones of social interaction would result in paths that are shorter but "rude" or even frightening.

Fig. 12 shows the two experiments that were performed. The image shows roughly the same initial configuration for the robot and the interacting persons, as well as the same goal. The only difference is that, in the left column the social filter has been disabled while in the right one it is active, which is illustrated by the point cloud between persons.

Due to the absence of a social space, in the left column images, the planning algorithm treats the humans are simple obstacles, and the chosen path is the one that moves straight to the goal. However, when the social filter is active, nodes that are generated inside the interaction zone are penalized with a high risk, and then are excluded during the path search.

This example clearly shows that although a straight path to the goal can be considered to be more efficient in terms of energy and total distance that was traveled, it moves in such a way that it causes discomfort to interacting groups of people in the environment. On the other hand, the example shown in the right column, manages to avoid the zone of interaction, at the cost of traveling a longer distance.

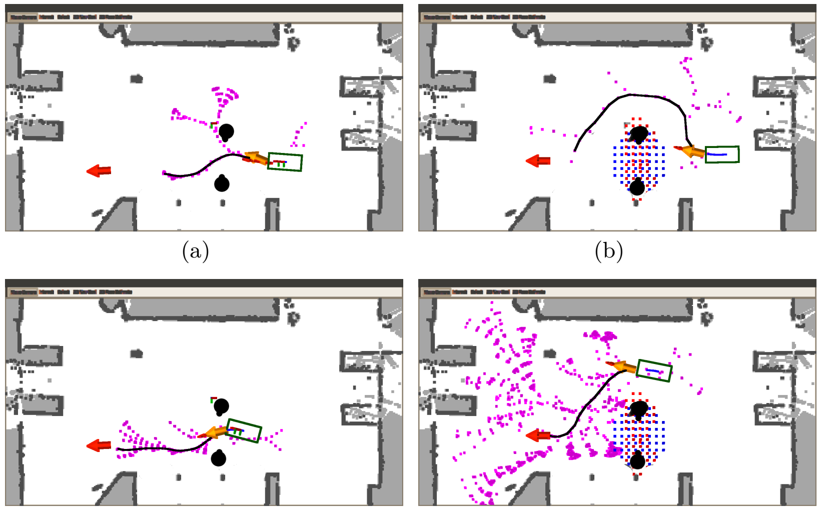

(a) (b) (c) (d)

Fig. 12 The robot is represented by a rectangle, the goal by the leftmost arrow, and interacting people by black circles. Images (a,c,e) show the social filter module deactivated, the resulting trajectory is shorter but socially unacceptable; (b,d,f) images shows a trajectory that is longer but respects the social interaction zones displayed as clouds of blue points.

4 Conclusions and future work

As we have mentioned above, the platform presented in this paper should be considered work in progress. Nevertheless, we consider that the results we have obtained until now are both relevant and promising and had been instructive in relation to several aspects of the problem at the application and the technical level:

- *Socially acceptable behavior is very important.* Even in our scripted tests, both interacting people and the wheelchair's user reported that they felt very uncomfortable when the robot passed right through the middle of a talking group.
- *Predictive behavior leads to socially acceptable behavior.* For example, when pedestrians were passing through the robot's path, it often happened that it stopped (knowing that the path was going to be free) to let the person pass. This seems to indicate that in many cases, knowing how people will move, the most reasonable thing to do is to be polite. It also suggests game theory as a possible way to analyze these interactions.

On the other hand, there are several open fundamental issues that need to be addressed, in particular, the problem of defining proper ways of evaluating comfort and social compliance has not been tackled here. The reason lies in the difficulty to put together experiments which really factor out all those variables that are not being studied. For example, during our experiments, we were applying questionnaires to the wheelchair passenger with very inconclusive results because the environment contained a flight of stairs going down. The result was that people were too frightened about the wheelchair falling there to be able to consider social discomfort.

As future work we plan to have the help of sociologists to aid in the formulation of questionnaires that can better capture the variables we want to study, as the comfort, for example. We also noticed that the reduced size of the useful space of our test environment (approximately $70m^2$) posed limitations to the variety of tests that we could perform. So future tests will be conducted in a larger environment, free of risk factors (as stairs), with a larger number of humans and more free space for the robot to maneuver, so we can better explore the limitations and advantages of our techniques.

References

[1] Richard C. Simpson. Smart wheelchairs: A literature review. *The Journal of Rehabilitation Research and Development*, 42(4):423, 2005.

[2] Thomas Kollar, Stefanie Tellex, Deb Roy, and Nicholas Roy. Grounding Verbs of Motion in Natural Language Commands to Robots. In *International Symposium on Experimental Robotics*, 2010.

[3] Sarangi P. Parikh, Valdir Grassi Jr, Vijay Kumar, and Jun Okamoto Jr. Usability study of a control framework for an intelligent wheelchair. In *IEEE International Conference on Robotics and Automation*, page 4745–4750, 2005.

[4] Rachel Kirby, Reid Simmons, and Jodi Forlizzi. Companion: A constraint-optimizing method for person acceptable navigation. *The 18th IEEE International Symposium on Robot and Human Interactive Communication*, 2009.

[5] Jorge Rios-Martinez, Anne Spalanzani, and Christian Laugier. Understanding human interaction for probabilistic autonomous navigation using Risk-RRT approach. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2014–2019, September 2011.

[6] Dizan Vasquez and Thierry Fraichard. A novel self organizing network to perform fast moving object extraction from video streams. In *IEEE-RSJ Int. Conf. on Intelligent Robots and Systems*, pages 4857–4862, Beijing (CN), October 2006.

[7] Dorinn Comaniciu and Peter Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5):603–619, 2002.

[8] Soh De Loong. People detection using the Microsoft Kinect, December 2011. Retrieved February, 2012, from `http://www.ros.org/wiki/ppl_detection`.

[9] Dizan Vasquez, Thierry Fraichard, and Christian Laugier. Growing hidden markov models: a tool for incremental learning and prediction of motion. *International Journal of Robotics Research*, 28(11–12):1486–1506, 2009.

[10] Hans Moravec. Sensor fusion in certainty grids for mobile robots. *AI magazine*, 9(2):61, 1988.

[11] Chiara Fulgenzi, Anne Spalanzani, and Christian Laugier. Probabilistic motion planning among moving obstacles following typical motion patterns. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009.

[12] Edward T. Hall. *The hidden Dimension: Man's Use of Space in Public and Private*. The Bodley Head Ltd, London, UK, 1966.

[13] Adam Kendon. Spacing and orientation in co-present interaction. In *Development of Multimodal Interfaces: Active Listening and Synchrony*, volume 5967 of *Lecture Notes in Computer Science*, pages 1–15. Springer Berlin / Heidelberg, 2010.

[14] Morgan Quigley, Brian Gerkeyy, Ken Conleyy, Josh Fausty, Tully Footey, Jeremy Leibsz, Eric Bergery, Rob Wheelery, and Andrew Ng. ROS: an open-source robot operating system. In *ICRA Workshop on Open Source Software*, 2009.