

On-board Human-aware Navigation for Indoor Resource-constrained Robots: A Case-study with the Ranger

Zeynab Talebpour

Iñaki Navarro

Alcherio Martinoli

Abstract—Introducing simple robotic platforms into domestic environments is faced with the challenge of social acceptability. Therefore human-aware navigation is a must for robots operating in environments shared with human users. In this work, we focus on the human-aware navigation problem in a structured environment for a robot with limited sensing and constrained maneuvering called Ranger. The Ranger is a simple domestic robotic platform designed for interacting with children. The system combines person detection and tracking—which is the result of fusing laser-scan and depth-image based detectors provided by an RGB-D camera—, basic autonomous navigation and the concept of personal space. We rely only on the on-board sensors for mapping, localization, human tracking, and navigation. Systematic experiments are carried out with a real robot in the presence of a human in order to compare our human-aware navigation with a non human-aware simple approach. The results show that human-aware navigation is able to achieve trajectories which are respecting the personal spaces of the human and are thus more acceptable for the users.

I. INTRODUCTION

Robots will progressively become part of the work spaces and habitats of humans. Despite where they are located, in home environments or hospitals as assistants, in factories as co-workers or as guides in supermarkets, a key behavior which they all share is navigation. The quality of robots' movement influences strongly how their intelligence is perceived [1]. Safety, comfort, legibility and other factors related to social acceptance of a robot all depend on the way that the robot moves. Hence one objective of human-robot interaction is to develop methods for making robot navigation socially acceptable.

Human-aware navigation focuses on the interaction dynamics between humans and robots that occur as a result of navigation [2]. *Human-robot* awareness is defined as the understanding that the humans have of the aspects related to the robots such as the status, activities, locations, identities and surroundings of the robots. On the other hand, *robot-human* awareness is the knowledge that the robot has of the commands of the humans that are needed for directing its activities and also any human-delineated constraints which may require a modified course of action and disobeying commands.

In this article, we focus on the combination of both aspects, implementing a reciprocal human-robot aware navigation on the ranger robot [3], a small limited robotic platform designed to interact with children. To our knowledge, this

is the first work in which both human-robot and robot-human awareness are implemented on-board on a resource-constrained robot without the need of external devices such as external cameras, special trackers, etc. In the literature, we can find several strategies for comfort ranging from appropriate approaching strategy [4], maintaining appropriate distance [5], control strategies to avoid being noisy [6] and use of planning for avoiding interference [7]. Additionally, several approaches to generate legibility for robot navigation, have been reported in [8].

The key factor we address on our human-aware navigation controller is *comfort*. One important concept which is used in numerous studies [9], [5], [10] in this area is that of the virtual space around a person mutually respected by other humans, called *proxemics* [11]. Based on this concept, depending on the relationships and the interactions that exist among humans, people choose different social distances relating to intimate, personal, social or public contexts. Changes in the expected distance may indicate dislike if it is too large or cause discomfort if it is too small. We base our controller on proxemics principles, implementing the model proposed by [12].

Furthermore, we consider legibility an important aspect of human-robot awareness and therefore take it into account in our navigation approach. Legibility means that a person intuitively understands the intentions of a robot [13]. It is shown in [14] that legibility and intent-expressiveness should be the main focus of motion planning as opposed to the predictability of the motion. However, a quantitative evaluation of the robot awareness by the human, is not carried out in this work and should be targeted in the future.

In this article, we also define a set of metrics for human-aware navigation, reflecting comfort, which allow us to compare experimentally our human-aware controller with a non-human-aware one.

The remainder of this article is organized as follows. In Section II, we describe the robotic platform used in this work. Section III provides an explanation of the software architecture and the navigation controllers. Section IV defines the metrics used to measure human-aware navigation based on comfort. In Section V, we show the results of real robot experiments and compare two navigation approaches. Finally, Section VI concludes this work.

II. THE RANGER ROBOT

The robotic platform used in this article is the Ranger robot [3] (see Fig. 1), designed at the Laboratoire de Systèmes Robotiques (LSRO) of EPFL. It is inspired by a

The authors are with the Distributed Intelligent Systems and Algorithms Laboratory, School of Architecture, Civil and Environmental Engineering, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.




Fig. 1: The Ranger robot. a) Complementary gestures added to the robot for increasing the legibility of navigation. b) Ranger augmented with a RGB-D camera to detect people and obstacles. The picture is taken when the robot decided to turn to its right side.

common object found in many children rooms: a wooden storage box for toys. This box is augmented with robotic capabilities in order to interact with children and motivate them to tidy up their room [15].

The Ranger has a body based on a wooden box but is equipped with wheels, mechanical eyes, inertial sensors, three infra-red distance sensors, ground sensors, a bumper, an inside balance, capacitive external touch sensors, LED panels behind the wooden surface, sound, eyeglasses, a detachable pacifier and a relative positioning system to detect other robots. The robot has a square footprint of 30 cm wide and it weights approximately 3.5 kg .

This robot has been designed to support a very specific interaction aimed at encouraging children to tidy up their room. As a result, autonomous navigation with this platform is a challenging problem for two reasons. First, its kinematic configuration consisting of two differential wheels in the front, and two castor wheels in the back, together with its squared shape, make the motion planning complicated. Second, the obstacle detection capabilities are limited to the three infra-red sensors placed in the front of the robot. The detection range of the sensor placed in the middle is $0.2 - 1.50\text{ m}$, while for the other two, it is $0.04 - 0.30\text{ m}$, all of them having very narrow cones. For this reason, we have augmented the sensing capabilities with a depth camera *Primesense Carmine 1.09*, placed in the front upper part of the robot (see Fig. 1b). This depth sensor has a resolution of $640 \times 480\text{ pixels}$, a horizontal field of view of 1 rad and a nominal range of $0.35 - 1.4\text{ m}$.

The Ranger robot has four processors: three micro-controllers, connected using a CAN bus, which manage the real-time part of the robot, and a full embedded computer running Linux Ubuntu 10.04. The ASEBA framework [16] is used to control the whole system, making the link with the higher-level controllers. These high-level controllers are executed on a separate laptop placed on top of the robot, running Linux Ubuntu 14.04.

III. SYSTEM COMPONENTS

In order to generate human-aware autonomous navigation, we have taken an incremental approach of initially enabling

basic autonomous navigation and augmenting it with human-awareness consecutively. Our system is comprised of a number of nodes running within the Robot Operating System (ROS) framework [17] taken from the ROS navigation stack. As a result we are able to compare and evaluate the behaviors of two robotic controllers responsible for the non human-aware and human-aware navigation. Figure 2 shows the different components of our system and how they are connected.

A. Basic Navigation

There are three main components required for enabling the robot to autonomously navigate in an environment: mapping, localization, and path planning. We will explain each part of the system shortly in the following.

1) *Mapping*: A spatial model of the environment surrounding the robot is obtained using its sensors which will then be used for localization and navigation. The map is created by applying a grid-based SLAM using Rao-Blackwellized particle filters [18] implemented by the *gmap-mapping* package.

2) *Localization*: AMCL (adaptive or KLD-sampling Monte Carlo localization) [19] which uses a particle filter based localization method, implemented by the ROS *amcl* package. It tracks the pose of the robot and provides localization using a given map, odometry and a single row of depth readings typically obtained by a laser scan. We have converted the readings of the depth camera to laser-like readings, by taking the closest depth reading from each column in the 10 central lines of the depth image.

3) *Path Planning*: The motor commands required to take the robot to its final goal without colliding with static or moving obstacles are provided by a local and a global planner. We associated a *costmap* to each of these planners. Costmap assign a cost to each cell of a given map. The details of the navigation approach used here, are explained more extensively in [20]. The global planner uses an A*-search algorithm to find the optimal path to the goal and the local planner is based on the Dynamic Window Approach (DWA)[21].

B. Human Tracking

To add human-awareness to a robot capable of autonomous navigation, the robot needs to detect the people present in its surroundings. It is common in the literature of human-aware navigation that detection and tracking of the people is performed or jointly performed by an external tracking system [22].

On-board person detection and tracking becomes an even harder problem for a small-footprint mobile robot with sensors that lie close to the ground due to limited perception. The RGB-D camera of the ranger is positioned such that the robot has a low-lying viewpoint. We perform on-board detection and tracking using two different algorithms that we fuse with means of a Kalman filter to improve the existing person detection and tracking.

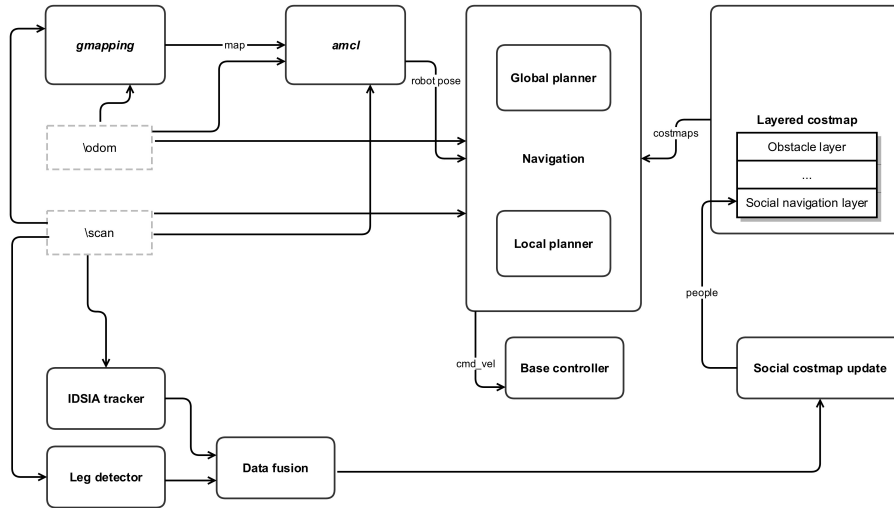


Fig. 2: Connection of different system components.

The first detection and tracking algorithm used is the one by Pesenti Gritti et al. [23] for small-footprint ground robots, which has proven to be robust for cluttered indoor environments. This tracker returns the position of the person relative to the robot and its approximate orientation. To explain their method in short, the point cloud is initially down sampled and the floor is detected as the most prominent horizontal plane. Then the points emerging from the floor are expanded to knee level and potential leg candidates are formed. Thereafter a trained SVM classifier based on Histogram of Oriented Gradients (HOG) features, calculates the probability of each candidate being a leg. Additionally, tracking is done using a Kalman filter.

The second tracker, by Arras et al. [24], is a simple leg detector which uses two dimensional range scans. In that work, supervised learning is used for creating a classifier that simplifies the detection of people. *AdaBoost* is applied for training a strong classifier from simple features of groups of nearby beams corresponding to legs in the range data. Figure 3a shows a sample visualization of the laser scan for the leg detector and Fig. 3b is depicting how the final detection of the person is visualized upon leg matching and fusion.

Pesenti Gritti's tracker detects and tracks people robustly with a small number of false-positive detections. However, in our experiments the speed of detection for this tracker was smaller than Arras' tracker. Arras' tracker has a very high speed of detection but false positive detections can be reported due to limited information used for detecting legs. We have fused both trackers using a Kalman filter that is initialized with the first detection of Pesenti Gritti's tracker and is updated with data coming from both trackers which have different information publishing rates.

C. Human-Awareness

Once the robot knows where the people are in the environment it should consider them differently from other

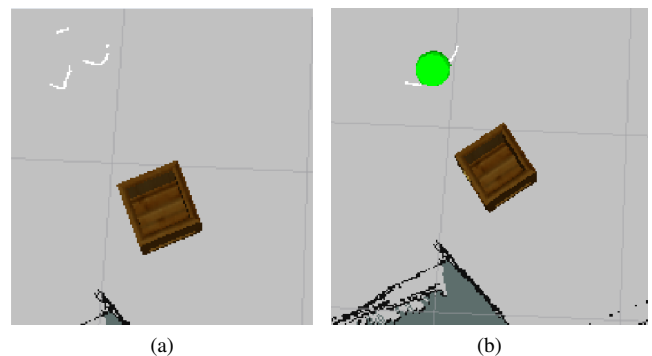


Fig. 3: (a) Robot's Laser scan when encountering a person. (b) Final detection and tracking of the person is the result of fusing the two trackers.

obstacles that it perceives and should plan accordingly. We rely on proxemics to make this difference and assign a cost to the personal space of the human which the robot should avoid intruding. A two-dimensional Gaussian cost function is centered around each person with its variance proportional to the relative velocity of the person, similar to the approach of [12]. These social costs are considered in an additional ROS costmap layer similar to [25] which will influence path planning. Figures 4a and 4b show how the social costmap looks like around a person in different situations. This social navigation layer is the final part required for our human-aware navigation system.

There exist a number of difficulties that we have faced, trying to achieve this final navigation behavior. The limited range and field of view of the depth camera cause limited perception of the environment which affects different system components, particularly localization. This results in a degradation of the overall navigation performance.

Additionally, it is not generally very clear how the social costmaps should be preserved over time. Considering the high dynamicity of the human movement, a common as-




Fig. 4: (a) The path of the robot given by the planner when encountering a social costmap centered around a static human. (b) The social costmap of a person moving towards the robot.

sumption is that the social cost assigned to the area surrounding a person is only valid as long as the human is detected and is within the perception range of the robot. However, this creates problems for robots with limited perception such as the Ranger, when the people present in the environment are static or do not change their positions considerably. We encountered this problem during our experiments when the robot re-planned upon detecting a person moving in a socially acceptable manner, but once the person was out of reach of its sensing range, it re-planned again trying to minimize the distance needed to travel to reach the goal in a human-agnostic manner.

We have chosen to update the social costmap for a predefined period which should be tuned according to the level of dynamicity of the human movements in the environment. In our experiments, we update the social costmap for a period of 30 s upon receiving no data from the people trackers. After this period, a timeout clears the social costmap and erases all the past data. It is arguable which approach is better in the general case, but as mentioned before, we believe this decision of how long to preserve a social costmap is tied to the type of the environment and the perception capabilities of the robot. If the robot is able to observe the people at all times it makes sense not to update the social costmaps but if this is not the case, erasing it immediately as soon as the human is out of sensing range, would be questionable.

D. Robot-Human Awareness

To increase the awareness that the human has of the robot and for having a more legible navigation we have added complementary gestures to the robot as advised by [8]. This is done for clarifying intention which in our case is the direction towards which the robot wants to move. The pupils of the robots eyes indicate the direction of the movement and the LEDs on the side to which the robot wants to turn are activated to improve the navigation legibility. Figure 1b shows the robot before wanting to turn to its right.

E. Reciprocal Human-Robot Awareness

In basic navigation, only the global and local costmaps used for assigning costs to areas corresponding to obstacles

are taken into account for path planning. However, when augmented with human detection and tracking modules, the robot is able to have a social consideration for people when performing this task. This is done by assigning costs to areas related to people in an additional costmap which is combined with global and local costmaps for taking the final decision about the path that the robot has to take. This social costmap is responsible for making the navigation human-aware. Another aspect of importance here is the ability of the robot to make its intentions known which calls for increasing legibility. This results in the awareness of the robots actions by the human.

IV. METRICS

In order to evaluate the performance of our controllers, we introduce three metrics.

A. Minimum Distance to the Human

The minimum distance between the robot and the human during the experiment is selected.

$$m_1 = \min_{\forall t \in T} \text{dist}(\vec{x}_r[t], \vec{x}_h[t]) \quad (1)$$

where dist is a function which returns the Euclidean distance between two points in meters. For each measurement $\vec{x}_r[t]$ obtained at time t during an experiment of length T , \vec{x}_r denotes the position of the robot and \vec{x}_h denotes the position of the human.

B. Time Spent in Areas Associated to the Social Costmap

The average time spent in areas with social costs assigned to them in seconds. We define the area around a human which should be avoided as a circle of radius r defined as follows.

$$r = (-2\sigma^2 \cdot \log(\frac{\text{cutoff}}{A}))^{\frac{1}{2}} \quad (2)$$

This is because in our experiments the person is always static. A is the amplitude and σ^2 is the variance used for the two-dimensional Gaussian cost function and cutoff is a threshold which is used for limiting the area of the costmap. This radius is also applied in the implementation to specify the area of the social costmap.

$$m_2 = \frac{\sum_{t=1}^T f[t]}{T} \quad (3)$$

$$f[t] = \begin{cases} 1 & \text{if } \text{dist}(\vec{x}_r[t], \vec{x}_h[t]) \leq r \\ 0 & \text{if } \text{dist}(\vec{x}_r[t], \vec{x}_h[t]) > r \end{cases} \quad (4)$$

C. Accumulated Social Cost

This metric assigns the social cost used in navigation by the planners to the points traversed by the robot which intrude the areas around the human.

$$m_3 = \frac{\sum_{t=1}^T C[t] \cdot f[t]}{T} \quad (5)$$

$$\begin{cases} m_x = \text{dist}(\vec{x}_r[t], \vec{x}_h[t]) \cos(\theta) \\ m_y = \text{dist}(\vec{x}_r[t], \vec{x}_h[t]) \sin(\theta) \end{cases} \quad (6)$$




Fig. 5: On the left: The view given by the overhead camera of the ground truth system of the arena where our experiments were conducted. The starting position of the robot, position of the human and the goal are marked on the image. On the right: The map of the environment used for navigation.

$$C[t] = A.e^{-\left(\frac{m_x^2}{2\sigma_x^2} + \frac{m_y^2}{2\sigma_y^2}\right)} \quad (7)$$

where θ is the bearing of the robot relative to the person, σ_x^2 is the variance of the Gaussian in x dimension, and σ_y^2 is the variance of the Gaussian in y dimension. The values of σ_x , σ_y and A are taken from commonly used values in the literature for a static person. However, the model allows for considering velocity in x and y directions for calculating the sigmas. These values can be tuned for the desired behavior of the robot, e.g., a more conservative approach uses larger costmaps whereas smaller costmaps can be used for more populated environments or for cultures that demand a smaller personal space around individuals.

V. EXPERIMENTAL RESULTS

To test our human-aware navigation in a real scenario we did a number of experiments in the environment depicted in Figure 5. The approximate area of our arena is $6 \times 4 m^2$. We used an overhead tracking system as the provider of the ground truth to keep record of the pose of the robot during our experiments. The positioning error is $2\text{-}3\text{cm}$ in the central area of the arena and $5\text{-}7\text{cm}$ further away from the camera on the borders the environment. We investigated the performance of our human-aware navigation method in a simple move-to-goal scenario in the presence of a static person. We are interested in systematically comparing the performances of both navigation algorithms in the mentioned scenario. Each experiment has been repeated 10 times for a given algorithm. In all experiments, the human is standing in the arena and the trajectory of the robot is recorded by tracking an active marker placed on the robot. The human is placed directly under the camera to minimize occlusions. The main reason for having the human static is the difficulties arising from tracking both the robot and the person with the overhead camera given the small height of the robot. Since we have this simple case of static human, we did not evaluate the robot-awareness gained by adding features that increase legibility. Although it is perceived that signaling intentions of the robot results in better acceptability even for a static human, we need to further study this aspect in the future.




Fig. 6: Sample Ranger trajectories captured by the ground truth system for both navigation approaches. Positions are given in the camera coordinate system.

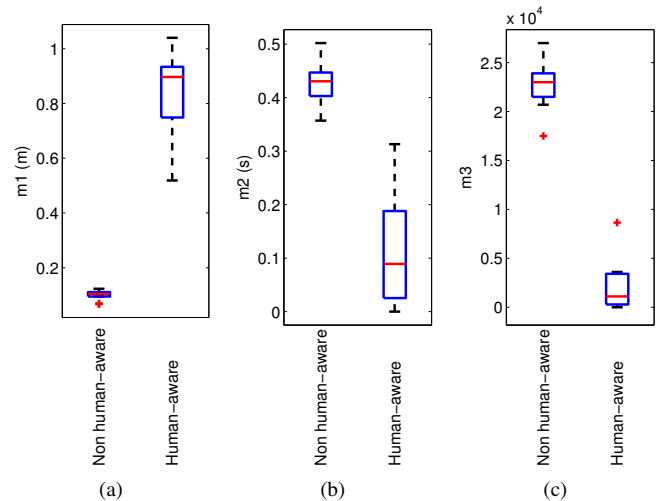


Fig. 7: Boxplots for the three metrics acquired for non human-aware and human-aware navigation approaches. (a) Minimum distance to the human in meters (m_1). (b) Time spent in areas associated to the social costmap in seconds (m_2). (c) Accumulated social cost (m_3).

Figure 6 shows two sample trajectories from our experiments. The human-aware approach (blue) is clearly avoiding to enter into the areas close to the human whereas the basic navigation approach (green) favors the shortest path between the starting point to the target position while considering people as obstacles which should be avoided to ensure safety.

Figure 7 shows the result of our metric evaluation. As expected we can see that there is a large difference between the minimum distance of the robot to the human (m_1) in the two approaches and the human-aware approach keeps a much bigger distance. At the same time the associated costs in terms of the average time spent close to the human (m_2) and the accumulated social cost (m_3) are also considerably higher for the non human-aware approach showing the effectiveness of our human-aware navigation.

Results show that the human-aware approach is more sensitive to variations in different experiments. We suspect that this might be due to the higher sensitivity of the human-aware approach to the position of the person that

corresponds to the center of the social costmap. This costmap is constantly updated and used for replanning, while the other approach considers the human merely as an obstacle and accounts for it in collision avoidance.

VI. CONCLUSIONS AND FUTURE WORK

We have demonstrated the effectiveness of autonomous human-aware navigation for a simple domestic robotic platform designed for children, relying only on its on-board sensors and computation capabilities. We have also defined a set of metrics for measuring human-awareness. Different system components, arising challenges and limitations have been explained. A simple real-world experiments show how our system is able to perform the navigation task in a socially acceptable manner.

For fully understanding the effectiveness of our reciprocal human-robot aware navigation algorithm we need to perform experiments to assess the robot-awareness gained by the human. In addition, improving human-awareness by detecting ongoing interactions present in the environment, predicting human motions and future interactions, and human motion model development are further steps that lead to a higher social acceptance. However, although Ranger is a good robot for studying human-robot interactions targeted to children, the limitations of this simple platform result in difficulties for more complex navigation algorithms. Therefore, our future research aims to study the mentioned aspects by means of mBot, a robotic platform developed within the FP7 European project MONarCH [26].

ACKNOWLEDGMENT

This research was partially supported by the Swiss National Science Foundation through the National Center of Competence in Research Robotics. We would like to express special thanks to Francesco Mondada, Daniel Burnier, Norbert Crot, Florian Vaussard and Philippe Rétornaz from the Laboratoire de Systèmes Robotiques at EPFL, to Séverin Lemaignan from the Computer-Human Interaction in Learning and Instruction lab at EPFL, and to David Mansolino from our lab for their great help with the Ranger robot. Without their support this work would have not been possible. We would also like to thank Gianni di Caro, Armando Pesenti Gritti and Alessandro Giusti for their valuable support on using and integrating their people tracker software.

REFERENCES

- [1] P. Althaus, H. Ishiguro, T. Kanda, T. Miyashita, and H. Christensen, "Navigation for human-robot interaction tasks," in *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 2, 2004, pp. 1894–1900.
- [2] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch, "Human-aware robot navigation: A survey," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1726 – 1743, 2013.
- [3] F. Mondada, J. Fink, S. Lemaignan, D. Mansolino, F. Wille, and K. Franinović, "Ranger, an Example of Integration of Robotics into the Home Ecosystem," in *International Workshop and Summer School on Medical and Service Robotics*, 2014.
- [4] K. Dautenhahn, M. Walters, S. Woods, K. L. Koay, and C. L. Nehaniv, "How may i serve you?: a robot companion approaching a seated person in a helping context," in *1st ACM SIGCHI/SIGART conference on Humanrobot interaction*. ACM Press, 2006, pp. 172–179.

- [5] L. Takayama and C. Pantofaru, "Influences on proxemic behaviors in human-robot interaction," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 5495–5502.
- [6] E. Martinson, "Hiding the acoustic signature of a mobile robot," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007, pp. 985–990.
- [7] D. Vasquez, P. Stein, J. Rios-Martinez, A. Escobedo, A. Spalanzani, and C. Laugier, "Human Aware Navigation for Assistive Robotics," in *13th International Symposium on Experimental Robotics*, 2012.
- [8] C. Lichtenthäler and A. Kirsch, "Towards legible robot navigationhow to increase the intend expressiveness of robot navigation behavior," in *International conference on social roboticsworkshop embodied communication of goals and intentions*, 2013.
- [9] J. Mumm and B. Mutlu, "Human-robot proxemics: Physical and psychological distancing in human-robot interaction," in *6th ACM/IEEE International Conference on Human-Robot Interaction*, 2011, pp. 331–338.
- [10] M. Walters, M. Oskoei, D. Syrdal, and K. Dautenhahn, "A long-term human-robot proxemic study," in *IEEE RO-MAN*, 2011, pp. 137–142.
- [11] E. T. Hall, "The hidden dimension," 1966.
- [12] R. Kirby, R. Simmons, and J. Forlizzi, "Companion: A constraint-optimizing method for person-acceptable navigation," in *The 18th IEEE International Symposium on Robot and Human Interactive Communication*, 2009., Sept 2009, pp. 607–612.
- [13] T. Kruse, P. Basili, S. Glasauer, and A. Kirsch, "Legible robot navigation in the proximity of moving humans," in *2012 IEEE Workshop on Advanced Robotics and its Social Impacts.*, 2012, pp. 83–88.
- [14] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," in *2013 8th ACM/IEEE International Conference on Human-Robot Interaction.*, 2013, pp. 301–308.
- [15] J. Fink, S. Lemaignan, P. Dillenbourg, P. Rétornaz, F. Vaussard, A. Berthoud, F. Mondada, F. Wille, and K. Franinović, "Which robot behavior can motivate children to tidy up their toys?: Design and evaluation of ranger," in *Proceedings of the ACM/IEEE International Conference on Human-robot Interaction*, 2014, pp. 439–446.
- [16] S. Magnenat, P. Rétornaz, M. Bonani, V. Longchamp, and F. Mondada, "ASEBA: A Modular Architecture for Event-Based Control of Complex Robots," *IEEE/ASME Transactions on Mechatronics*, vol. 16, no. 2, pp. 321–329, 2011.
- [17] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "Ros: an open-source robot operating system," in *ICRA Workshop on Open Source Software*, 2009.
- [18] G. Grisetti, C. Stachniss, and W. Burgard, "Improved techniques for grid mapping with rao-blackwellized particle filters," *IEEE Transactions on Robotics.*, vol. 23, no. 1, pp. 34–46, 2007.
- [19] D. Fox, "Adapting the sample size in particle filters through kld-sampling," *The international Journal of robotics research*, vol. 22, no. 12, pp. 985–1003, 2003.
- [20] E. Marder-Eppstein, E. Berger, T. Foote, B. Gerkey, and K. Konolige, "The office marathon: Robust navigation in an indoor office environment," in *International Conference on Robotics and Automation*, 2010.
- [21] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.
- [22] D. Vasquez, P. Stein, J. Rios-Martinez, A. Escobedo, A. Spalanzani, and C. Laugier, "Human aware navigation for assistive robotics," in *Experimental Robotics*, 2013, pp. 449–462.
- [23] A. P. Gritti, O. Tarabini, J. Guzzi, G. Di Caro, V. Caglioti, L. M. Gambardella, A. Giusti, *et al.*, "Kinect-based people detection and tracking from small-footprint ground robots," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems.*, 2014, pp. 4096–4103.
- [24] K. O. Arras, O. M. Mozos, and W. Burgard, "Using boosted features for the detection of people in 2d range data," in *2007 IEEE International Conference on Robotics and Automation.*, 2007, pp. 3402–3407.
- [25] D. V. Lu, D. Hershberger, and W. D. Smart, "Layered costmaps for context-sensitive navigation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 709–715.
- [26] J. Messias, R. Ventura, P. Lima, J. Sequeira, P. Alvito, C. Marques, and P. Carriço, "A robotic platform for edutainment activities in a pediatric hospital," in *IEEE International Conference on Autonomous Robot Systems and Competitions.*, 2014, pp. 193–198.