A person follower mobile robot system for indoor environments

Ninad Pradhan, Stan Birchfield, and Timothy Burg
Dept. of Electrical and Computer Engineering, Clemson University

***Abstract*— Person following is integral to robot companion systems and in service and assistive robots operating in various other scenarios. Human populated environments challenge these robots to navigate around highly dynamic workplaces. The primary challenges for these robotic systems, navigation through uncertainty and occlusion detection and handling, are addressed in this paper. Appearance descriptors built around pose information from person-tracking framework assist in occlusion inference. Predictive fields path planning models obstacle motion and is useful in navigating dynamic environments. Person followers need to map and localize their environment to be able to accomplish their primary tasks. A mapping method based on robot odometry and Manhattan constraints achieves drift-free mapping in feature-poor, homogeneous hallways typical to indoor environments. A taxonomy is proposed to be able to classify person followers relative to expectations from the human leader interacting with the robot.**

I. INTRODUCTION

Robots are beginning to make a gradual expansion from industrial environments into everyday life, participating in household or personal activities. Such robots are called 'service robots' [1], and are charged with operating in human populated environments. Various applications of mobile service robots have been investigated over the past few years, from 'BIRON' [2] the robot companion, to 'Johnny' [3] the robot butler.

Person following is either a module in a multi-functional service robot or an independent service robot class. Service robots with person following have typically been used to provide service and care in indoor environments, and have been called 'socially assistive robots' [4]. Such robots have helped the elderly suffering from cognitive impairment [5], [6] and carried around oxygen therapy tanks for their human companions [7].

The importance of person following in assistive robotics has driven the development of many person follower systems. These systems have highlighted and addressed many of the individual challenges in person following. Leader detection and tracking has been demonstrated using stereo and feature tracking [8], [9], mean-shift color histograms [10], and machine learning [11]. Occlusion detection and inference has been implemented using motion models with Extended Kalman Filter (EKF) [12], or purely based on appearance information [13]. Navigation in dynamic environments has been implemented by learning human motion patterns [14], avoiding typically crowded areas in familiar environments [15], and tracking human motion to avoid their projected real-time positions [16].

The conceptual development and practical implementation of a person follower is presented in this paper. The system-level contributions of this work are threefold: an occlusion inference system based on appearance and pose information is developed (Section III-A); a Simultaneous Localization and Mapping (SLAM) approach robust to rotational drift and homogeneity of typical indoor environments is demonstrated (Section III-B); and our work in predictive fields path planning [17], [18], [19] is applied to the challenging problem of robot navigation in dynamic, populated environments (Section III-C). The Kinect RGBD is used for sensing the robot environment. Results which highlight each contribution are given in Section IV.

The importance of particular contributions to person following literature has thus far been hard to gauge. In previous work on person followers, expectations from the robot have been well defined [20] and system development has largely been guided by robot-centric goals. However, no explicit effort has been made so far to define the person follower problem in terms of expectations from the human leader. An human-centric taxonomy of person follower systems is proposed in Section II, with the intention of providing context for person follower systems to both end-users and researchers.

II. PERSON FOLLOWER CLASSIFICATION SYSTEM

Yoshimi et al. [20] clearly outlined the functions of person following robots as follows:

- The robot should initialize to its leader,
- It should follow the leader at his/her pace,
- It should avoid obstacles, and
- Contact with the leader should be resumed after occlusions

Person followers have implicitly or explicitly been designed with these considerations in mind. While these tasks capture robot-specific requirements quite well, they do not address an equally important question - what is the expected role of the leader in the human-robot interaction for a specific person following robot?

The expected role of the human leader is an useful perspective for users of person following robots. Moreover, it provides an human-centric basis to compare and contrast person following systems. A taxonomy for person followers, which could lead to a standardized metric for assessing such systems, is proposed in this section.

The proposed classification system assigns four 'levels' of sophistication for the person follower. The levels are assigned titles appropriate to leader expectations: 'Fully cooperative leader' (Level 1), 'Partially cooperative leader' (Level 2), 'Patient leader' (Level 3), and 'Independent leader' (Level

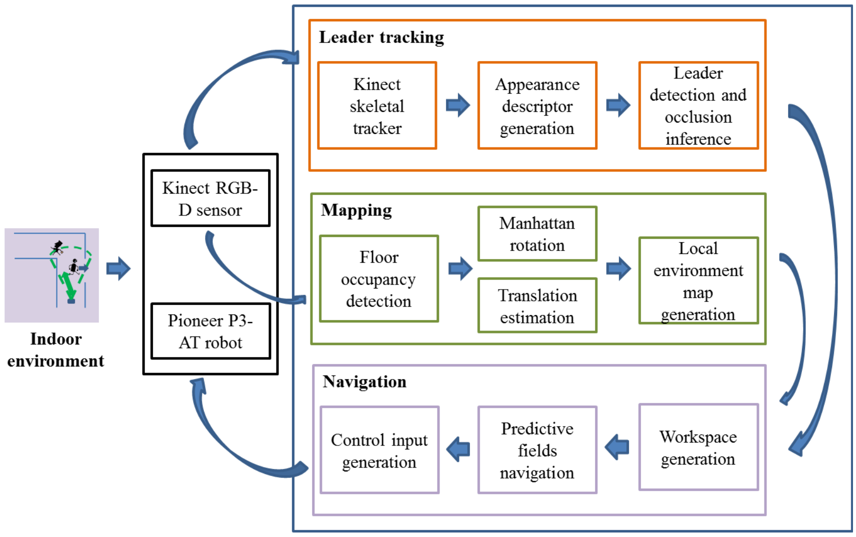

Fig. 1. Block diagram of the person follower system.

4). Every level contains at least the functionality of all levels preceding it, and additional capabilities which separate it from preceding levels.

A. Level 1 - Fully cooperative leader

In Level 1, the leader:

- Is aware of the physical limitations of the robot, e.g. maximum speed, sensor range, physical dimensions.
- Creates an ‘exact’ path for the robot to follow.
- Stays within the maximum range of the robot sensor.
- Ensures that neither partial nor complete occlusion occurs due to obstacles in the environment.

Thus, a fully cooperative leader guides the robot to a position close to him/her at all times, and maintains a direct line of sight with respect to the robot.

The Level 1 paradigm has been used in the past to test algorithms such as stereo based leader detection [9], and stereo-based outdoor person following [8].

B. Level 2 - Partially cooperative leader

In Level 2, the leader:

- Is aware of the physical limitations of the robot.
- Expects robot path planning around environmental or stationary obstacles.
- Ensures complete occlusion does not occur; however, allows partial occlusions such as low-lying stationary obstacles.

Level 2 differs from Level 1 in terms of the expectations from the robot to plan its own path around stationary obstacles. Hence, the leader is said to be ‘partially cooperative’. Some form of map generation may be required for path planning.

The Level 2 paradigm has been used in demonstrating stationary obstacle avoidance within the person following problem definition [20], [21].

C. Level 3 - Patient leader

In Level 3, the leader:

- Expects the robot to create its own path around stationary or moving obstacles.
- Makes no effort to stay within robot sensor range.
- Makes no effort to avoid partial or complete occlusions.
- Waits for the robot to overcome occlusions before resuming his/her motion.

Level 3 differs from Level 2 in that an occlusion detection and handling module is mandated by system requirements. A sophisticated path planning mechanism with mapping and localization is required for path planning in dynamic environments. The ‘patient leader’ requirement is that the leader allow the robot to execute its planned route fully before resuming motion.

The proposed system [17] is fully compliant with Level 3 requirements. Other Level 3 approaches have used motion-based tracking to recover leader after occlusions[11].

D. Level 4 - Independent leader

In Level 4, the leader:

- Expects the robot to create its own path around stationary or moving obstacles.
- Does not wait for occlusion handling to complete before resuming motion.
- Expects high level communication when the robot is either distressed (e.g. mechanical breakdown) or lost.

Level 4 differs from Level 3 in the relaxation of the requirement for the leader to wait for the robot. The implies that the robot should be capable of reestablishing leader location relative to a global map, and should be able to communicate with the leader.

Level 4 systems, which place very limited expectations on the leader, are thus the least intrusive to humans in a human populated environment. A system which realizes all of the requirements for Level 4 person followers has not been realized, though recent work on socially assistive robots comes close [6].

As service robots become common, the proposed taxonomy could be standardized to act as a guide for customers looking for the right robot for household or office use.

III. Person follower system

Designed in accordance with the requirements of a 'Level 3 - Patient Leader' person follower, the system block diagram is shown in Figure 1. The system comprises of three modules:

- Leader tracking
- Mapping and localization
- Navigation

'Leader tracking' leverages skeletal information provided by the Kinect skeletal tracker [22] to solve the challenging problem of keeping track of the leader through short term and long term occlusions. Skeletal tracking output is augmented with appearance descriptors to develop an appearance and pose based occlusion inference algorithm.

'Mapping and localization' for an indoor-navigating robot has to work in typical hallway environments. Typical simultaneous localization and mapping (SLAM) approaches depend on availability of workspaces with an abundance of color or depth features. The technique developed for our mapping module requires neither; it uses robot odometry and Manhattan constraints to give a reliable, drift-free mapping output.

The output of leader tracking and mapping modules feeds the 'navigation' module. Robot navigation algorithms need to have guarantees of convergence to goal, and need to incorporate obstacle motion to work in human-populated, dynamic environments. Our recent work on 'predictive fields path planning' [18], [19] extends the classic navigation function technique [23] to practical environments, and is a natural fit for the navigation requirements of a robotic person follower.

A. Leader tracking

The leader tracking module is designed to be capable of tracking the leader through short-term and long-term occlusions. Without this capability, a person follower would be over-dependent on its leader for occlusion avoidance. Previous work on occlusion avoidance has focused on depth [12] or appearance [13] as separate cues for occlusion inference. The system presented here combines both into a pose and appearance based occlusion inference algorithm.

Pose information is extracted from a depth image using the Microsoft Kinect SDK skeletal tracker [22]. This part-based algorithm represents the human figure as an ensemble of 20 joints (19 bones), and can track up to 6 skeletons per depth image, providing access to relevant information such as joint 3D positions and tracking status per skeleton. It is capable of robustly tracking unoccluded skeletons. However, the tracker does not have the capability to maintain skeletal tracks through occlusions, assigning a new identity to a skeleton reappearing after occlusion.

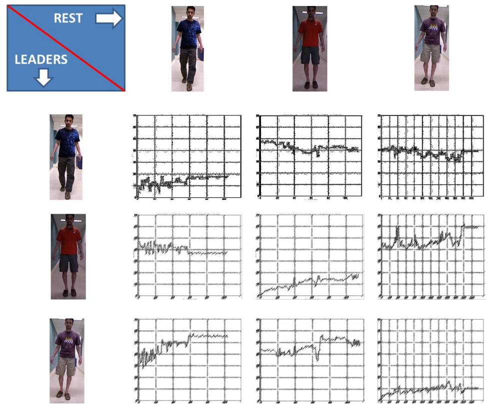

Fig. 2. Tabulated results from descriptor comparisons

In each cell, the row image is the leader, against which all test descriptors of the trial corresponding to the column image are compared. The x axis is frame number and y axis is comparison score. Leader descriptors from frames in the same trial show low comparison scores along the diagonal. Off-diagonal graphs show higher comparison scores, as desired.

This gap in the tracking framweork is addressed by augmenting pose information for the detected skeletons with a novel color-based appearance descriptor. To generate this descriptor, the skeletal outline from the depth image provided by the tracker is superimposed on the RGB image. 'Bone patches', sections of the RGB image along bone orientation, are extracted and converted to two different colorspaces: Hue-Saturation-Intensity (HSI) and L*a*b* (CIE1976). These colorspaces provide separate access to chrominance components of color. Concatenating their mean chromaticity components (H,S and a*,b*) provides an illumination invariant color representation for each bone patch. This pose-based skeleton appearance descriptor comprises of 76 elements:

$$D_n = [\ H_1 \quad S_1 \quad a_1^* \quad b_1^* \quad \dots \quad H_{19} \quad S_{19} \quad a_{19}^* \quad b_{19}^*\ ]$$

A leader appearance descriptor is generated during system initialization. Subsequent descriptors, e.g. descriptor i in frame j (D_i^j), are compared to the leader descriptor, D_{leader}, to infer whether the leader has been occluded. The Euclidean distance between these 76d descriptors gives a match score $S_{i,leader}^j$:

$$S_{i,leader}^j = \left\| D_{leader}^0 - D_i^j \right\| . \tag{1}$$

A threshold is set for the match score, and scores exceeding this threshold are indicative of an occlusion having occurred. Figure 2 shows that generated descriptors are sensitive to appearance variations. To ensure that the skeletal tracker and appearance descriptor don't work at cross purposes, the following rules are observed:

- When the skeletal tracker infers that the leader is tracked, high confidence is placed in its tracking ability. This is done by setting a less strict match score threshold.

- When the leader skeleton has been lost, reidentification is confirmed by setting a strict match score threshold for the appearance descriptor. When this is satisfied, the skeleton is labeled as the leader.

The advantage of appearance based descriptors over temporal occlusion handling methods [11] is that they place no constraints on the time between occlusion and reappearance. Thus, both short and long term occlusions can be inferred.

B. Mapping and localization

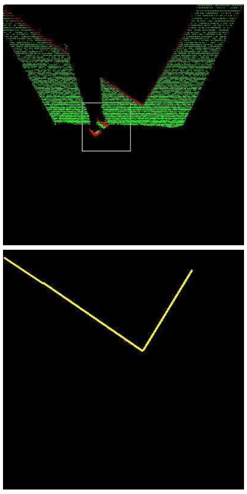

Fig. 3. Multi-line RANSAC estimates from scan data

The top image shows raw occupancy image data. This is converted to a scan representation and multi-line RANSAC operates on scan data to estimate wall lines, shown in the bottom image using yellow lines.

To overcome occlusions, a person follower needs to build a local map and localize itself so as to plan a path within its environment. This requires the use of SLAM techniques, which commonly use 3D points corresponding to image features [24], 3D point clouds (Iterative Closest Point [25]), or 2D range scans [26] to estimate inter-frame robot transformations. These techniques work best in feature rich environments, where visual or depth cues can be leveraged to accurately build a local map. Their performance degrades in homogeneous hallways typical to indoor environments.

The mapping approach for the system presented here overcomes this vulnerability of standard techniques by combining robot odometry and Manhattan rotation estimates. Manhattan rotation estimates [27] are based on the assumption that planes in an indoor environment are mutually orthogonal. With this assumption, which holds true for many common indoor environments, it is possible to get rotation estimates free from rotational drift which inevitably accumulates when feature or scan based inter-frame estimates are used.

The raw data for mapping is an 'occupancy image', which uses a least square estimated floor plane equation to identify non-floor points. This is converted to a polar (r, θ) representation, where, for each angle θ_i, r_i is the radial distance to the shortest occupied point in the occupancy map. These 'scan points' are used in multi-line RANSAC estimation [17], shown in Figure 3.

During map initialization, a reference line orientation is detected and mutually orthogonal orientation reference bins are created. Difference in orientation, or 'offset', of lines in each frame relative to reference orientation bins is computed. Inter-frame rotation can then be computed using the offset values, leading to total rotation estimate relative to map origin:

$$\theta_{n_{total}} = \theta_{n-1_{total}} + \theta_{n-1}^{n} \tag{2}$$

To get an estimate of the translation, the robot motion pattern is controlled such that it rotates first and then translates at a fixed velocity $\dot{r}$ along its new orientation vector. If Δt_{n-1}^{n} is the time elapsed since the last translation estimate, then the robot has translated:

$$r_{n-1}^{n} = \dot{r}_{n-1}^{n} \cdot \Delta t_{n-1}^{n} \tag{3}$$

along its direction vector.

If the inter-frame translation of the robot is given by $(\Delta x_{n-1}^{n}, \Delta z_{n-1}^{n})$, then the current position of the robot relative to the map origin in the (x, z) plane is given by:

$$\begin{aligned} x_{n_{total}} &= x_{n-1_{total}} + \Delta x_{n-1}^{n} \\ z_{n_{total}} &= z_{n-1_{total}} + \Delta z_{n-1}^{n} \end{aligned} \tag{4}$$

To generate a map using this data, the location of each point in the occupancy map is transformed using Eq. 2 and Eq. 4 and the transformed location indicated on the global map.

C. Navigation

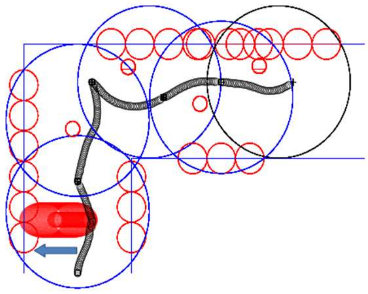

Fig. 4. Predictive field navigation over a L-shaped hallway

Overlapping black circles show robot trajectory across workspaces represented by large blue circles. The robot avoids a moving obstacle whose direction is indicated using a blue arrow and path using overlapping red envelopes.

The robot follows its leader using a combination of 'exact person following' and path planning. As long as the person is visible, and the straight line trajectory to the person is clear, the robot is commanded to move to its desired position and orientation using a control input similar to a simple proportional controller:

$$\begin{aligned} \dot{r} &= \min(r_{observed} - r_{desired}, \dot{r}_{max})(mm/sec) \\ \dot{\theta} &= \theta_{observed} - \theta_{desired}(deg/sec) \end{aligned}$$

where $(r_{observed}, \theta_{observed})$ are robot-relative polar coordinates of the leader, $r_{desired}$ is desired distance to leader. The robot rotates to place the leader at the center of the Kinect's field of view.

Leader occlusion or obstacle detection activates the path planning algorithm. This algorithm, called 'predictive fields path planning' was developed in our previous work [17], [18], [19] as an extension of navigation function path planning [23] to dynamic environments. It preserves navigation function's mathematical guarantees of convergence to goal and adds the following features:

- Encapsulation of obstacle motion using predictive fields.
- Normalized direction inputs for robot velocity.
- A simple, practical workspace representation
- Moving goal representation in the form of waypoints.

The robot is given a directional control input $\dot{q} = u$ [18], where:

$$u = -K \frac{\left(\frac{\partial \varphi}{\partial q}\right)^T}{\left\| \frac{\partial \varphi}{\partial q} \right\| + \epsilon}, \tag{5}$$

K being the matrix of control gains and ϵ a small positive constant. The term φ is called the navigation function [23]. It incorporates terms attractive to goal position and repulsive terms for moving and stationary obstacles which encapsulate obstacle motion patterns using elliptical predictive fields [19]. A simulated indoor navigation route is shown in Figure 4.

IV. Results

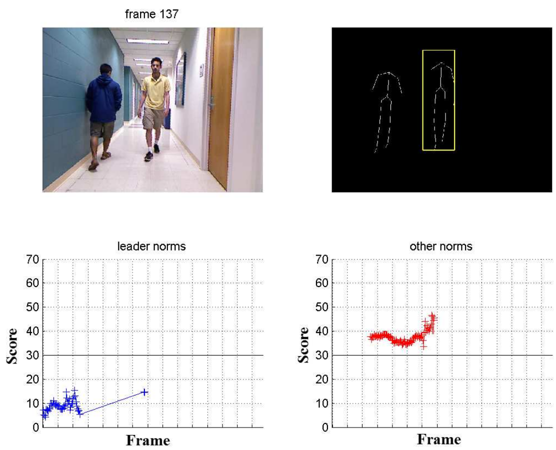

Fig. 5. Leader recovery after occusion.

Plot on the bottom left shows leader match scores and that on the bottom right shows non-leader skeleton scores. The leader is identified after occlusion as shown by the reappearance of a blue mark and connecting line on the leader plot on the bottom left. The non-leader skeleton consistently returned high comparison scores, as evidenced by the bottom-right plot.

The system was implemented using a Pioneer P3-AT mobile robot, interfaced to the controlling laptop using Aria libraries. A tripod mounted Kinect was used as the forward-looking RGBD sensor. The Kinect was interfaced using the Microsoft Kinect SDK, whereas computer vision components of the algorithm were implemented using the open source Blepo vision library.

Leader tracking: Occlusion inference and leader detection after reappearance was tested using a case where the leader walked in front of the robot, was occluded, and the occluding person then moved out of the way. Figure 5 shows the state of the occlusion inference module just after the occlusion had passed. The bottom left and right plots have frame number (or time) as their x axis. Inferred leader detections are plotted on the bottom left and occluding person detections on the bottom right plot. As can be seen, the plots are consistent with the setup of the trial.

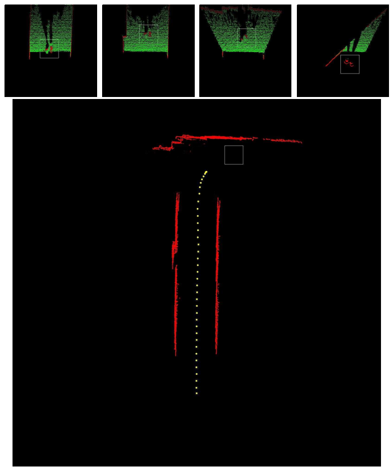

Fig. 6. Map output at L-junction.

Occupancy images are stitched together using the mapping algorithm to yield the map image. Red pixels in the map image are transformed scan points from all frames, yellow dots show the motion of the robot as it turns around the L-shaped hallway. The white rectangle is the position of the leader in the current frame.

Mapping and localization: The robot's environment, and the path traced by the robot in it, can be mapped relative to the local map generated in Figure 6. Some of the occupancy maps which are stitched together to get the local map are shown in the top panel of this figure. Green areas in the occupancy map are floor pixels. Manhattan estimates allow the mapping module to cope with turns around corners with very little drift.

Navigation: The leader moved unoccluded for some time, during which the robot operated in exact person following mode. Sensor range or obstacle occlusions triggered a switch to predictive path planning, as shown in Figure 7. After overcoming the occlusion, the robot switched back to its default 'exact following' state.

V. Conclusions and Future Work

A person follower system was presented, motivated by the desire to develop a practical and useful service robot. It was shown that the person following classification system provides a comparative basis for assessing person followers.

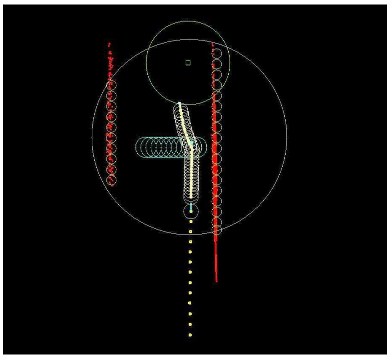

Fig. 7. Path planning around a moving occluding person.

Wall obstacles are represented using overlapping yellow circles, and the single moving person obstacle, moving right to left, is represented by overlapping blue circles. Overlapping white circles show the trajectory of the robot. To overcome occlusions, the robot uses predictive fields path planning. When the robot reaches the large circular region around goal, it resumes exact following.

Future iterations of this system will benchmark robot performance within each class. The appearance-based occlusion handling technique could be augmented using other modalities such as face or gait recognition, to make it more robust to occlusions. In indoor environments, landmark or event detection, such as the opening of a door or the entrance to a stairwell, could provide useful heuristic information for the navigation module and greatly improve the utility of these service robots.

REFERENCES

[1] R. Schraft and G. Schmierer, *Service robots: products, scenarios, visions*. A K Peters, 2000.

[2] A. Haasch, S. Hohenner, S. Huwel, M. Kleinehagenbrock, S. Lang, I. Toptsis, G. A. Fink, J. Fritsch, B. Wrede, and G. Sagerer, "Biron - the bielefeld robot companion," in *in Proc. Int. Workshop on Advances in Service Robotics*, 2004, pp. 27–32.

[3] T. Breuer, G. Giorgana Macedo, R. Hartanto, N. Hochgeschwender, D. Holz, F. Hegger, Z. Jin, C. Mller, J. Paulus, M. Reckhaus, J. lvarez Ruiz, P. Plger, and G. Kraetzschmar, "Johnny: An autonomous service robot for domestic environments," *Journal of Intelligent and Robotic Systems*, vol. 66, no. 1-2, pp. 245–272, 2012.

[4] D. Feil-Seifer and M. Mataric, "Defining socially assistive robotics," in *Proceedings of the 9th International Conference on Rehabilitation Robotics*, 2005, pp. 465–468.

[5] N. Roy, G. Baltus, D. Fox, F. Gemperle, J. Goetz, T. Hirsch, D. Margaritis, M. Montemerlo, J. Pineau, J. Schulte, and S. Thrun, "Towards personal service robots for the elderly," in *Proceedings of the Workshop on Interactive Robots and Entertainment (WIRE 2000)*, Pittsburgh, PA, May 2000.

[6] H.-M. Gross, C. Schröter, S. Müller, M. Volkhardt, E. Einhorn, A. Bley, C. Martin, T. Langner, and M. Merten, "Progress in developing a socially assistive mobile home robot companion for the elderly with mild cognitive impairment," in *Proceedings of the IEEE Conference on Intelligent Robots and Systems (IROS)*, 2011, pp. 2430–2437.

[7] A. Tani, G. Endo, E. F. Fukushima, S. Hirose, M. Iribe, and T. Takubo, "Study on a practical robotic follower to support home oxygen therapy patients-development and control of a mobile platform-," in *Proceedings of the IEEE Conference on Intelligent Robots and Systems (IROS)*, sept. 2011, pp. 2423 –2429.

[8] J. Brookshire, "Person following using histograms of oriented gradients," *International Journal of Social Robotics*, vol. 2, pp. 137–146, 2010.

[9] Z. Chen and S. T. Birchfield, "Person following with a mobile robot using binocular feature-based tracking," in *Proceedings of the IEEE Conference on Intelligent Robots and Systems (IROS)*, Oct. 2007.

[10] W.-H. Yun, D. Kim, and J. Lee, "Person following with obstacle avoidance based on multi-layered mean shift and force field method," in *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, oct. 2010, pp. 3813 –3816.

[11] J. Miura, J. Satake, M. Chiba, K. Ishikawa, K. Kitajima, and H. Masuzawa, "Development of a person following robot and its experimental evaluation," in *Proceedings of the 11th International Conference on Intelligent Autonomous Systems*, 2010, pp. 89–98.

[12] J. Satake and J. Miura, "Robust stereo-based person detection and tracking for a person following robot," in *ICRA Workshop on Person Detection and Tracking*, 2009.

[13] M. Tarokh and P. Merloti, "Vision-based robotic person following under light variations and difficult walking maneuvers," *Journal of Field Robotics*, vol. 27, no. 4, pp. 387–398, 2010.

[14] M. Bennewitz, W. Burgard, G. Cielniak, and S. Thrun, "Learning motion patterns of people for compliant robot motion," *International Journal of Robotic Research*, vol. 24, no. 1, pp. 31–48, 2005.

[15] B. D. Ziebart, N. Ratliff, G. Gallagher, C. Mertz, K. Peterson, J. A. D. Bagnell, M. Hebert, A. Dey, and S. Srinivasa, "Planning-based prediction for pedestrians," in *Proceedings of the IEEE Conference on Intelligent Robots and Systems (IROS)*, October 2009.

[16] C. Weinrich, M. Volkhardt, M. Einhorn, and H.-M. Gross, "Prediction of human avoidance behavior by lifelong learning for socially compliant robot navigation," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2013, pp. 376–381.

[17] N. Pradhan, "Mobile robot navigation for person following in indoor environments," Ph.D. dissertation, Clemson University, Clemson, SC, USA, 2013.

[18] N. Pradhan, T. C. Burg, S. T. Birchfield, and U. Hasirci, "Indoor navigation for mobile robots using predictive fields," in *Proceedings of the American Control Conference (ACC)*, June 2013.

[19] N. Pradhan, T. C. Burg, and S. T. Birchfield, "Robot crowd navigation using predictive position fields in the potential function framework," in *Proceedings of the American Control Conference (ACC)*, June 2011, pp. 4628–4633.

[20] T. Yoshimi, M. Nishiyama, T. Sonoura, H. Nakamoto, S. Tokura, H. Sato, F. Ozaki, N. Matsuhira, and H. Mizoguchi, "Development of a person following robot with vision based target detection," in *Proceedings of the IEEE Conference on Intelligent Robots and Systems (IROS)*, 2006, pp. 5286–5291.

[21] G. Doisy, A. Jevtic, E. Lucet, and Y. Edan, "Adaptive person-following algorithm based on depth images and mapping," in *Proceedings of the Workshop on Robot Motion Planning: Online, Reactive, and in Real-time, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2012)*, 2012.

[22] J. Shotton, A. W. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, and A. Blake, "Real-time human pose recognition in parts from single depth images," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 1297–1304.

[23] E. Rimon and D. Koditschek, "Exact robot navigation using artificial potential functions," *IEEE Transactions on Robotics and Automation*, vol. 8, no. 5, pp. 501–518, 1992.

[24] D. Nistér, O. Naroditsky, and J. R. Bergen, "Visual odometry," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004, pp. 652–659.

[25] P. J. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, Feb. 1992.

[26] A. Diosi and L. Kleeman, "Fast laser scan matching using polar coordinates," *International Journal of Robotic Research*, vol. 26, no. 10, pp. 1125–1153, Oct 2007.

[27] B. Peasley, S. T. Birchfield, A. Cunningham, and F. Dellaert, "Accurate on-line 3D occupancy grids using Manhattan world constraints," in *Proceedings of the IEEE Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 5283–5290.