in: Proc. KI 2005, 28th German Conference on Artificial Intelligence, Koblenz, 2005
LNAI 3698, pp. 320-334, © Springer-Verlag Heidelberg 2005, ISBN 3-540-28761-2

A Probabilistic Multimodal Sensor Aggregation Scheme Applied for a Mobile Robot *

Erik Schaffernicht, Christian Martin,
Andrea Scheidig, and Horst-Michael Gross

Ilmenau Technical University
Department of Neuroinformatics and Cognitive Robotics
andrea.scheidig@tu-ilmenau.de

Abstract. Dealing with methods of human-robot interaction and using a real mobile robot, stable methods for people detection and tracking are fundamental features of such a system and require information from different sensory. In this paper, we discuss a new approach for integrating several sensor modalities and we present a multimodal people detection and tracking system and its application using the different sensory systems of our mobile interaction robot Horos working in a real office environment. These include a laser-range-finder, a sonar system, and a fisheye-based omnidirectional camera. For each of these sensory information, a separate Gaussian probability distribution is generated to model the belief of the observation of a person. These probability distributions are further combined using a flexible probabilistic aggregation scheme. The main advantages of this approach are a simple integration of further sensory channels, even with different update frequencies and the usability in real-world environments. Finally, promising experimental results achieved in a real office environment will be presented.

1 Introduction

Dealing with Human-Robot-Interaction (HRI) in real-world environments, one of the general tasks is the realization of a stable people detection and the respective tracking functions. Depending on the specific application that integrates a person detection, different approaches are possible. Typical approaches use visual cues for face detection, a laser-range-finder for detection of moving objects, like legs, or acoustical cues for sound source detection.

Projects like Embassi [1], which aim to detect only the users' faces, usually in front of a stationary station like a PC, typically use visual cues (skin-color-based approaches, sometimes in combination with the detection of edge oriented features). Therefore, these approaches cannot be applied for a mobile robot which has to deal with moving people with faces not always perceivable.

* This work is partially supported by TMWFK-Grant #B509-03007 to H.-M. Gross and a HWP-Grant to A. Scheidig

In [2] a skin-color-based approach for a mobile robot is presented using an extension of particle filters to generate object configurations which represent more then one person in the image.

Other approaches trying to perceive the whole person rather than only her face use laser-range-finders to detect people as moving objects or directly by their legs, e.g. GRACE [3] or TOURBOT [4]. In [5] a approach based on particle representations in joint probabilistic data association filters is presented. Drawbacks of these approaches occur, for instance, in situations where a person stands near a wall and cannot be distinguished, in scenarios with objects yielding leg-like scans, like table-legs or chair-legs, or if the laser-range-finder does not cover 360 degrees of the robot space.

For real-world scenarios, more promising approaches combine more than one sensory channel,like visual cues and the scan of the laser-range-finder. An example for these approaches is the SIG robot [6], which combines visual and auditory cues. People are detected by a face detection system and tracked by using stereo vision and sound source detection. This approach is especially useful for scenarios like face-to-face interaction. Further examples are the EXPO-ROBOTS [7], where people are detected as moving objects by a laser-range-finder (resulting from differences from a given static environment map) firstly. After that, these hypotheses are verified by visual cues. Other projects like BIRON [8] detect people by using the laser-range-finder for detecting leg-profiles and combine these information with visual and auditory cues (anchoring). The essential drawback of these approaches is the sequential processing of the sensory cues. People are detected by laser information only and are subsequently verified by visual cues. These approaches fail, if the laser-range-finder yields no information, for instance, in situations when only the face of a person is perceivable because of leg occlusions.

Therefore, we propose a multimodal approach to realize a robust detection and tracking of people. Compared to other approaches, all used sensory cues are concurrently processed using a probabilistic aggregation scheme, that scales very well with the number of sensors and modalities used in terms of computational complexity. This way people are not only detected by only one feature. They can be detected by their legs and their faces or by only one of this features, respectively. The main advantage of our approach is the simple expandability by integrating further sensory channels, like sound sources, because of the used aggregation scheme.

The structure of this paper is as follows: first we present the employed different sensory modalities of our mobile robot for people detection and tracking: the omnidirectional camera, the laser-range-finder and the sonar sensors (section 2). Using these modalities, we generate specific probability-based hypotheses about the positions of detected people and combine these probability distributions by covariance intersection (section 3). Respective experimental results are presented in section 4 followed by a short summary and an outlook in section 5.

2 Mobile Interaction Robot Horos

To investigate respective methods, we use the mobile interaction robot Horos as an information system for employees, students, and guests of our institute. The system's task includes that Horos autonomously moves in the institute, detects people as possible interaction partners and interacts with them, for example, to answer questions like the current whereabouts of specific people.

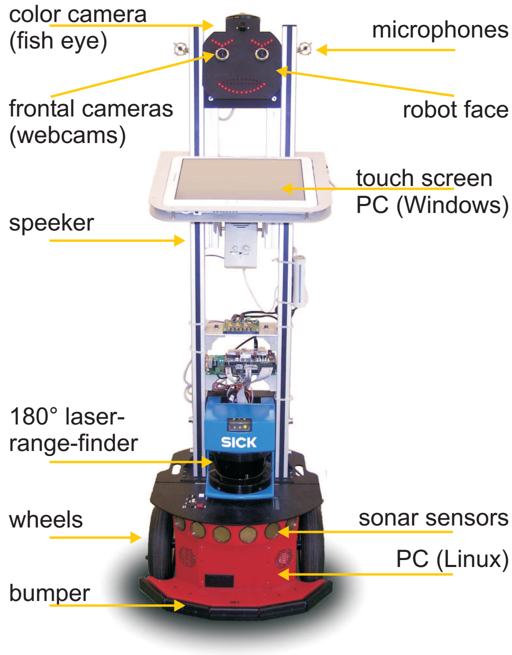

Fig. 1. Sensory and motory modalities of the mobile interaction robot Horos (HOme RObot System).

The hardware platform of Horos is an extended Pioneer-based robot from ActiveMedia. It integrates an on-board PC (Pentium M, 1.6 GHz, 512MB) and is equipped with a laser-range-finder (Sick) and sonar sensors. For the purpose of HRI, this platform was mounted with different interaction-oriented modalities (see Figure 1).

This includes a tablet PC (Pentium M, 1.1 GHz, 256MB) running under Windows XP for pen-based interaction, speech recognition and speech generation. It was further extended by a robot face which includes an omnidirectional fisheye camera and two microphones. Moreover, we integrated two frontal webcams for the visual analysis of dialog-relevant user features (e.g. age, gender, emotions).

Subsequently, only the omnidirectional camera, the laser-range-finder and the sonar sensors are discussed in the context of the people detection and tracking task.

2.1 Laser-based Information

The laser-range-finder of Horos is a very precise sensor with a resolution of one degree, perceiving the frontal 180 degree field of Horos (see Figure 2, left). It is fixed on the robot approximately 30 cm above the ground. Therefore, it can only perceive the legs of people (see Figure 2, right).

Based on the approach presented in [9], we also analyze the scan of the laser-range-finder for leg-pairs using a heuristic method. The measurements are segmented into local groups of similar distance values. Then each segment is checked for different conditions like width, deviation and others that are common for legs. The distance between segments classified as legs is pairwise computed to determine whether this could be a human pair of legs. For each pair found, the distance and direction is extracted.

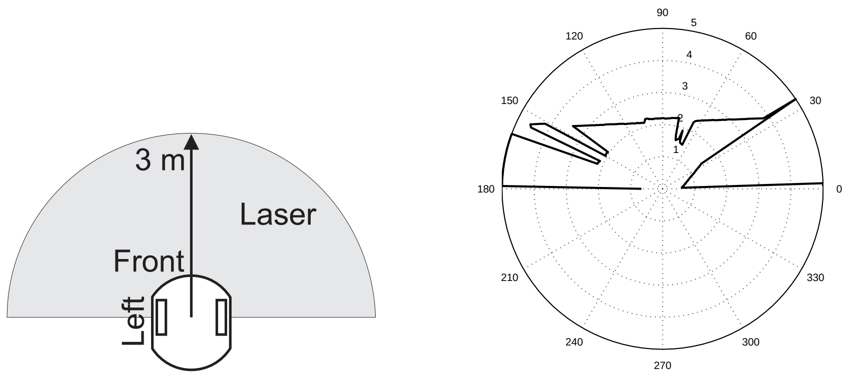

Fig. 2. Left: Top view of the schematic Horos. The sensory range of the laser-range-finder used to detect people is depicted grey. **Right:** Real scan of the laser-range-finder, depicted in polar coordinate system. In this situation, the robot is standing in a door directed to the top, perceiving a wall in front of it and the opened door to its right. Further, it senses a pair of legs in front of it (at 70°) and another one to its left (at 155°).

This yields very good results for distances of people which stand less than 3 meters away. In a greater distance legs will be missed due to the limited resolution of the laser-range-finder[1]. The strongest disadvantage of this approach is its false-positive classification detecting table-legs, chair-legs and also waste-paper baskets as legs. Also people standing sideway to the robot or wearing long

[1] At a distance of 3 meters the laser beams have a gap of more than 5cm between each other. In greater distances some legs are missed.

skirts do not yield appropriate values of the laser-range sensor to detect their legs.

2.2 Sonar Information

Horos is equipped with 16 sonar sensors, arranged at the Pioneer platform approximately 20 cm above the ground. The sound cones have an aperture angle of about 15 degrees. Because of this, a person detection using the sonar sensors does only work by analyzing the sonar scan for leg profiles (see Figure 3 right).

The disadvantage of these sonar sensors is their high inaccuracy. The measurement depends not only on the distance to an object, but also on the object's material, the direction of the reflecting surface, crosstalk effects when using several sonar sensors, and the absorption of the broadcasted sound. Because of these disadvantages, only distances of at most 2 meters can be considered for person detection using these sonar sensors (see Figure 3, left). This means the sonar sensors yield pretty unreliable and inaccurate values, a fact which has to be considered in the generation of a hypothesis of a person detection. For the purpose of a very simple person detection, we assume that all measurements less than 2 meters could be hypotheses for a person. These hypotheses could be further refined by comparing the position of each hypothesis with a given map of the environment. If the hypothesis would correspond to an obstacle in the map, it could be dismissed. The disadvantage of this refinement is, that people standing near by an obstacle would not be considered as a valid hypothesis.

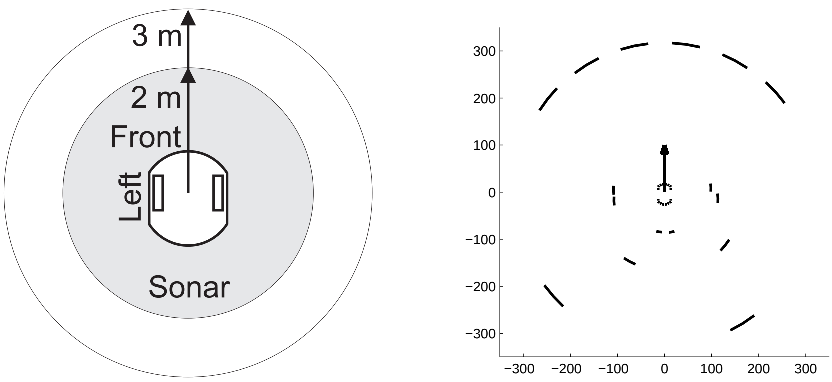

Fig. 3. Left: Top view of the schematic Horos. The sensoric range of the sonar sensors used to detect people is depicted grey. **Right:** Real values of the sonar sensors. In this case, the robot is standing in the middle of a floor directed to the top, sensing walls to the left and to the right and a person directly behind it. The distance values in front of the robot (dashed curve) are the result of our range limitation to a maximum of 2 meters.

2.3 Fisheye Camera

As third sensory system, we use an omnidirectional camera with a fisheye lens yielding a 360 degree view around the robot (see Figure 4 left). Because of the task of person detection, the usage of such a camera requires that the position of this camera is lower than the position of the faces. An example of an image resulting from the camera is depicted in Figure 4 (right). To detect people in the omnidirectional camera image, a skin-color-based multi-target-tracker [10] is used. This tracker is based on the condensation algorithm [11] which has been extended, so that the visual tracking of multiple people at the same time is now possible. The particle clouds used to estimate the probability of people in the omnidirectional image will concentrate on the different skin-colored objects. A problem is the possible detection and tracking of non-human skin-color-based objects, like wooden objects or cork pinboards. An essential advantage of this simple approach is, however, that it is much faster than subsampling the whole image trying to find regions of interest and its resistance to minor interferences, like partial occlusions.

A person detection using omnidirectional camera images yields good hypotheses about the direction of a person but not about his distance. Therefore, the integration of the different information from the camera with the data from the laser-range-finder and the sonar sensors should result in a more powerful and robust people detection and tracking system. Subsequently, the developed method for the combination of the several sensory systems will be introduced and discussed.

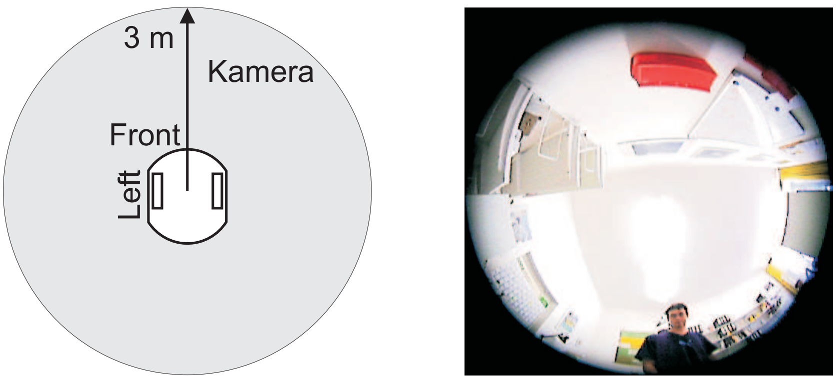

Fig. 4. Left: Top view of the schematic Horos. The sensoric range of the omnidirectional camera used to detect people is depicted grey. **Right:** Real image of the omnidirectional camera with a fisheye lens. A person to be detected is standing in front of the robot and can be seen at the bottom of the image.

3 The Multi-Modal Aggregation Scheme

At first, a suitable data representation for the aggregation of the multimodal hypotheses had to be choosen. The possibilities ranged from simple central point representation to probability distributions approximated by particles. The used aggregation scheme is based on Gaussian distributions, see section 3.1. Because of the unknown correlations between the different sensor readings, a *Kalman Filter* based approach was not used to combine these hypotheses. Instead *Covariance Intersection* is applied (section 3.2).

3.1 User Modeling Considering the Different Sensor Information

For the purpose of tracking, the information about detected humans is converted into Gaussian distributions $\phi(\mu, C)$. The mean μ equals the position of the detection in polar coordinates and the covariance matrix C represents the uncertainty about this position. The form of the covariance matrix is sensor-dependent due to the different sensor characteristics described in section 2. Furthermore, the sensors have different error rates of misdetections that have to be taken into account. All computation is done in the defolded cartesian r, φ space, see Figure 5. Examples for the resulting distributions are shown in Fig. 6.

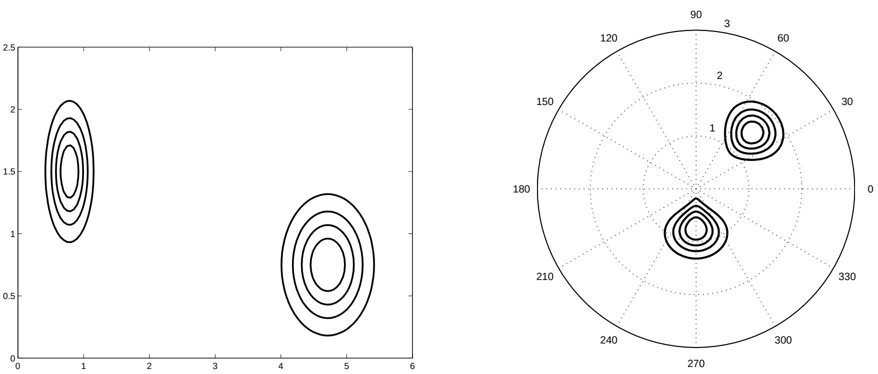

Fig. 5. Left: Two Gaussian hypotheses shown in a Cartesian r, φ system. **Right:** The same hypotheses in polar r, φ coordinates, the center indicates the position of the robot (Remark: computation is done in the cartesian space, while the polar r, φ space is used for better illustration).

Laser-based Information: Laser-range-finders yield a very precise measure, hence the corresponding covariances are small and the distribution is narrow (see Figure 6 left). The radial variance is fixed for all possible positions, but the variance of the angular coordinate is distance dependent. A sideways step of a person standing directly in front of the robot changes the angle more than the same movement in a distance of 2 meters. The smaller the distance of the detection the larger the variance has to be. The probability of a misdetection is the lowest of all used sensors, but the laser-range-finder only covers the front area of the robot due to sensor arrangement.

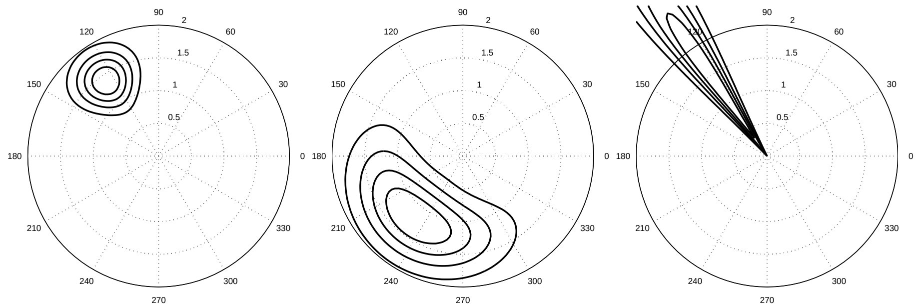

Fig. 6. Examples for generated hypotheses. The center of each plot represents the robot. **Left:** Hypothesis generated by laser showing a person left in front of the robot. **Middle:** Sonar-based information showing a hypothesis behind the robot. **Right:** Camera-based hypothesis without depth information showing the same person as in the left figure.

Sonar Information: Information from the sonar tends to be very noisy, imprecise und unreliable. Therefore, the variances are large and the impact on the certainty of a hypothesis is minimal (see Figure 6 middle). Nevertheless, the sonar is included to support people tracking behind the robot. So we are able to form an estimate of the distance in vision-based hypotheses.

Fisheye Camera: In contrast to the other sensors, the camera can only provide information about the angle of a detection, but not about the distance of a person. Therefore, for the radial variance of the distance coordinate a very large value was selected, with a fixed mean value (see Figure 6 right). The angular variance is determined by the angular variance of the particle distribution used in the visual skin-color based person tracker (see Section 2.3). The information content of a detection in the image of the fisheye camera is controlled by the position of the detection. In the front area of the robot, the influence is lower, because of the available laser as reliable sensor. Behind the robot, the image is the only source to get information about the presence of a person, the sonar has only supporting character. Thus, the relative weight of a visual hypothesis should be higher behind the robot.

The modeling and integration of additional sensor modalities, like sound localization or other features from the camera image, could be done in a similar way as described above.

3.2 Multi-Hypotheses Tracking Using Covariance Intersection

Tracking based on probabilistic methods attempts to improve the estimate x_t of the position of the people at time t. These estimates x_t are part of a local map M that contains all hypotheses around the robot. This map is used to aggregate the sensor hypotheses. Therefore, the movements of the robot $\{u_1, ..., u_t\}$ and the observations about humans $\{z_1, ..., z_t\}$ have to be taken into account. In other words, the posterior $p(x_t|u_1, z_1, ..., u_t, z_t)$ is estimated. This process is assumed to be Markovian. Then the probability can be computed from the previous state probability $p(x_{t-1})$, the last executed action u_t and the current observation z_t. The posterior is simplified to $p(x_t|u_t, z_t)$. After applying the Bayes rule, we get

$$p(x_t|u_t, z_t) \propto p(z_t|x_t)p(x_t|u_t) \ . \tag{1}$$

where $p(x_t|u_t)$ can be updated from $p(x_{t-1}|u_{t-1}, z_{t-1})$ using the motion model of the robot and the assumptions about typical movements of people.

A Gaussian mixture $M = \{\mu_i, C_i, w_i | i \in [1, n]\}$ is used to represent the positions of people, where each Gaussian is the estimate for one person. $\phi_i(\mu_i, C_i)$ is a Gaussian centered at μ_i with the covariance matrix C_i. The weight w_i ($0 < w_i \leq 1$) contains information about the contribution of the corresponding Gaussian to the total estimate.

Next, the current sensor readings z_t have to be integrated, after they have been preprocessed as described earlier. If M does not contain any element at time t, all generated hypotheses from z_t are copied to M. Otherwise data association has to be done to determine which elements from z_t and M refer to the same hypothesis. The Mahalanobis distance d_m between two Gaussians $\phi_i \in z_t$ and $\phi_j \in M$ is used as association criterion.

$$\begin{aligned} \mu &= \mu_i - \mu_j \\ C &= C_i + C_j \\ d_m &= \mu C^{-1} \mu^T \end{aligned} \tag{2}$$

This distance is compared to a threshold. As long as there are distances lower than the threshold, the sensor reading i and the hypothesis j with the minimum distance are merged. The problem of merging hypotheses in case two people pass near each other has to be tackled seperatly, confer e.g. [12]. The update is done via the *Covariance Intersection* rule (see [13] and [14]), a technique very similar to the *Kalman Filter*. As an advantage of this approach, the unknown correlations between the different sensor readings can be integrated, since this data fusion algorithm does not use any information about the cross-correlation of the inputs. A non-linear convex combination of the means and covariances is computed as follows

$$\begin{aligned} C_{new}^{-1} &= (1-\omega)C_i^{-1} + \omega C_j^{-1} \ . \\ \mu_{new}^{-1} &= C_{new} \left[(1-\omega)C_i^{-1}\mu_i + \omega C_j^{-1}\mu_j \right] \ . \end{aligned} \tag{3}$$

The weight ω is chosen to minimize the determinant as

$$\omega = \frac{|C_i|}{|C_i| + |C_j|} \ . \tag{4}$$

The more reliable distribution (that with the smaller determinant of the covariance matrix) is weighted higher in the update. If the current sensor reading is more certain than the current one, the resulting covariance of the hypothesis in M is reduced.

Sensor readings not matching with a hypothesis of M are introduced as new hypothesis in M. The weight w_i is representing the certainty of the corresponding Gaussian. The more sensors support a hypothesis, the higher its weight should be. If the weight passes a threshold, the corresponding hypothesis is considered to be a person. The weight is increased as

$$w_i(t+1) = w_i(t) + \alpha(1 - w_i(t)) \ , \tag{5}$$

if that hypothesis has been matched with a sensor reading. The time constant $\alpha \in [0, 1]$ is chosen with respect to the certainty of the current sensor (see section 3.1); the more reliable the sensor, the higher the α-weight is. In the case of an unmatched hypothesis, the weight has to be decreased.

$$w_i(t+1) = w_i(t) - (1 - \theta)\frac{t_{new} - t_{old}}{t_v} \ . \tag{6}$$

The term t_{new} is the current point of time and t_{old} the moment the last sensory input was processed. A person is considered to be lost if t_v seconds passed and no sensor has made a new detection that can be associated with this hypothesis. This temporal control regime is sensor dependent, too. Hypotheses with a weight lower than the threshold θ are deleted.

4 Application

The presented system is in use on the Horos robot in a real-world office environment. The fact of a changing illumination in different rooms and numerous distractions in form of chairs and tables is quite challenging.

Figure 7 shows a typical aggregation example. In this experiment, the robot was standing in the middle of an office room and did not move. Up to three people were moving around the robot. The enviroment contained several distracting objects, like table legs and skin-colored objects. No sensor modality was able to detect the people correctly. Only aggregation over sensors modalities and time led to the proper result.

The system was able to track multiple people correctly with an accuracy of 93 percent in the experiment. In most cases false negative detections occured behind the robot. The rate of false positive detections is higher, about every forth hypothesis was a misdetection. This is due to the simple cues integrated into the system. But for the intended task of Horos, the interactive office robot, it is considered to be more important not to miss to many people than finding to many. But there are ways to reduce the amount of false positive detections. Most misdetections are static in the environment, so based on the movement trajectories created by the tracker they can be identified (see section 4.1).

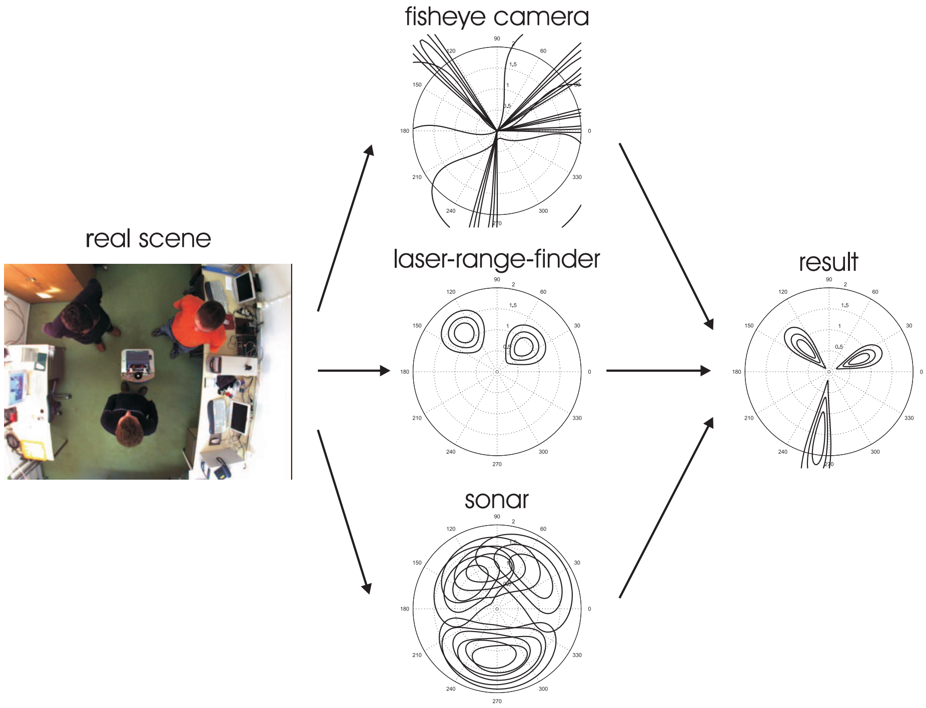

Fig. 7. Aggregation example. The left picture shows the real office scene from a bird's eye view. Three people are surrounding the robot, standing in the middle. The three figures in the middle show the current hypotheses generated by fisheye camera, laser-range-finder, and sonar from top to bottom. No sensor on its own can represent the scene correctly. The final picture displays the aggregated result from the sensors and the previous timestep. This is a correct and sharpened representation of the current situation.

Overall, the presented system improved the performance in the area behind the robot only slightly compared to a simple skin-color tracker. This is, because the sonar-based sensors do not provide many useful information for the tracking task. The main contribution of the sonar sensors is the addition of distance information to existing hypotheses extracted from the fisheye camera and preventing a precipitate extinction of hypotheses in cases of sudden changes in the illumination. In this case, the skin-color tracker will presumably fail, but if the sonar-based information still confirms the presence of the person at the respec-

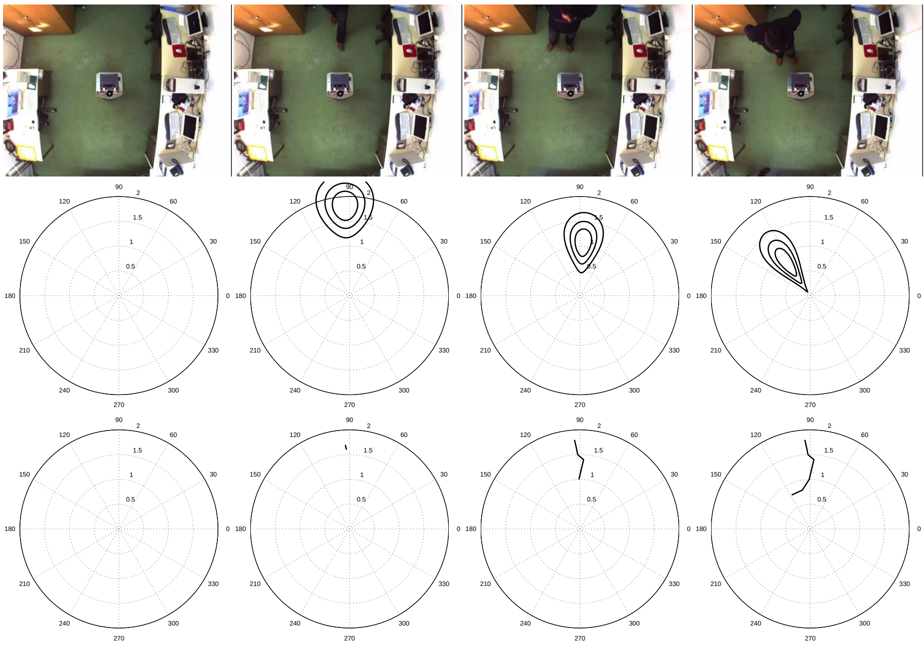

Fig. 8. Trajectory generation. The first row shows the real scene for four discrete time steps from a bird's eye view. In the second row the corresponding results of the tracking system are displayed. In the bottom row the generated trajectories are shown.

tive position, the hypothesis will not be deleted until the skin-color tracker has recovered.

In the front area of the robot, the system clearly outperforms single sensor-based tracking. Here the influence of the sonar on the result is not observable, because in most cases the laser-range-finder generates hypotheses more precisely. The laser reduces the deficiency of the skin-color tracker, while the skin-color based information compensates the shortcomings of the laser. These results are observable in Fig. 7. This leads to the assumption that the inclusion of additional sensory systems generating hypotheses about people (e.g. sound source hypotheses) will further improve the performance of this tracking system.

4.1 Trajectory Generation

The system was practically tested in the context of a survey task. Horos was standing in a hallway in our institute building. His task was to attract attention of people that came by. As soon as the system recognized a person near him, the robot addressed the visitor to come nearer. He then offered to participate in a survey about desired future functionality of Horos. The people tracking module was used to detect break offs, thus if the user was leaving before finishing the survey. The robot tried to fetch them back and finalize the survey. After the successfull completion of the interaction or a defined time interval with no person coming back, the cycle began again with Horos waiting for the next interaction partner. The experiment was made in the absence of any visible staff members, so the people could interact more unbiased.

These efforts are repeated from time to time to gather more information, and there is a second, not obvious, intention. The tracking module was used to generate typical movement trajectories of the users. Figure 8 shows the generation of such trajectories. In our future work, we will attempt to classify the path of movement to gain more knowledge about the potential user. In the context of adaptive robot behavior and user models, it is an important issue to assess the interaction partner. The users' movements and the positions relative to the robot are a fundamental step in this direction. If the robot can distinguish between people who are curious, but don't have the heart to step nearer, people who are in a rush and those how just want to interact with the robot, an appropriate reaction can be learned. The use of a multi-person-tracker is a prerequisite, since the experiments show visitors often appearing in groups of two or more people.

Examples for different trajectories are shown in Fig. 9. The most challenging aspects for a classification of trajectories are in our opinion the varying speed of the people and the search for typical movement schemes describing the interest of potential users. Based on the trajectories longtime immovable hypotheses can be discarded with respect to position and interaction status as a false detection.

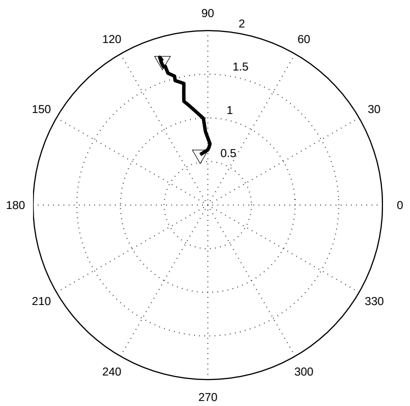

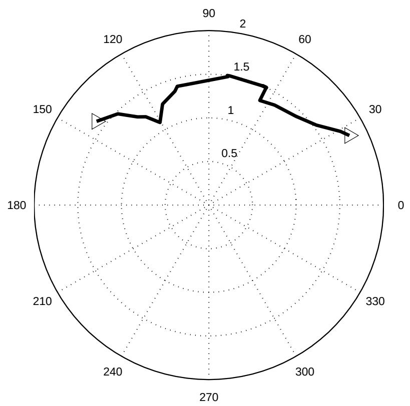

Fig. 9. Left: A trajectory showing a person coming straight towards the robot. **Right:** The person is crossing from left to right. In doing so, the robot is avoided. The varying time intervals between the movements aren't visible in the figure.

4.2 Guiding and Following

Another application of the tracking system in the context of our office scenario is the guidance function. The robot can guide visitors from the entrance of the building to the rooms of staff members. If the visitor is leaving sensor range, the robot stops and asks him to return. As soon as the tracker confirms the presence of the user the tour is continued. For this task, a multi-person-tracker is not mandatory, but it allows additional functionality, e.g., a group of people who are unintentionally blocking the robots' path can be detetected and can be asked to clear the path.

The system is able to master the inverted situation, to follow a person, too. This task is not difficult if the user faces the robot. In this case, however, the user has to move backwards, which is unnatural and possibly dangerous. Therefore, the task includes following the person even if the user turns around and no more skin-color is observable by the robot. Without the helpful information of the skin-color tracker, the system successfully follows the user using the hypotheses based on laser and sonar data.

5 Summary and Outlook

We presented a flexible multimodal probability-based approach for detecting and tracking people. It is implemented on our mobile office robot Horos and is functioning in real-time. Because of the sensor fusion and the probabilistic aggregation, its results are significantly improved compared to a single sensor tracking system. It can be easily extended with other sensors, because there is only the need for a new preprocessing module that produces appropriate Gaussian distributions based on the new sensor readings and an adaption of the weights that model the respective sensor characteristics. The system is able to aggregate data from input modalities with different update frequencies.

In our future work, we will extend the system with additional cues to further increase robustness and reliability for real-world environments. Currently, we are working on the integration of an audio-based speaker localization. In addition, it will be investigated if the face detector by Viola and Jones (see [15] and [16]) can be used for the verification of hypotheses (see [2]) and if it could be integrated into the aggregation scheme itself as an additional cue. Furthermore, we will study the behavior of our system compared to other known approaches and investigate the localization accuracy using labeled data of reference movement trajectories.

References

1. Froeba, B., Kueblbeck, C.: Real-time face detection using edge-orientation matching. In: Audio- and Video-based Biometric Person Authentication (AVBPA'2001). (2001) 78–83

2. Martin, C., Boehme, H.J., Gross, H.M.: Conception and realization of a multi-sensory interactive mobile office guide. In: IEEE Conference on Systems, Man and Cybernetics. (2004) 5368–5373
3. Simmons, R., Goldberg, D., Goode, A., Montemerlo, M., Roy, N., Sellner, B., Urmson, C., Schultz, A., Abramson, M., Adams, W., Atrash, A., Bugajska, M., Coblenz, M., MacMahon, M., Perzanowski, D., Horswill, I., Zubek, R., Kortenkamp, D., Wolfe, B., Milam, T., Maxwell, B.: Grace: An autonomouse robot for AAAI robot challenge. AAAI Magazine **24** (2003) 51–72
4. Schulz, D., Burgard, W., Fox, D., Cremers, A.: Tracking multiple moving objects with a mobile robot. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2001) 371–377
5. Schulz, D., Burgard, W., Fox, D., Cremers, A.: People tracking with a mobile robot using sample-based joint probabilistic data association filters. International Journal of Robotics Research (IJRR) **22** (2003) 99–116
6. Nakadai, K., Okuno, H., Kitano, H.: Auditory fovea based speech separation and its application to dialog system. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-2002). Volume 2. (2002) 1320–1325
7. Siegwart, R., Arras, K.O., Bouabdallah, S., Burnier, D., Froidevaux, G., Greppin, X., Jensen, B., Lorotte, A., Mayor, L., Meisser, M., Philippsen, R., Piguet, R., Ramel, G., Terrien, G., Tomatis, N.: Robox at expo.02: A large scale installation of personal robots. Special issue on Socially Interactive Robots, Robotics and Autonomous Systems **42** (2003) 203–222
8. Fritsch, J., Kleinehagenbrock, M., Lang, S., Fink, G., Sagerer, G.: Audiovisual person tracking with a mobile robot. In Groen, F., Amato, N., Bonarini, A., Yoshida, E., Krse, B., eds.: Int. Conf. on Intelligent Autonomous Systems, IOS Press (2004) 898–906
9. Fritsch, J., Kleinehagenbrock, M., Lang, S., Ploetz, T., Fink, G., Sagerer, G.: Multimodal anchoring for human-robot-interaction. Robotics and Autonomous Systems, Special issue on Anchoring Symbols to Sensor Data in Single and Multiple Robot Systems **43** (2003) 133–147
10. Wilhelm, T., Boehme, H.J., Gross, H.M.: A multi-modal system for tracking and analyzing faces on a mobile robot. In: Robotics and Autonomous Systems. Volume 48. (2004) 31–40
11. Isard, M., Blake, A.: Condensation - conditional density propagation for visual tracking. International Journal on Computer Vision **29** (1998) 5–28
12. MacCormick, J., Blake, A.: Probabilistic exclusion and partitioned sampling for multiple object tracking. International Journal on Computer Vision **39** (2000) 57–71
13. Julier, S., Uhlmann, J.: A nondivergent estimation algorithm in the presence of unknown correlations. In: Proceedings of the 1997 American Control Conference, IEEE (1997) 2369–2373 vol.4
14. Chen, L., Arambel, P., Mehra, R.: Estimation under unknown correlation: Covariance intersection revisited. IEEE Transactions on Automatic Control **47** (2002) 1879–1882
15. Viola, P., Jones, M.: Fast and robust classication using asymmetric adaboost and a detector cascade. In: NIPS 2001. (2001) 1311–1318
16. Viola, P., Jones, M.: Robust real-time object detection. In: Proc. of IEEE Workshop on Statistical and Computational Theories of Vision. (2001)