

Real-Time Fusion of Multimodal Tracking Data and Generalization of Motion Patterns for Trajectory Prediction

Martin Weser, Daniel Westhoff, Markus Hüser, Jianwei Zhang

Institute Technical Aspects of Multimodal Systems

Dept. of Informatics, University of Hamburg

Hamburg, Germany

{weser,westhoff,hueser,zhang}@informatik.uni-hamburg.de

Abstract—A sensor-based model of a service robot’s environment is a prerequisite for interaction. Such a model should contain the positions of the robot’s interaction partners. Many reasonable applications require this knowledge in real-time. It could for example be used to realize efficient path planning for delivery tasks. Additionally to the actual positions of the partners it is important for the service robot to predict their possible future positions. In this paper we propose an extensible framework that combines different sensor modalities in a general real-time tracking system. Exemplarily, a tracking system is implemented that fuses tracking algorithms in laser range scans as well as in camera images by a particle filter. Furthermore, human trajectories are predicted by deducing them from learned motion patterns. The observed trajectories are generalized to trajectory patterns by a novel method which uses Self Organizing Maps. Those patterns are used to predict trajectories of the currently observed persons. Practical experiments show that multimodality increases the system’s robustness to incorrect measurements of single sensors. It is also demonstrated that a Self Organizing Map is suitable for learning and generalizing trajectories. Convenient predictions of future trajectories are presented which are deduced from these generalizations.

I. INTRODUCTION

Service robots will increasingly support everyday work in business or home environments in the near future. Possible services are delivery tasks, cleaning services or home care. Recent developments in mobile robotics emphasize this [1], [2], [3]. However, the distribution and thus further development of mobile robots is mainly dependent on the acceptance of society. An important criteria for this acceptance is the robot’s ability to interact with the environment. Therefore it is essential to give the robot a detailed model of its environment, i. e. the location of its interaction partners. In general, this knowledge can only be generated using sensory input. An explicit specification of a dynamic environment is usually impossible.

The generation of such an environment model is a nontrivial task since the sensor calibration and measurements are susceptible to errors. In addition, the observed objects may be occluded and therefore not detectable by some sensors. We argue that this uncertainty about the environment decreases if different sensor-modalities with diverse qualities are used. The advantages of the different sensors can complement one another using appropriate fusing methods. In the following we present a scalable robust system that fuses arbitrary multimodal data containing trajectories of moving objects

using a particle filter. Furthermore we describe the different sensors’ preprocessing algorithms and explain their real-time capabilities.

Besides the interaction partner’s actual position, another important ability for a service robot is to predict people’s future positions. For example, the robot could intercept people for delivery tasks. Predictions about motions of dynamic objects are based on a motion model, which can be specified explicitly or learned from previous observed motions. As a reasonable application of our tracking system we developed a novel approach to predict people’s motion based on motion patterns learned by a Self Organizing Map (SOM).

The remainder of this paper is organized as follows: In section II an overview of existing research on trajectory prediction and multimodal tracking is given. Section III introduces the proposed framework for multimodal tracking. The laser- and camera-based tracking algorithms used in our implementation of the framework are presented. Preliminary experiments show the usefulness of the framework. The results of the experiments are used in section IV to explain the generalization of motion patterns and the experience-based prediction. Section V gives a conclusion.

II. RELATED RESEARCH

The research related to the proposed system for people tracking and trajectory prediction falls into three areas: people tracking in different sensor modalities, filtering and fusion of multimodal tracking data and motion prediction.

A. People Tracking

To increase the quality of the robot’s observations it is possible to fuse data of multiple sensor modalities. Two frequently used sensors for tracking applications are cameras and laser range finders. Due to the great amount of publications this paper cannot give a complete overview. We exemplarily present some selected papers which describe frequently used methods.

1) *Camera Tracking*: If static cameras are used, background subtraction is a common technique to separate foreground objects. In [12] every pixel is assigned a statistical color probability of the observed background by a mixture of gaussians. This probability is used to determine the pixel’s background membership in each frame. In [13] an occurrence model is used to track certain objects. This model contains

the color appearance and a probability mask, which represents the probability of each pixel to be part of the object. Another approach to track certain objects is presented in [14]: Here a color histogram represents the model. This approach uses no background subtraction, it is suitable for tracking people in images recorded by mobile cameras. Since this approach provides real-time tracking capabilities for non-stationary cameras it is applied in our system and described in section III in detail.

2) *Laser Tracking*: In [15], [16] and [17] background subtraction is used with laser range scans for object tracking. The systems differ in the way they generate the background. They all have in common that the background is modeled as a probability density function over the range measurements. For laser range finders mounted on a mobile robot it is not suitable to use such background models since the background measurements are changing permanently. Therefore, in [18] human legs are registered only via their size and shape.

B. Filtering and Fusion of Multimodal Tracking Data

Both the tracking algorithms in laser range scanners and in camera images mostly include filtering of the position estimates. This is reasonable since all measurements contain errors. The filtering becomes more important if several sensors are used. The integration and weighting of the different sensors is mostly included in the filtering algorithm. In [18] a Kalman filter is used to integrate tracking algorithms in laser range scans and camera images. A particle filter is presented in [19] which fuses audio and video information. [20] applies another approach which deals with multimodality. The different tracks gained by different sensor modalities are connected via an anchoring method.

C. Motion Prediction

Predicting the motion of objects is commonly used to avoid collisions in path planning tasks. There are numerous approaches present in the literature to predict obstacle motion in this context. Until a few years ago motion prediction was synonymous with the prediction of the object's position at the immediately following time step. This short-term prediction was mainly approached by modeling the object motion by a statistical process [4], [5], [6]. In [7] a Kalman filter is used to predict immediate following positions.

Even though several of these approaches might give accurate predictions, their actual benefit in relation to collision avoidance is questionable [8]. Short term predictions are also not sufficient to be used by higher level task planning processes. To predict object motions in a longer period of time, it is assumed that the motions are following determinate motion patterns.

Based on this idea of observable motion patterns, [9] clusters similar trajectories to trajectory prototypes by an expectation maximization algorithm. Partially observed trajectories are compared with these prototypes, to predict future motions of objects. It is assumed that the object is moving on the prototype trajectory which is most similar to the partially observed trajectory. Other methods [10], [11]

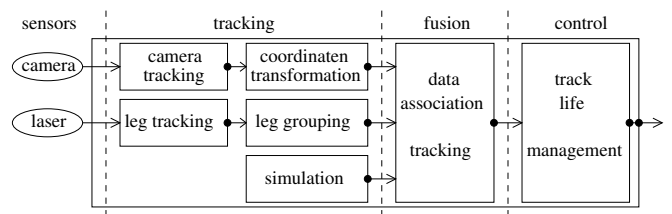


Fig. 1. Implemented components of the tracking system. The black dots symbolize the common interface described in section III-A.

use comparable techniques. The main difference lies in the method which is used to generate the trajectory prototypes. In [10] a pairwise clustering process is used. In [11] the target of the observed object's motion is estimated and the predicted path is generated by a path planning algorithm. The major drawback of that kind of systems is the impossibility to predict unusual trajectories.

III. PEOPLE TRACKING AND SENSOR FUSION

Both camera tracking as well as laser tracking have their own specific advantages and drawbacks. To build a robust and accurate tracking system, it is necessary to integrate several independent tracking algorithms working on the different sensors. With an appropriate fusion algorithm the specific advantages of the sensors could complement one another to decrease the overall error.

In the following, we introduce a conceptual framework that is able to deal with an arbitrary number of sensors. Subsequently, we describe how a particle filter algorithm is applied to fuse and filter data of different sensor modalities. Finally, the tracking algorithms for the measurements of laser range scanners as well as for camera images are presented. To provide the results in real-time both the filter algorithm and all used tracking algorithms need to be real-time capable.

A. Conceptual Framework for Multimodal Tracking

A typical multiple target tracking system consists of four blocks: sensor hardware, sensor processing and single sensor tracking, track fusion and association and track life management. A tracking system should be modular to allow the addition, removal and exchange of sensors and sensor processing algorithms. Therefore, the most important aspect of a tracking system is its ability to filter and fuse the results from individual sensors.

We developed a framework that contains the above-mentioned blocks. The schematic block diagram for this framework is shown in figure 1. Our implementation of the data association block includes filtering and data fusion using a particle filter. To ensure the extensibility of our system, we defined a general interface that has to be implemented by the modules for low-level sensor processing as well as higher-level modules. Through this interface an arbitrary number of tracks is provided where each track consists of

- the current position and velocity of the tracked object,
- the uncertainty about the position and velocity and
- a unique identity number.

The structure of this framework is mainly motivated by software technical encapsulation, substitutability of algorithms and extensibility concerning further sensor modalities. This structure provides efficient development, maintenance and testing capabilities.

B. Sensor Fusion and Filtering

We consider the problem of tracking as the detection of a state of a target. Therefore, we model the state x_t of a tracked person at time t as a four-dimensional vector $(x, y, \delta x, \delta y)^T$. This vector describes not only the position on the ground plain but also the velocity of the person. Since measurements of sensors contain errors it is impossible to derive the actual state of observed persons in a non-probabilistic way. Generally, a probability density function (*pdf*) is used to represent the state. Nonlinear Bayesian filtering can be applied to determine this *pdf* taking every previous measurement into account. The Bayesian solution to deriving a belief about the current state is a recursive discrete time approach. Since optimal solutions for nonlinear Bayesian tracking can only be applied when certain constraints hold we used the particle filter in our implementation. The particle filter is an approximate nonlinear Bayesian filter.

The particle filter is used in two stages of our tracking system. At first, it is used in the tracking algorithms applied on the data of a single sensor. After that, it is used for the fusion of tracking results made on different sensors. Other nonlinear filters like the extended Kalman filter can replace the individual particle filters in our system since each module is encapsulated and does not demand a particular algorithm.

The basis of the particle filter is the *importance sampling*. A multidimensional function $g(x)$ is factorized into two functions $g(x) = f(x)\pi(x)$, where $\pi(x)$ is interpreted as a probability density function with $\pi(x) \geq 0$ and $\int \pi(x) dx = 1$.

If a set of samples $\{x^i | i = 1, \dots, i = N\}$ with $N \gg 1$ and distributed according to $\pi(x)$ is generated, the integral of the function $g(x)$ can numerically be approximated as

$$\int g(x) dx \approx \frac{1}{N} \sum_{i=1}^N f(x^i) \quad (1)$$

When the density function $\pi(x)$ is unknown, $g(x)$ will be approximated by a function $\tau(x)$ similar to $\pi(x)$ regarding that $\forall x : \pi(x) > 0 \Rightarrow \tau(x) > 0$. Equation (1) reformulates to

$$\int g(x) dx = \int f(x) \frac{\pi(x)}{\tau(x)} \tau(x) dx \approx \frac{1}{N} \sum_{i=1}^N f(x^i) \omega(x^i) \quad (2)$$

with

$$\omega(x^i) = \frac{\pi(x^i)}{\tau(x^i)} \quad (3)$$

when the sample set x^i is distributed according to $\tau(x)$

Importance sampling is applicable to nonlinear estimation if a set of samples $\{x_t^i, \omega_t^i\}$ is chosen at each time t where $\omega_t^i \triangleq \omega(x_t^i)$ describes a weighting of the samples and $\sum_{i=1}^N \omega_t^i = 1$. If the sequence of measurements

$(z_0, z_1, \dots, z_t) = Z_t$ is known and $f(x) = 1$ this yields to

$$p(x_t | Z_t) \approx \frac{1}{N} \sum_{i=1}^N \omega_t^i \delta(x_t - x_t^i). \quad (4)$$

The weights ω_t^i follow from

$$\omega_t^i \triangleq \omega(x_t^i) = \frac{p(x_t^i | Z_t)}{\tau(x_t^i | Z_t)}. \quad (5)$$

For each new measurement the *pdf* describing the state of a tracked person is approximated by displacing each sample according to a probability function $\tau(x_t^i | x_{t-1}^i, Z_t)$. It is incidental that $\tau(x_t^i | Z_t) = \tau(x_{t-1}^i | Z_{t-1}) \tau(x_t^i | x_{t-1}^i, Z_t)$. The approximation is defined recursive as

$$\begin{aligned} p(x_t^i | Z_t) &= \frac{p(z_t | x_t^i, Z_{t-1}) p(x_t^i | Z_{t-1})}{p(z_t | Z_{t-1})} \\ &= \frac{p(z_t | x_t^i) p(x_t^i | x_{t-1}^i) p(x_{t-1}^i | Z_{t-1})}{p(z_t | Z_{t-1})} \end{aligned} \quad (6)$$

The weightings ω_t^i follow by insertion of (6) in (5).

$$\begin{aligned} \omega_t^i &= \frac{p(z_t | x_t^i) p(x_t^i | x_{t-1}^i) p(x_{t-1}^i | Z_{t-1})}{p(z_t | Z_{t-1}) \tau(x_{t-1}^i | Z_{t-1}) \tau(x_t^i | x_{t-1}^i, Z_t)} \\ &= \omega_{t-1}^i \frac{p(z_t | x_t^i) p(x_t^i | x_{t-1}^i)}{p(z_t | Z_{t-1}) \tau(x_t^i | x_{t-1}^i, Z_t)} \end{aligned} \quad (7)$$

The normalizing factor $p(z_t | Z_{t-1})$ is constant and can be precalculated for each measurement.

The use of measurements gathered on multiple sensors results in different measurement models. This affects $p(z_t | x_t^i)$ and $p(z_t | Z_{t-1})$ for each sensor.

See [21] and [22] for more detailed descriptions on Bayesian filtering.

C. Laser-based Tracking

As presented in section II, tracking algorithms which use laser range finders are often divided into two steps. In the first instance they generate a background model and then they determine whether the measurement belongs to the background. The background model is often represented by a histogram over the range measurements at each angle. It is assumed that the maximum of the histogram is caused by the background. This model is usually calculated in advance. The range measurements similar to the background distance are classified as background and discarded. This has the following drawbacks:

- The prior generation of the background model is time-consuming.
- If there are foreground objects present during the calculation of the background they are included in the background.
- These systems are unable to handle alterations of the background.

We developed a novel method to calculate the background's distances which is updated with each measurement.

The background distance $h_i(t)$ at time t and angle i is given by the following recursive equation:

$$h_i(t) = h_i(t-1) + \begin{cases} \varepsilon_1 & \text{if } h_i(t-1) < m_i(t) \\ (-\varepsilon_2) & \text{else} \end{cases} \quad (8)$$

The values of the increments ε_1 and ε_2 determine the adaptivity of the background model.

After background measurements are removed, groups of foreground measurements are tracked with a particle filter.

D. Camera-based Tracking

An overview of common methods for camera-based tracking is given in [23]. In our system we use the approach presented in [14] since it is suitable for cameras mounted on a mobile robot. In future work this property will be used.

People tracked in the camera image are represented by a weighted color histogram. Pixels are weighted with a monotone decreasing kernel function $K : \mathbb{R}^2 \rightarrow \mathbb{R}$ which assigns smaller weights to the pixels which are farther from the center of a detected person. If the size of a person is denoted by $2h^*$, the probability of the object's color u can be calculated as follows:

$$\hat{q}_u = C \sum_{x \in X^*} K\left(\frac{x}{h^*}\right) \delta(b(x) - u), \quad (9)$$

where C denotes a normalization constant. The function $b(x)$ assigns the pixel to an index of the histogram's color bin. Therefore, a person located at the coordinate y in the image plane is represented by a color histogram:

$$\hat{p}_u(y) = C_h \sum_{x \in X_h^y} K\left(\frac{y-x}{h}\right) \delta(b(x) - u) \quad (10)$$

where h is the size of the target candidate. As a measure of similarity between two color histograms we chose the *Bhattacharyya coefficient* [24].

The goal of the tracking algorithm is to find the location y with the highest similarity between the color histogram of a person and a candidate located at y . This is achieved using the sample mean shift method described in [14].

E. Experimental Results

For our experiments we used two SICK laser range finders mounted on a mobile service robot and a stationary camera at the laboratory of the TAMS institute. Due to the uncertainty of the camera tracking, which is caused by noisy measurements and changing illumination conditions we weighted the outcome of the laser tracking higher. In figure 2 the observed person's true trajectory is assumed as linear. Although the greater variance of the trajectory computed by the camera algorithm is obvious, the fused result has improved compared to the laser tracking result.

The camera-based tracking algorithm runs with the resolution of 640x480 pixel to achieve the aspired real-time capabilities. With a standard pc our implementation achieves 25 fps while tracking 3-4 targets. The laser-based algorithm reaches up to 30 fps due to the lower amount of data. Since the used particle filter is an efficient approximation of the

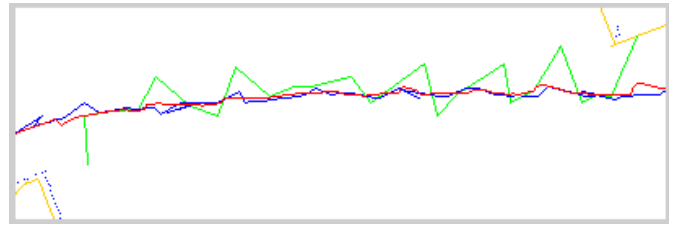


Fig. 2. Comparison of the sensor modalities: camera tracking (green) and laser tracking (blue) are fused by a particle filter (red). The greater variance of the camera tracking is obvious.

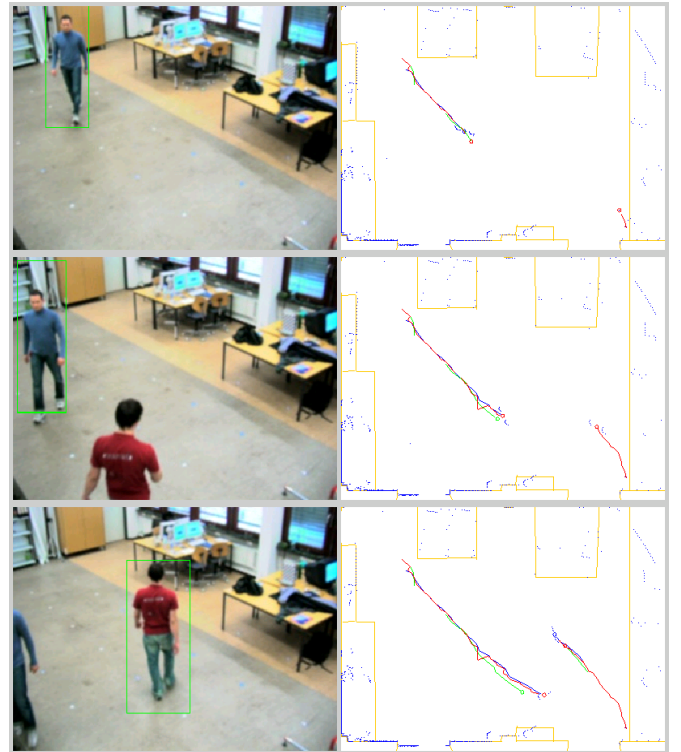


Fig. 3. Comparison of the sensor modalities: camera tracking (green) and laser tracking (blue) are fused by a particle filter (red). The additional blue points are the current measurements of the laser range finders.

bayesian filtering our system should be able to ensure real-time constraints if an appropriate environment i.e. RTLinux is used.

In a first experiment we used both sensor modalities to increase the accuracy and robustness of the tracking algorithm. Figure 3 shows a multimodal tracking of two persons.

In a second experiment the system observed 31 trajectories during a two-hour period of time. These trajectories are shown in figure 4. They will be used in the following section to explain the prediction with real-world examples.

IV. GENERALIZATION AND PREDICTION OF TRAJECTORIES

In recent years experience-based long term prediction of trajectories has become popular [9], [10], [11]. The systems



Fig. 4. The 31 trajectories observed during a two-hours experiment. The tracking starts when persons become visible to the sensors and ends when they leave the range of the sensors.

have in common that motions are first generalized to subsequently predict future trajectories. To generalize trajectories completely unsupervised we propose a novel approach using a SOM.

The following section describes the basics of the SOM algorithm and extensions needed to learn trajectory patterns. Afterwards, a prediction algorithm is presented, which is based on the learned trajectory patterns.

A. Learning of Trajectories using a Self Organizing Map

A SOM is a kind of artificial neuronal net. It is usually used to map statistical data from a high-dimensional input space onto a set of reference vectors of a usually lower-dimensional topological space. A major feature of a SOM is topology conservation with respect to the neighborhood of the input set. That means that similar input vectors (IV) are mapped to the same reference vector (RV).

In our system the SOM is used to learn the observed trajectories of people. The topology conservation is directly used to generalize these trajectories to motion patterns. A subsequent clustering of trajectories is omitted. However, it is necessary to make the motion patterns which are inherent in the SOM explicit to be usable for the prediction.

A SOM consists of a set of RV sometimes referred to as *nodes*. The RV are ordered in a topological space. In general, the alignment is arbitrary. We chose a quadratic structure to permit simple visualization and efficient storage in a two-dimensional array. The set of RV is depicted by $M = \{m^{ij} | m^{ij} \in \mathbb{R}^d\}_{i=1 \dots m_x, j=1 \dots m_y}$, where m_x is the width and m_y the height of the map. The dimension d of the RV and IV is equivalent.

During the learning phase the SOM is iteratively trained with the IV x^t . The RV $m^{cd} \in M$ most similar to x^t is referred to as response node. As a degree of similarity we chose the euclidian distance in the input space. The response node and nodes in its topological neighborhood adept to the IV. We chose the Manhattan distance as the distance function for the topological space.

The learning step for each node $m^{ij}(t)$ at the time t is defined by the following learning rule:

$$m^{ij}(t+1) = m^{ij}(t) + n_t(d)l(t)(x^t - m^{ij}(t)),$$

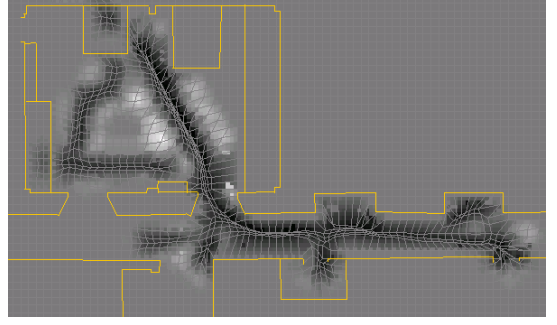


Fig. 5. This figure shows the distribution of the MAAR values calculated on the RV of converged SOM trained with real tracking results (shown in figure 4). The smaller the MAAR value the darker the color.

where $l(t)$ is the learning rate and $n_t(d)$ is a neighborhood function depending on the Manhattan distance between m^{ij} and the response node m^{cd} .

The neighborhood function is a kernel function which specifies a smoothing factor. To enable local learning it is necessary that $n_t(d) \rightarrow 0$ for all $d > k_t$, where k_t is the size of the neighborhood. For the convergence of the SOM it is necessary that $l(t) \rightarrow 0$ or $k_t \rightarrow 0$ for $t \rightarrow \infty$.

B. Extraction of learned Motion Patterns

After the SOM has converged the RV approximate the input data, i.e. the observed trajectories. To characterize frequently used paths it is self-evident to use a density estimation of the RV in the input space. Since the SOM algorithm does not generate new nodes, the environment of every local density maximum contains a local minimum. Transferred to the application of tracking, this means that it is unlikely for a person to stand beside a frequently used path. Since this does not agree with reality, we developed a novel method to characterize frequently used paths based on a SOM. Our method assigns to every node the value corresponding to the size of the minimum adjacent area MAAR which is spanned by four adjacent nodes. To provide the MAAR distribution the values are normalized and bilinearly interpolated in the ground plane. The smaller the MAAR value the more frequently a person was observed in this area. Figure 5 shows the MAAR distribution on the SOM trained with the real tracking results shown in figure 4.

To extract the motion patterns from the MAAR distribution we used non-maxima suppression with hysteresis threshold. The detected edges are transferred into a graph representation to permit the prediction described in subsection IV-C.

C. Experience-based Trajectory Prediction

The graph extracted from the MAAR distribution represents frequently used trajectories in a generalized form. After that the trajectory prediction assumes that tracked persons normally move along the graph segments.

1) *Mapping Observed Persons to Graph Segments*: To decide which motion pattern best matches an observed person a measure is required which defines a correlation between a

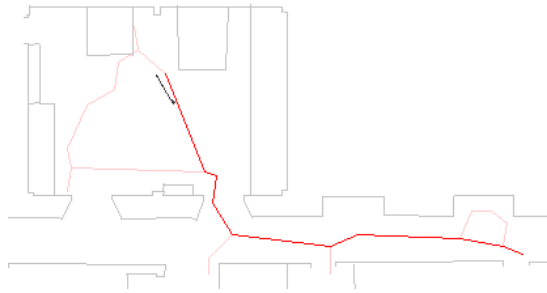


Fig. 6. The generalized motion graph extracted from the MAAR distribution of figure 5 using non-maxima suppression with hysteresis threshold (pink) and the prediction (red) of the trajectory of an observed person (black).

person and graph segment. We chose the euclidian distance between a person's position and the graph segments. The direction of the observed person's motion is not considered for mapping but is used later to determine the person's direction on the mapped motion pattern.

2) *Prediction*: Since experiments showed that the mapping is suitable, the motion pattern, i.e. the graph segment which the observed person is moving along, is known. With this, the long-term prediction is to determine the sequence of path segments the person will be following subsequently.

To predict this sequence, the probabilities for turns at graph branchings are learned. The long-term prediction of a trajectory adds adjacent path segments to the mapped segment with respect to the probability of a turn at a graph branching. The prediction terminates if no subsequent path segment is available or the overall probability, i.e. the product of the probabilities for each prior turn, comes under a certain threshold. Figure 6 exemplarily shows a prediction of the trajectory of an observed person.

V. CONCLUSION

In this paper we presented a framework where different sensor modalities can be integrated to build a multimodal tracking system. We implemented this framework for people tracking algorithms by processing laser range measurements and camera images. The individual tracking algorithms for the different sensors as well as the fusion module use particle filters and provide real-time capabilities. Due to the modular design other filter algorithms can be used. With some preliminary experiments we showed that the robustness of the tracking can be increased when camera- and laser-based tracking is combined.

As an example application we presented a new approach to generalize observed trajectories to motion patterns using a SOM in order to predict future trajectories. We tested our algorithm with more than 30 trajectories collected during a two-hour observation session and showed that our prediction technique achieves convenient trajectory predictions.

REFERENCES

[1] R. D. Schraft, J. Neugebauer, and C. Schaeffer and T. May. *Care-O-bot: Ein technisches Hilfesystem für unterstützungs- und pflegebedürftige Personen im häuslichen Bereich*, pages 234–244. Springer, 1998.

[2] S. Thrun, M. Bennewitz, W. Burgard, A.B. Cremers, F. Dellaert, D. Fox, D. Hähnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz. MINERVA: A second generation mobile tour-guide robot. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 1999.

[3] 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 *Proc. of the Workshop on Interactive Robotics and Entertainment (WIRE)*, Pittsburgh, 2000.

[4] Y. S. Nam, B. H. Lee, and M. S. Kim. View time based moving obstacle avoidance using stochastic prediction of obstacle motion. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 1081–1086, Minneapolis, 1996.

[5] S. Tadokoro, M. Hayashi, Y. Manabe, Y. Nakami, and T. Takamori. On motion planning of mobile robots which coexist and cooperate with human. In *Proc. of the IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 518–523 vol.2, August 1995.

[6] N. H. C. Yung and C. Ye. Avoidance of moving obstacles through behavior fusion and motion prediction. In *Proc. of the IEEE Int. Conf. on Systems, Man, and Cybernetics*, pages 3424–3429, San Diego, 1998.

[7] R. Madhavan and C. Schlenov. Moving object prediction for off-road autonomous navigation. In *Proceedings of the SPIE Aerosense Conference*, 2003.

[8] Amalia Foka. *Predictive Autonomous Robot Navigation*. PhD thesis, University of Crete, May 2005.

[9] Maren Bennewitz. *Mobile Robot Navigation in Dynamic Environments*. PhD thesis, University of Freiburg, Department of Computer Science, 2004.

[10] A. D. Vasquez and Th. Fraichard. Motion prediction for moving objects: a statistical approach. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 3931–3936, New Orleans, LA (US), April 2004.

[11] A. Bruce and G. Gordon. Better motion prediction for people-tracking. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2004.

[12] W. E. L. Grimson, C. Stauffer, R. Romano, and L. Lee. Using adaptive tracking to classify and monitor activities in a site. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 1998.

[13] A. W. Senior. Tracking with probabilistic appearance models. In *International Workshop on Performance Evaluation of Tracking and Surveillance Systems (ECCV)*, pages 48–55, June 2002.

[14] D. Comaniciu, V. Ramesh, and P. Meer. Real-time tracking of non-rigid objects using mean shift. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, Hilton Head Island, South Carolina, 2000.

[15] H. Zhao, R. Shibasaki, and N. Ishihara. Pedestrian tracking using single-row laser range scanners. In *Proc. of Computers on Urban Planning and Urban Management*, 2003.

[16] Ajo Fod, Andrew Howard, and Maja J. Mataric. Laser-based people tracking. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2002.

[17] H. Zhao and R. Shibasaki. Pedestrian tracking using multiple laser range scanners. In *Proc. of Computers on Urban Planning and Urban Management*, Sendai, Japan, May 2003.

[18] S. Feyrer and A. Zell. Robust real-time pursuit of persons with a mobile robot using multisensor fusion. In *Proc. of the 6th Int. Conf. on Intelligent Autonomous Systems (IAS-6)*, 2000.

[19] H. Asoh, F. Asano, T. Yoshimura, K. Yamamoto, Y. Motomura, N. Ichimura, I. Hara, and J.ire Ogata. An application of a particle filter to bayesian multiple sound source tracking with audio and video information fusion. In *Int. Conf. on Information Fusion*, pages 805–812, 2004.

[20] M. Kleinhagenbrock, S. Lang, J. Fritsch, F. Lömker, G. A. Fink, and G. Sagerer. Person tracking with a mobile robot based on multi-modal anchoring. In *Proc. of the IEEE Int. Workshop on Robot and Human Interactive Communication (ROMAN)*, pages 423–429, 2002.

[21] B. Ristic, S. Arulampalam, and N. Gordon. *Beyond the Kalman Filter*. Artech House, 2004.

[22] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, Cambridge, Massachusetts, London, England, 2005.

[23] James Russel. *Detecting Humans in Video Footage using Multiple Classifiers*. PhD thesis, University of Western Australia, 2004.

[24] T. Kailath. The divergence and bhattacharyya distance measure in signal selection. In *Proc. of the IEEE Int. Conf. on Trans. Commun. Tech. (COM-15)*, 1967.