Robust People Tracking and SHMM learning using SHMMs

Stephan Sehestedt, Sarath Kodagoda and Gamini Dissanayake
Centre for Autonomous Systems
Faculty of Engineering and IT
The University of Technology, Sydney. Australia
{s.sehestedt, s.kodagoda, g.dissanayake}@cas.edu.au

Abstract

For effective Human Robot Interaction (HRI), it is necessary for the robot to be aware of its human peers and be able to anticipate and predict their actions. This paper explores an improved strategy for people tracking using Sampled Hidden Markov Models (SHMM) for capturing common human motion patterns. Such an SHMM contains rich information about human spatial behavior and it can be learned online during robot operation. The proposed integration of people tracking and learning offers significant improvements to the outcomes when compared to existing techniques. Real world experiments that demonstrate the viability and effectiveness of the approach are presented.

1 Introduction

Successful Human Robot Interaction (HRI) requires a robot to have advanced abilities for undertaking complex tasks. One such ability is robust people tracking. People tracking has been identified as an important tool in HRI not only for safe operation [Schulz *et al.*, 2003] but also for collision avoidance [Bennewitz *et al.*, 2005] or to implement following behaviors [Gockley *et al.*, 2007], [Prassler *et al.*, 2002].

Predicting a person's future pose is a central problem in people tracking as it is difficult to model human motion sufficiently well. However, it has been observed that human motion follows place dependent patterns as it is influenced by a combination of social, psychological and physiological constraints [Hall, 1969] [Altman *et al.*, 1980] [Dean, 1996] [Arechavaleta *et al.*, 2006]. Therefore, a human motion model could be learned by a robot observing its environment. The learned model could subsequently be used to improve the robustness of people tracking.

Approaches to the learning of place dependent motion pattern models have been proposed in the past. Bennewitz et al. [Bennewitz *et al.*, 2005] used a network of laser scanners to learn the motion patterns of individual occupants of an office environment. After collecting a data set, EM was used to extract and cluster trajectories and the result was used to build a Hidden Markov Model (HMM). Using the model, the robot could anticipate future motion and improve its collision avoidance behaviors. Also using multiple laser scanners and off-line learning, in [Kanda *et al.*, 2009] it was proposed to learn activity patterns of people in a shopping mall to identify people who might be interested in entering a shop. A robot could then approach potential customers. Using stationary laser scanners, in [Luber *et al.*, 2011] it is proposed to learn a spatial affordance map based on a Poisson process. This is then exploited in a Multi Hypothesis Tracking (MHT) offering improved tracking capabilities and robustness. Vasquez et al. [Vasquez *et al.*, 2009] proposed Growing Hidden Markov Models to learn motion patterns. While on-line learning was shown to be possible, the observer had to be stationary. Lookingbill et al. [Lookingbill *et al.*, 2005] used a small helicopter with a camera as a semi-stationary observer to monitor a single roundabout to build an activity histogram. This histogram was shown to notably improve the effectiveness of their target tracking algorithm. Bruce and Gordon [Bruce and Gordon, 2004] used training data for their approach to learn goal locations in a small environment. A path planning algorithm provided a particle filter based tracker with an improved prediction model,

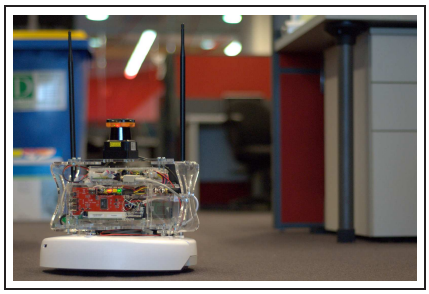

Figure 1: The robot LISA.

which was shown to perform better than using a Brownian motion model for the prediction step.

In [Altman *et al.*, 1980], however, it was noted that human motion patterns may change significantly over time, for example when there are changes in the environment. This led us to the conclusion that on-line learning using a mobile robot is preferable as this avoids the need for the installation of many sensors in the environment.

Therefore, here we propose an extension to motion pattern learning based on SHMMs [Sehestedt *et al.*, 2010] to simultaneously include people tracking to further exploit the potential of this idea. The effectiveness of people tracking algorithms can be greatly improved if a learned motion pattern model is available, therefore on-line learning of such a model allows a robot to improve its performance over time. Moreover, changes to the environment and resulting changes in the behavior of people who inhabit this environment can be incorporated into the model and thus, the robot can easily adapt to such changes. Furthermore, improved tracking robustness also leads to an improved learning performance. The ability to adapt is one step towards the ability of a robot to be integrated with humans as a co-worker. Ideas presented in this paper are evaluated by experiments conducted using our robot LISA shown in Fig. 1.

This paper is organized as follows. Section 2 gives a brief introduction to SHMMs for on-line learning of common human motion patterns. Section 3 introduces an alternative formulation for SHMM learning, leading to simultaneous people tracking and motion pattern model learning to take advantage of on-line learning during tracking. Section 4 presents experimental results showing the viability and effectiveness of the proposed approach. Finally, Section 5 summarizes the contributions and briefly discusses current and future work.

2 Sampled Hidden Markov Models

In this section a brief introduction to SHMM learning is given. SHMMs provide a sparse representation of common motion patterns and can be learned on-line only using a mobile robot's on-board sensors. The importance of this can be seen in Fig. 2, which illustrates a robot as a red circle and its current laser scan is shown as a red outline. It can be observed how limited its perceptive field is compared to the size of the environment it operates in. Thus, it is essential to incrementally learn and be able to use partial observation to update the model.

An SHMM is defined by its states S and state transition matrix A where each state is represented by a set of weighted samples. Although, these sample sets can represent arbitrary probability distributions, for easier notation we define states by means of their mean μ and covariance Σ

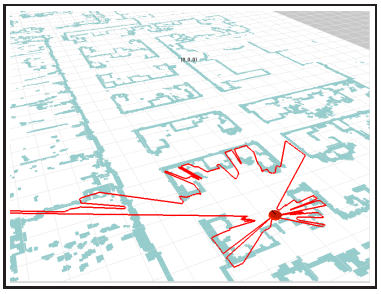

Figure 2: A small robot observing its environment.

$$S = s^{(i)} = \begin{bmatrix} \mu^{(i)} \\ \Sigma^{(i)} \end{bmatrix} \quad 1 \leq i \leq N \tag{1}$$

where N is the number of states in the model. Note that the model is actually time dependent as learning is done incrementally, however for easier notation this is omitted. The state transition matrix contains the probabilities for transitions from state i to state j as

$$A = a^{(ij)} = \begin{bmatrix} K^{(ij)} \\ P(s^{(j)}|s^{(i)}) \end{bmatrix} \begin{matrix} 1 \leq i \leq N \\ 1 \leq j \leq N \end{matrix} \tag{2}$$

where $K^{(ij)}$ is the number of times a transition was observed, from which the probability of the transition $P(s^{(j)}|s^{(i)})$ can be calculated. Learning an SHMM is done using a particle filter based people tracker. Consider the situation in Fig. 3(a) where a person walked along the trajectory indicated by the arrow. The tracking algorithm produces a series of sample clusters, each of which could be interpreted as one state of an SHMM. To obtain a more sparse representation of the observed trajectory a subset of those sample clusters was used to obtain the model shown in Fig. 3(b), where the states are shown as red covariance ellipses with red squares showing the means and the red lines illustrate the state transitions. The transitions are derived from the sequence of the sample clusters.

As the goal is to model common motion patterns, information has to be added to the model. Suppose another person is observed walking along the trajectory shown by the arrow in Fig. 3(c). Part of this trajectory overlaps with what was observed before. Using the symmetric Kullback-Leibler divergence (KLD) [Kullback and Leibler, 1951] is used to find associations between states in the SHMM and sample clusters obtained from people tracking as

$$\begin{aligned} KLD(s^{(i)} \parallel s^{(j)-}) = \; & KLD_s(s^{(i)} \parallel s^{(j)-}) \\ & + KLD_{s^-}(s^{(j)-} \parallel s^{(i)}) \end{aligned} \tag{3}$$

with $1 \leq i \leq N$ and $1 \leq j \leq K$, where K is the number of sample clusters (candidate states) $s^{(j)-}$ reported by the tracker. Whenever an association is found, the new information can be added to the according state and state transitions can be updated by counting. If for part of a trajectory no associations are found, new states may be added as can be seen in the figure.

A probabilistic formulation of the SHMM learning approach can be given as follows. A model of motion patterns D_t at time t is to be estimated based on the robot's localization estimate and people tracking. That is the belief

$$Bel(D_t) = P(D_t|\xi_t, \zeta_t, z_t, ..., \xi_0, \zeta_0, z_0) \tag{4}$$

is to be computed, where ζ_t is the localization hypothesis, ξ_t is the tracking result and z_t the sensor reading at time t.

From this an incremental update rule can be derived using the Bayes theorem

$$\begin{aligned} Bel(D_t) = \quad & P(\xi_t|D_t, \zeta_t, z_t, \xi_{t-1}, ..., \xi_0, \zeta_0, z_0) \\ & \cdot \frac{P(D_t|\zeta_t, z_t, \xi_{t-1}, ..., \xi_0, \zeta_0, z_0)}{P(\xi_t|\zeta_t, z_t, \xi_{t-1}, ..., \xi_0, \zeta_0, z_0)} \end{aligned} \tag{5}$$

Since the denominator is independent of D_t it is constant and can be substituted by

$$\eta = P(\xi_t|\zeta_t, z_t, \xi_{t-1}, ..., \xi_0, \zeta_0, z_0) \tag{6}$$

Thus,

$$\begin{aligned} Bel(D_t) = \quad & \eta P(\xi_t|D_t, \zeta_t, z_t, \xi_{t-1}, ..., \xi_0, \zeta_0, z_0) \\ & \cdot P(D_t|\zeta_t, z_t, \xi_{t-1}, ..., \xi_0, \zeta_0, z_0) \end{aligned} \tag{7}$$

This is the belief of D at time t given all past observations, sensor readings and observer poses. Obviously, this is not an efficient solution for updating the belief since all observations of moving people, sensor data and observer poses up until time t would have to be remembered. Therefore, it is assumed that observations and poses are conditionally independent of past observations and poses given ζ_t and D_t, i.e. the system is Markov

$$Bel(D_t) = \eta P(\xi_t|D_t, \zeta_t, z_t) \overbrace{P(D_t|\zeta_t, z_t, \xi_{t-1}, ..., \xi_0, \zeta_0, z_0)}^{\text{prior belief}} \tag{8}$$

Finally, the last term of this equation is the belief at time $t-1$ and thus the final update rule is written as

$$Bel(D_t) = \eta \overbrace{P(\xi_t|D_t, \zeta_t, z_t)}^{\text{People Tracking}} Bel(D_{t-1}) \tag{9}$$

This result allows the update $Bel(D_t)$ using the most recent observations of moving people.

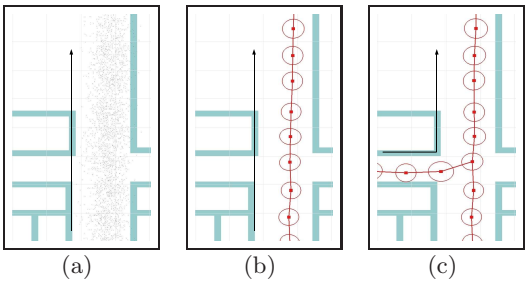

Figure 3: A basic example of SHMM learning.

3 Simultaneous People Tracking and Model Learning

In this section an alternative formulation for Simultaneous People Tracking and Motion Pattern Model Learning is presented. The idea is to tightly integrate learning and tracking in order to improve both. It is now the goal to estimate

$$P(D_t, \xi_t, \zeta_t|z_t) \tag{10}$$

which is the joint probability of the motion pattern model D_t, the position estimates of people in the robot's field of view ξ_t and the robot's location estimate ζ_t at time t given the latest sensor reading z_t. Assuming independence between these variables one can state

$$P(D_t, \xi_t, \zeta_t|z_t) = P(D_t|\xi_t, z_t)P(\zeta_t|z_t) \prod_{k=1}^{K} P(\xi_t^{(k)}|z_t) \tag{11}$$

Where K denotes the number of tracked people. The first term on the right hand side is also conditioned on ξ_t as this is what is needed to learn the model. Furthermore, it can be observed that the estimate of D_t is conditionally independent of z_t given ξ_t which is estimated from the latest sensor reading, leading to

$$P(D_t, \xi_t, \zeta_t|z_t) = P(D_t|\xi_t)P(\zeta_t|z_t) \prod_{k=1}^{K} P(\xi_t^{(k)}|z_t) \tag{12}$$

The second term on the right hand side defines robot localization and the estimate is commonly conditioned on control input given as u_t resulting in

$$P(D_t, \xi_t, \zeta_t|z_t, u_t) = P(D_t|\xi_t)P(\zeta_t|z_t, u_t) \prod_{k=1}^{K} P(\xi_t^{(k)}|z_t) \tag{13}$$

Intuitively, this equation could be solved from the right to the left, i.e. after people's states in the sensory field of the robot have been estimated this result could be used to improve localization. The localization data in turn would be used to determine the global positions of detected people. Moreover, the result could be used in the first term to update the model of motion patterns. However, to fully exploit the idea of an incrementally learned model of motion patterns D_t the simultaneous utilization and learning of the model is desirable. To accomplish this, the last term of above equation could be made dependent on D_t.

$$P(D_t, \xi_t, \zeta_t | z_t, u_t) = P(D_t|\xi_t) P(\zeta_t | z_t, u_t) \prod_{k=1}^{K} P(\xi_t^{(k)} | z_t, D_t) \tag{14}$$

Note that when people tracking is performed, although t has already incremented to the current discrete time step, D_t is at this point still $D_t = D_{t-1}$, which is accounted for by the notation as D_t.

We now have a formulation of SHMM learning which explicitly accounts for a mobile observer and takes advantage of improved tracking while learning. It remains to be shown how to use D_t to improve the people tracking performance.

3.1 Motion Prediction

For the lack of a general model of human motion, very conservative prediction models like Brownian motion or the constant velocity model are commonly used [Luber *et al.*, 2011]. Less conservative, learned models have been presented (e.g. [Bennewitz *et al.*, 2005] [Luber *et al.*, 2011] [Bruce and Gordon, 2004]), however, often the use of stationary observers or off-line learning limit the potential of the idea for mobile robots.

Here, we propose using SHMMs at the prediction stage of a particle filter [Arulampalam *et al.*, 2002] based people tracker as

$$P(\xi_t | \xi_{t-1}, D_t) \tag{15}$$

If a person's motion can be associated with an existing SHMM, then based on this the velocity of this person can be predicted by first calculating the time it would take to reach the predicted next state

$$\Delta T_p = \frac{\sqrt{(dx^2 + dy^2)}}{s} \frac{m}{m/s} \tag{16}$$

where dx and dy are the x and y distance from the current location to the predicted next state. With this result a prediction can be made based on the predicted velocities

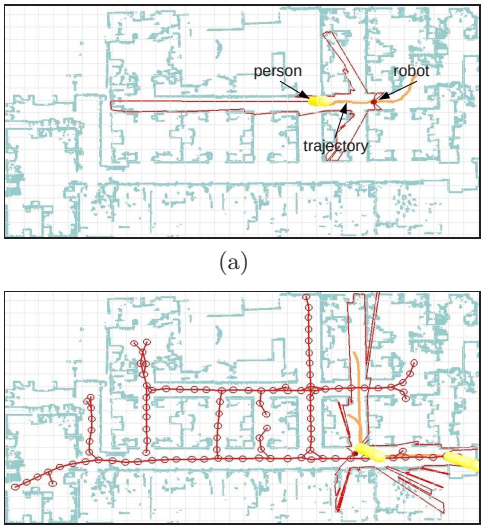

(a)

(b)

Figure 4: Learning a large model. (a) LISA moves in its environment while observing walking people. People tracking was accomplished using the presented method. (b) The observations were used to learn an SHMM of people's motion in the office. The SHMM eventually incorporated the information of more than 80 observed people.

$$\begin{aligned} v_x &= \frac{dx}{\Delta T_p} \frac{m}{s} \\ v_y &= \frac{dy}{\Delta T_p} \frac{m}{s} \end{aligned} \tag{17}$$

For the particle filter (PF) based tracker, it is then necessary to use model based prediction for a portion of the sample set. For the rest of the samples the near constant velocity model is used to account for model uncertainty. This uncertainty arises from possible outliers which may have contributed to the model as well as the fact that the model may still be incomplete due to incremental learning. Furthermore, there may be multiple possible next states and therefore, samples need to be assigned to all of those. These can either be fixed quantities or it could be proportional to the probability of the state transitions.

4 Experimental Results

All experiments were conducted using the robot LISA shown in Fig. 1. LISA is a low cost platform designed for fast and easy deployment. The base is an IRobot create equipped with a small Intel Atom X86-64 computer together with a Hokuyo UTM-30LX laser range scanner for perception.

The first experiment briefly demonstrates the learning of a large model in Fig. 4. LISA moves around in the office space, while sometimes following people. During the

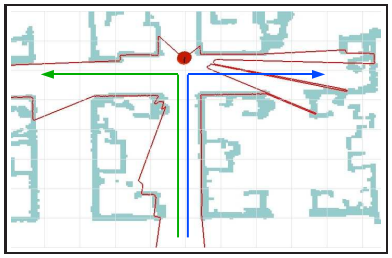

Figure 5: Experimental setup for the evaluation of the tracking performance.

experiment only two people could be observed at any one time, because of occlusions. This type of environment is difficult for many experiments and modifications cannot be done. The model's complexity grows as all available information is incorporated into the SHMM (Fig. 4(b)) and eventually contains the information of over 80 observed trajectories.

An experiment of this complexity, however, is unsuitable for the quantitative analysis of the tracking robustness and increased learning performance as a result of increased robustness. Therefore, the following experiments were designed around simple scenarios, excluding additional factors like localization errors. This brings the focus on the improved motion prediction based on learned SHMMs.

Fig. 5 shows the robot observing an intersection where people followed the indicated trajectories at a fast walking pace. This experiment was divided into four stages. In the first stage people followed the trajectory indicated by the blue arrow to evaluate the tracking performance. For the second stage a model of the blue trajectory was learned and used to predict motion. In the third stage the model from stage two was used, however, people now followed the green trajectory, i.e. they exhibited a drastically different behavior compared to the previous observations. In stage four of the experiment, the model was extended to include both trajectories and people could follow either trajectory. To evaluate the tracking performance 22 observed trajectories were used and the number of lost tracks was considered for each stage of the experiment.

Stage 1

The subjects were taking a sharp turn while being observed by LISA. As no prior learned SHMM was available, tracking was done using the near constant velocity model for motion prediction. Sample sizes of 50, 100, 200, 500 and 1000 were considered and the observed track loss rates are summarized in the first row of Table 1. Not surprisingly, it can be seen that the fewer samples were used, the more often the track was lost. While even with a low number of samples a person can be tracked when walking in a straight line, the error increases when the turn starts up to the point where a track may be lost. Only with many samples, the filter could keep track reliably.

Table 1: Number of track losses in 22 trials

Experiment	Sample Size 50	100	200	500	1000
Stage 1	13	8	5	2	0
Stage 2	3	0	0	0	0
Stage 3	22	12	9	4	1
Stage 4	2	0	0	0	0

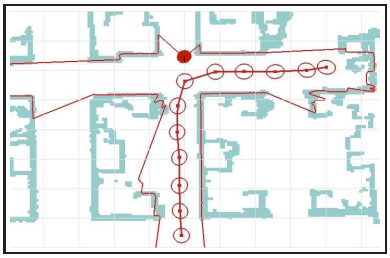

Figure 6: The SHMM in stage 2 of the experiment.

Stage 2

In stage 2 a model is learned first (see Fig. 6) and then used to improve motion prediction. Replaying the same data as in stage 1 of the experiment, it can be seen that the number of lost tracks is now much lower for all considered sample sizes in Table 1. With only 50 samples there were still some track losses. Generally it could be observed that during the turn the error was lower than in stage 1 and even tracking with only 100 samples produced good results.

Stage 3

In stage 3 of the experiment, the model from stage 2 was used, however, now people followed the green trajectory in Fig. 5. This resulted in people first behaving similar to the previously learned model, however, instead of the expected right turn a left turn occurred. Accordingly, there were more track losses again as shown in Table 1. In fact, it can be seen that 50 samples are not sufficient to handle such a situation. That there are more track losses compared to stage 1 is due to an obviously wrong prediction model, leaving the constant velocity model to try to recover. Furthermore, differences in people's

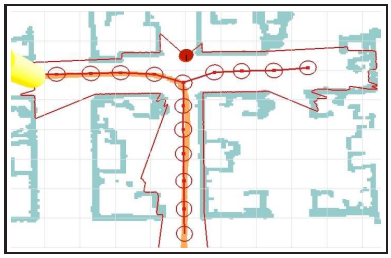

Figure 7: The SHMM in stage 4 of the experiment.

motion compared to stage 1 may have contributed.

Stage 4

Finally, in stage 4 the model from stage 2 was first extended to include the new behavior as shown in Fig. 7. Then 22 trajectories were used to evaluate the tracking performance with the updated SHMM. Table 1 shows that the number of track losses went down to the level of stage 2 again, indicating that even a PF with a low sample size is able to effectively use the more complex motion pattern model.

The findings in stage 1 and 3 of above experiment give rise to the next experiment: Learning in the presence of track losses. Starting with no motion pattern model, LISA observed people following the blue trajectory in Fig. 5 again. People tracking was done with a sample size of 100 in this experiment. Fig. 8(a) shows the tracking result for the first observed person, where the track was lost during the sharp turn. Fig. 8(b) illustrates the learning result, where it can be seen that the SHMM already covered most of the corner, albeit with some error (determined by visual inspection). Suspending model learning, the 22 test trajectories from the previous experiment's stage 1 were used again, finding that only 3 tracks were lost compared to 8 in the previous experiment. Hence, even this incomplete SHMM was already sufficient to improve the tracking performance.

Continuing learning, Fig. 8(c) shows a successful tracking result, which has some error during the turn. This result was used to update the model as shown in Fig. 8(d). By visual inspection, it was determined that the model exhibits some error compared to the real trajectories of the persons. However, with every following update the model converged closer to the average of the true trajectories of the test subjects as can be seen in Fig. 8(e) showing the model after 5 more observations.

For comparison, we configured the PF with 200 samples and disabled improved motion prediction. Six trajectories were replayed and successfully tracked while learning a model. The same was done with 100 samples and improved motion prediction enabled. The learned models are shown in Fig. 8(f), where visual inspection revealed that the latter model converged significantly closer to the average of the six trajectories, particularly in the section during and after the turn. Hence, SHMM based motion prediction is a great improvement over the constant velocity model in this experiment.

5 Conclusions

This paper presented a methodology for simultaneous people tracking and human motion pattern learning, based on on-line learning of Sampled Hidden Markov Models to represent human motion patterns. The integration of tracking and learning was shown to greatly improve both. The biggest improvement is the robustness of people tracking in difficult situations. The number of lost tracks could be greatly reduced even when using small sample sizes for tracking. These effects in turn led to a better convergence in motion pattern learning, which again highlights the benefits of on-line learning of human motion patterns.

It is difficult to collect enough data to make results statistically relevant. Therefore, more experiments have already been done in simulation and the above presented office environment (real world). The simulations serve to produce large amounts of data so that statistical relevance of the real world data can be inferred from this. Moreover, different kinds of motion with in complexity growing SHMMs have been studied. The outcomes of those experiments support the findings in this paper and will be published in follow up publications.

Future work encompasses a number of different aspects of the application of SHMMs as well as ongoing learning. Upcoming publications will also investigate SHMMs in people tracking as this is an important application with many different areas of interest in itself.

References

[Altman *et al.*, 1980] I. Altman, A. Rapoport, and J.F. Wohlwill. *Environment and culture*. Springer, 1980.

[Arechavaleta *et al.*, 2006] G. Arechavaleta, J-P. Laumond, H. Hicheur, and A. Berthoz. The nonholonomic nature of human locomotion: a modeling study. In *IEEE / RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics*, 2006.

[Arulampalam *et al.*, 2002] M. Sanjeev Arulampalam, Simon Maskell, and Neil Gordon. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50:174–188, 2002.

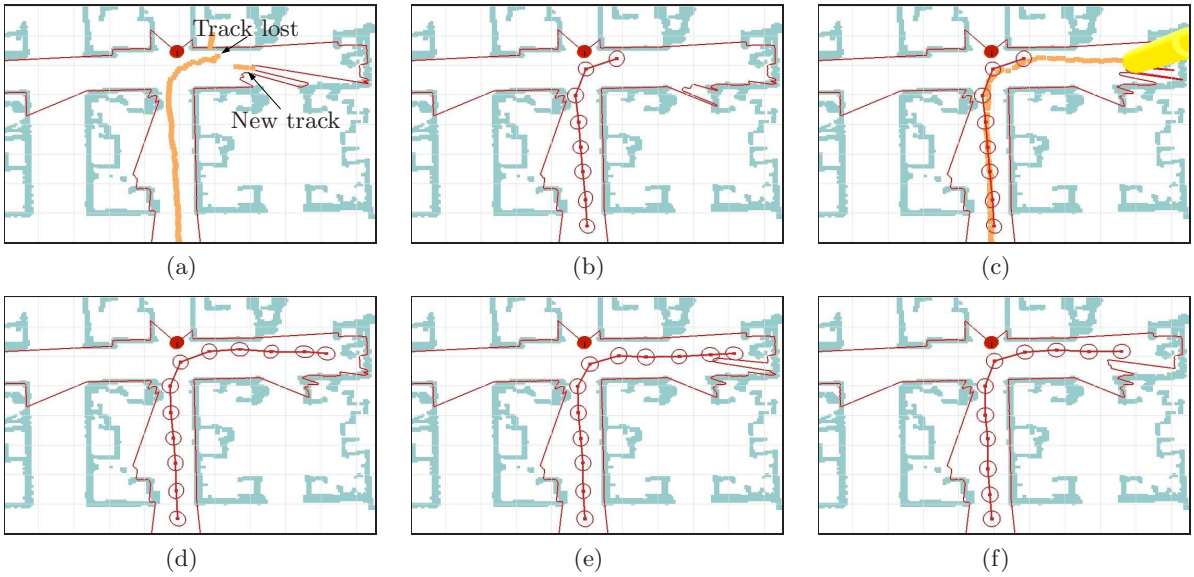

Figure 8: (a) - (d) Learning while track losses occur. (e) and (f) Model convergence.

[Bennewitz *et al.*, 2005] M. Bennewitz, Wolfram Burgard, G. Cielniak, and S. Thrun. Learning motion patterns of people for compliant robot motion. *International Journal of Robotics Research*, 24(1), 2005.

[Bruce and Gordon, 2004] A. Bruce and G. Gordon. Better motion prediction for people-tracking. In *IEEE International Conference on Robotics and Automation (ICRA)*. Citeseer, 2004.

[Dean, 1996] D. Dean. *Museum exhibition: theory and practice*. Routledge, 1996.

[Gockley *et al.*, 2007] Rachel Gockley, Jodi Forlizzi, and Reid Simmons. Natural person-following behavior for social robots. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 17–24, New York, NY, USA, 2007. ACM.

[Hall, 1969] Edward T. Hall. *The Hidden Dimension - Man's Use of Space in Public and Private*. The Bodley Head Ltd, London, 1969.

[Kanda *et al.*, 2009] T. Kanda, D.F. Glas, M. Shiomi, and N. Hagita. Abstracting people's trajectories for social robots to proactively approach customers. *IEEE Transactions on Robotics*, 25(6):1382–1396, dec. 2009.

[Kullback and Leibler, 1951] S. Kullback and R.A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, March 1951.

[Lookingbill *et al.*, 2005] A. Lookingbill, D. Lieb, D. Stavens, and S. Thrun. Learning activity-based ground models from a moving helicopter platform. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3948–3953, April 2005.

[Luber *et al.*, 2011] Matthias Luber, Gian Diego Tipaldi, and Kai O Arras. Place-dependent people tracking. *The International Journal of Robotics Research*, 2011.

[Prassler *et al.*, 2002] Erwin Prassler, Dirk Bank, and Boris Kluge. Motion coordination between a human and a mobile robot. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 2, pages 1228–1233, Dec. 2002.

[Schulz *et al.*, 2003] D. Schulz, W. Burgard, D. Fox, and A.B. Cremers. People tracking with mobile robots using sample-based joint probabilistic data association filters. *The International Journal of Robotics Research*, 22(2), 2003.

[Sehestedt *et al.*, 2010] S. Sehestedt, S. Kodagoda, and G. Dissanayake. Models of motion patterns for mobile robotic systems. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4127–4132, October 2010.

[Vasquez *et al.*, 2009] D. Vasquez, T. Fraichard, and C. Laugier. Growing hidden markov models: An incremental tool for learning and predicting human and vehicle motion. *The International Journal of Robotics Research*, 28(11-12):1486–1506, 2009.