The 2010 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 18-22, 2010, Taipei, Taiwan

Models of Motion Patterns for Mobile Robotic Systems

Stephan Sehestedt, Sarath Kodagoda and Gamini Dissanayake

Abstract— **Human robot interaction is an emerging area of research with many challenges. Knowledge about human behaviors could lead to more effective and efficient interactions of a robot in populated environments. This paper presents a probabilistic framework for the learning and representation of human motion patterns in an office environment. It is based on the observation that most human trajectories are not random. Instead people plan trajectories based on many considerations, such as social rules and path length. Motion patterns are learned using an incrementally growing Sampled Hidden Markov Model. This model has a number of interesting properties which can be of use in many applications. For example, the learned knowledge can be used to predict motion, infer social rules, thus improve a robot's operation and its interaction with people in a populated space. The proposed learning method is extensively validated in real world experiments.**

I. INTRODUCTION

Operating effectively in dynamic environments is one of the big challenges of mobile robotics with the unpredictability of human motion requiring sudden changes to planned tasks. Thus far, a common approach is to employ a method to minimize the impact of such events. This may be done by using sensors which are unaffected by moving objects, such as a camera which observes the ceiling [1]. Alternatively, tracking of dynamic objects allows segmentation of any sensor observations so that sensor data that is detrimental to the operation of tasks such as localization can be discarded [2]. This paper takes the view that prior knowledge of the motion of dynamic objects can be exploited in tasks such as path planning and human robot interaction.

Extracting motion patterns has attracted significant attention in the video surveillance literature where the interest is to identify suspicious behaviors by observing a scene. Here, one of the fundamental underlying assumptions is that the observer is stationary. Another common assumption is the observability of the whole trajectory. Algorithms based on these notions have been successfully implemented and presented in a range of publications including [3], [4] and [5].

In the field of mobile robotics the above assumptions usually do not hold thus requiring different strategies. The greater difficulty stems from the fact that mobile robots need to operate in expansive environments and are likely to encounter more diverse motion patterns. The problem will be further complicated by the location uncertainty of the moving observer. Furthermore, it is essential that incomplete observations of trajectories can be effectively incorporated.

Stephan Sehestedt, Sarath Kodagoda and Gamini Dissanayake are with the ARC Centre of Excellence for Autonomous Systems (CAS), The University of Technology, Sydney, Australia {S.Sehestedt, S.Kodagoda, G.Dissanayake}@cas.edu.au

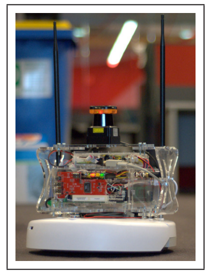

(a)

(b)

Fig. 1. a) The IRobot Create in its environment. b) The office space.

To our knowledge there are only a few publications that address these issues. Bennewitz et al. [6] developed a method to learn a model of dynamics in an office environment which was used for a mobile robot. It assumes the observer to be stationary and complete observability of trajectories. Furthermore, the models are learned off-line and knowledge of start and end locations is required. Vasquez et al. [5] proposed Growing Hidden Markov Models (GHMM) to incrementally learn motion patterns in an area. While this method allows for on-line learning, the method still relies on the strong assumption of utilizing a stationary observer.

This paper presents and discusses our approach to learn typical motion patterns in an environment based on Sampled Hidden Markov Models (SHMM). Here, a sample set is used to represent the dynamics in the environment, which is used to incrementally learn and dynamically update a Hidden Markov Model (HMM). In particular, we will focus on the SHMMs properties and possible applications. A strength of this method is that it is suitable for on-line learning on a moving observer, Consequently all presented experiments have been conducted using the mobile robot LISA (see Fig. 1(a)) in the office shown in 1(b).

The remainder of this publication is organized as follows. Section II briefly outlines a sampling procedure to learn a probability distribution of motion patterns. In section III we propose SHMMs to represent common motion patterns in an environment which can be learned on-line and unsupervised on a mobile robot. Key properties of the proposed model are discussed and experimentally demonstrated in Section IV. Finally, Section V presents a discussion, conclusions and future work.

II. SAMPLING MOTION PATTERNS

In a 2D environment, motion patterns can be described as a probability distribution over the x-y-θ location and velocity v. Discretizing the state space into a spatial grid followed by building a motion histogram [7] and then normalizing the values of the grid cells would result in an approximation of the joint probability distribution

$$P(x, y, \theta, v) \tag{1}$$

which represents the probability of the simultaneous occurrence of x-y-θ and v. This distribution constitutes the knowledge of all motion patterns in the environment independent of time. The distribution is very complex and thus a significant amount of data is required to estimate its parameters accurately using this simple strategy. Therefore, in [8] we proposed a sampling algorithm to incrementally learn an approximation of Eq. 1. Here we extend the idea to an efficient representation of motion patterns. In the following we briefly outline the proposed sampling procedure.

A mobile robot equipped with sensors for localization and object tracking, observes a person's trajectory. Tracking algorithms commonly represent each piece of a trajectory as a probability distribution from which it is possible to sample. In Fig. 2(a) a person (green rectangle) walks from the left to the right while being tracked. The samples are taken from the predictions of the tracking algorithm and are weighed according to the noisy observation. In Fig. 2(a), a 2D projection of the samples is shown along with the 95% confidence ellipses in x and y (green ellipses). Fig. 2(b) shows the sample set after more people moved along a similar trajectory.

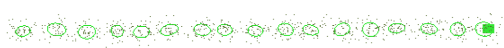

(a)

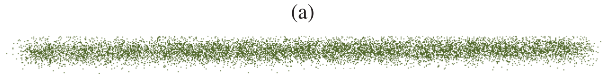

(b)

Fig. 2. a) The object (green rectangle) moved from the left to the right. The dark points denote samples generated from the tracker's prediction. The green ellipses denote the covariance after weighing the samples according to the most recent observation of the target b) The sample set after more objects were observed.

III. Sampled Hidden Markov Models

In this section we present our approach to learn Sampled Hidden Markov Models (SHMM) using the sampling algorithm outlined above. The challenge is to derive a model which can be learned and utilized by a mobile robot to improve its operation in a populated environment. Even though the sample based representation given in Section II is already computationally more efficient than a conventional grid based approach, the proposed SHMM reduces computational complexity even further by exploiting a sparse representation.

A. Hidden Markov Models

A Hidden Markov Model is a statistical model that represents a system as a directed graph. Here we briefly outline HMMs following the notation used by Rabiner [9]. HMMs are defined by N states of a system $S = s^1, s^2, ..., s^N$, observation symbols $V = v^1, v^2, ..., v^K$ with K being the number of symbols and state transition probability distribution $A = a^{ij}$, which is given as

$$a^{(ij)} = P(q_{t+1} = s^{(j)} | q_t = s^{(i)}), 1 \leq i \leq N \\ 1 \leq j \leq N \tag{2}$$

Furthermore, the observation probabilities in state j, $B = b^{ij}$ are formulated as

$$b^{(ij)} = P(v^{(i)} | s^{(j)}), 1 \leq i \leq K \\ 1 \leq j \leq N \tag{3}$$

Finally, the initial state distribution $\pi = \pi^i$ is defined as

$$\pi^{(i)} = P(q_1 = s^{(i)}), 1 \leq i \leq N \tag{4}$$

A common problem of HMMs is that there is no easy way to update these models over time [9]. A great variety of derivatives have been presented in the past to relax the issue and here we will briefly refer to the ones most relevant to the presented work

The idea of using HMMs to model trajectories is not new, however, only few publications are found in the domain of mobile robotics. The use of a hierarchy of HMMs to describe motion patterns on different levels was proposed by Liao et al. [10]. However, it requires the topology to be given and learning is done off-line. Vasquez et al. [5] propose Growing Hidden Markov Models for incremental learning of topology. However, its practical applicability in mobile robotics applications is limited due to the assumptions that are made. In particular, the method requires the observation of complete trajectories, meaning objects always have to be seen from the start of the path to the very end and the observer needs to be stationary at all times. In contrast, in the following section we will present an approach which allows to efficiently learn and update an HMM over time, which does not assume full observability of trajectories and can be used on mobile platforms.

B. Deriving a Sampled Hidden Markov Model

From the sampling algorithm given in Section II a particle cloud is obtained (as shown in Fig. 3(a)), which has the same temporal resolution as the sensor used for tracking, along with clustering information (i.e. it is a series of sample clusters, with each cluster representing the tracked objects pose and velocity at one point in time). This set of samples represents one persons trajectory as far as it has been observed. It is assumed that the observed process is a first order Markov process, i.e. motion at time t only depends on motion at $t - 1$.

1) Sampling The States and Transitions: From the algorithm given in Section II a vector of M clusters of weighted samples is obtained which describes an observed trajectory

$$C = \left[\ c^{(0)} \ c^{(1)} \ ... \ c^{(M)} \ \right] \tag{5}$$

To extract an HMM each of those clusters in C can be interpreted as a state of an HMM as

$$S = s^{(i)} = \left[\begin{array}{c} \mu^{(i)} \\ \Sigma^{(i)} \end{array} \right] \quad 1 \leq i \leq N \tag{6}$$

where $\mu^{(i)}$ and $\Sigma^{(i)}$ are mean and covariance of the i-th state and N is the number of states. Assuming zero states at the beginning, $N = M$ after adding C to the initially empty model. $\mu^{(i)}$ and $\Sigma^{(i)}$ are computed from the underlying sample set and thus represent a 4-dimensional distribution over x-y-$\theta - v$. In Fig. 3(a) a 2D projection of SHMM states is shown as red covariance ellipses in x and y. This figure also shows the learned model based on a single observed trajectory and the underlying samples.

The transition from state i to state j is given by the sequence of sample clusters and thus the transition matrix A consists of

$$A = a^{(ij)} = \left[\begin{array}{c} K^{(ij)} \\ P(s^{(j)}|s^{(i)}) \end{array} \right] \begin{array}{l} 1 \leq i \leq N \\ 1 \leq j \leq N \end{array} \tag{7}$$

where $K^{(ij)}$ is the number of times the transition was observed and $P(s^{(j)}|s^{(i)})$ is the probability of the transition. Naturally, the probabilities of the newly learned transitions in this example are 1.

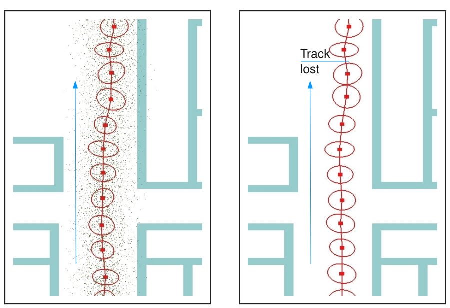

(a) (b)

Fig. 3. SHMM based model learning. a) Based on a single observed trajectory. b) Updated model after observing a second person. (real data)

2) Updating The Model: When observing another trajectory, a sequence of sample clusters is produced and a data association step will be carried out. This will associate the observations to the learned model, while initializing new nodes for previously unobserved parts. This data association is realized by comparing the pdfs in C to the model using the symmetrized Kullback-Leibler distance (KLD) [11]. The KLD is commonly used in literature to compare two probability distributions. The symmetrized KL-distance is defined as follows

$$\begin{aligned} KLD_{sym}(s^{(i)}|c^{(j)}) &= KLD(s^{(i)}|c^{(j)}) \\ &+ KLD(c^{(j)}|s^{(i)}), 1 \leq i \leq N \\ &\qquad\qquad 1 \leq j \leq M \end{aligned} \tag{8}$$

where $KLD_{sym}(s^{(i)}|c^{(j)})$ denotes the symmetrized KL-distance of state $s^{(i)}$ to cluster $c^{(j)}$ taking into account all N states and all M clusters of a trajectory. If an association is found between the i-th state and the j-th cluster, the cluster's samples belonging to j will be added to the state. To keep the number of samples used to model a state constant and to discard low weighted samples, a resampling procedure is employed. This is done similar to a normal particle filter with systematic resampling [12]. Finally, the transition probabilities are updated as

$$P(s^{(j)}|s^{(i)}) = \frac{K^{(ij)}}{\sum_{j=0}^{N} K^{(ij)}} \tag{9}$$

If a cluster could not be associated to an already existing state of the SHMM, it is added as a new state and the state transition matrix A gets extended accordingly.

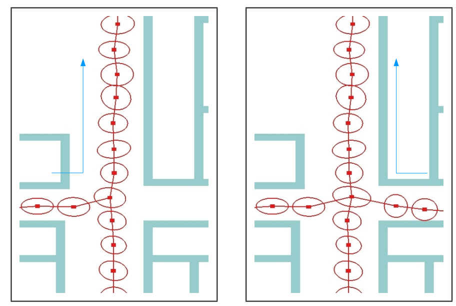

(a) (b)

Fig. 4. SHMM based model learning for trajectory adding. a) A person is tracked coming from the left and then walking along the previously learned path. b) Same situation with another person coming from the right. The robot (not visible) changed its location during observations. (real data)

To update the state transitions the knowledge about the sequence of newly added and associated states can be exploited. When a transition is already known $K^{(ij)}$ can be incremented and the transition probabilities can updated accordingly. It is to be noted that the second person was only observed until the track lost point given in Fig. 3(b). Therefore, only the observed states were updated.

In contrast, Fig. 4(a) shows the SHMM after another person was observed coming from the left, again following the trajectory indicated by the blue arrow. It can be seen that new states were added coming from the left and that a transition from the new part of the model to the former model was learned. The situation is similar in Fig. 4(b) where another person was tracked coming from the right and the model is updated accordingly.

C. Using the SHMM On a Mobile Robot

In order to implement the SHMM on a mobile robot, several other issues need careful investigation. Most importantly, computational complexity should be kept at a minimum and the model learning should not be affected by limited observability.

1) Complexity Considerations: Here we consider the efficient implementation strategies to manage the computational complexity. In the proposed SHMM the pdf is of 4 dimensions (x, y, θ, v). Processing Gaussian distributions with 4-

dimensions can be a computationally demanding task especially for a mobile robot. Therefore, it is proposed to exploit the structure of the SHMM to reduce the dimensionality to 2.

Firstly, there are two possible ways to exclude v from the state estimate. It is possible to either sample the trajectory data with a fixed frequency or by utilizing a binning approach. When using fixed frequency sampling, the distances between clusters inherently capture the speed of the observed object, thus implicitly representing v. In binning, speed domains can be chosen and for each domain a distance between successive clusters is defined. Our current implementation uses the latter method as it is less vulnerable to timing inaccuracies.

Secondly, the target orientation can be excluded from the state estimate. Considering the 2 dimensional structure of the SHMM as in Fig. 4(b), it can be seen that the expected orientation can be derived from the relative locations of successive states. Moreover, when there are multiple transitions a probability for the matching headings can be obtained using the transition probabilities. Hence, it is decided to drop the explicit use of θ without loss of information. This results in a 2-dimensional Gaussian distribution for describing a state.

2) Limited Observability: In general, mobile robots are not capable of observing the whole operating environment due to sensory limitations and occlusions. Depending on their routes, some parts of the environment can be over exposed, whereas the other parts can be poorly explored. Therefore, determination of transition probabilities based on observation alone are erroneous due to different exposure times of various parts of the environment. Further, the value $K^{(ij)}$ cannot be used as a measure of traffic density, as the relation between the values of different transitions is not known.

To overcome this problem, the overall time of observation of a part of the model has been added to the state's transitions as

$$a_{ij} = \begin{bmatrix} \Delta T^{(ij)} \\ K^{(ij)} \\ P(s^{(j)}|s^{(i)}) \end{bmatrix} \tag{10}$$

Where $\Delta T^{(ij)}$ denotes the total time in which this transition could have been observed, i.e. the time this area was inside the field of view of the observer. Consequently, the transition probabilities are computed as

$$P(s^{(j)}|s^{(i)}) = \frac{K^{(ij)}/\Delta T^{(ij)}}{\sum_{k=0}^{n} K^{(ik)}/\Delta T^{(ik)}} \tag{11}$$

Where n is the number of outgoing transitions from state i.

IV. Experimental Results

All Experiments were conducted using the LISA (Lightweight Integrated Social Autobot), based on an iRobot Create, which carries a Hokuyo UTM-30LX laser range finder and a small Intel Atom based computer (see Fig.1(a)) for localization and people tracking. The environment is an open office space of approximately $20 * 25m$ as shown in Fig. 1(b) and Fig. 8(a).

We first present experiments devised for analyzing the model learning capability of the proposed SHMM. Then, we briefly present possible applications of the learned models of motion patterns for completeness. In the following figures the pose of the robot is shown by a green arrow, states of an SHMM are shown as red ellipses and the state transitions as red lines between the means of states.

A SHMM model of a typical human trajectory while navigating at a corner is shown in Fig. 5. The trajectory is smooth and does not contain sharp corners, which agrees with the human navigation literature [13].

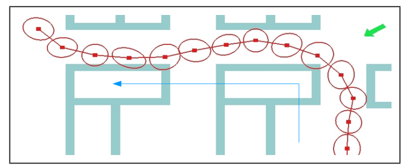

Fig. 5. A typical curved human trajectory modeled as an SHMM.

The next experiment was devised to assess the adaptability of the SHMM. In Fig. 6 the robot observed people walking from the bottom to the top of the image. After observing three people having a similar trajectory, an obstacle was placed on the way (Fig. 6(b). This caused people to slightly alter their trajectories. After observing 5 more persons, it was noted that the previously learned model was automatically adapted by the SHMM to accommodate the slight changes (Fig. 6(c)).

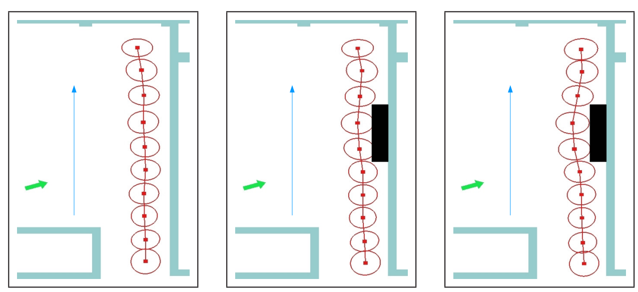

(a) (b) (c)

Fig. 6. People are moving from the bottom to the top while a robot learns the patterns. A) The initial model. b) An obstacle is introduced. c) The model converged to a slightly different shape.

Next a larger obstacle was placed on the general human trajectory, so that people have to take a substantial detour (Fig. 7). Due to the notable differences in the observations (decided by the KL-distance between the model and the observation), it can be seen that a new trajectory is added to the model as a separate branch (Fig. 7(b)). After a few more observations, it could be noted that the transition probabilities of A to C became larger than that of the previously learned A to B transition. This is graphically represented in Fig. 7(b) and Fig. 7(c), where the thickness of the line joining the states A-C is thicker than that of A-B.

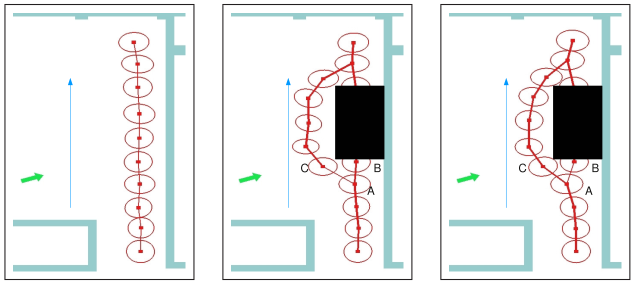

(a) (b) (c)

Fig. 7. People are moving from the bottom to the top while a robot learns the patterns. a) Initial trajectory without obstacles, b) An obstacle is introduced, c) The model adapted and converged according to the new information.

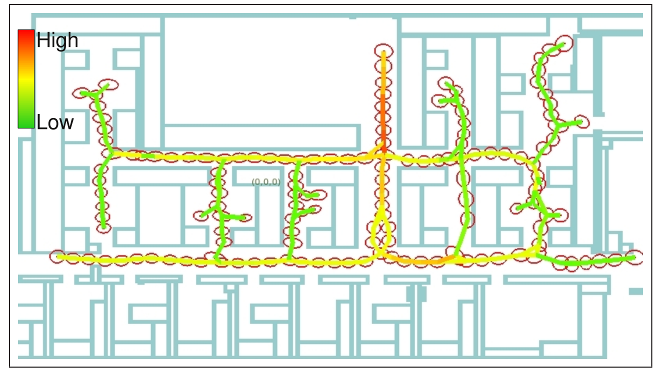

Fig. 9. The observed traffic density; colors range from green (low traffic density) to red (high traffic density).

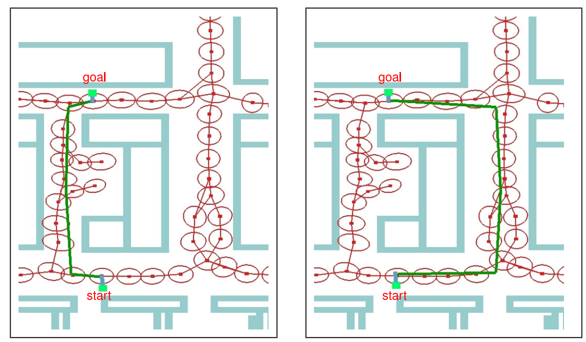

(a) (b)

Fig. 10. Path planning results, a) A path generated using basic $A*$ algorithm, b) Socially compliant path generated with the integration of the learnt model.

The next experiment was devised to demonstrate the capability of learning larger environments with complex motion patterns. In an office type of environment consisting of cubicals, LISA is used to observe various motion patterns while moving around. Starting from the first most observation (Fig. 8(a)), the SHMM encompassed all of its observations to learn a more complicated motion patterns as given in Fig. 8(c). Although the model is learned with respect to the observations at any time, it may lack completeness. It happens due to unobserved patterns or tracking failure. This phenomenon can be seen in Fig. 8(b), where there is a discontinuity in the model (inside the dashed rectangle). Once the unobserved part of a trajectory has been observed, the model becomes continuous as can be seen in Fig. 8(c). Another interesting observation can be made inside the dashed circle in Fig. 8(c). It is an intersection with people arriving from two directions leading to two clothoid trajectories. Although these two trajectories seem to have a more complex structure than necessary, it is a natural phenomenon which often occurs in such narrow corners due to the phenomenon described in the first experiment.

Fig. 8(d) shows the learned Gaussian distributions of trajectories with which the states and transitions are represented. As the constraints of the structure (map) of the environment is not taken into consideration, there are some apparent overlaps of the distributions with obstacles, such as walls. With more observations this effect would be reduced due to decreasing uncertainty. Fig. 9 shows the traffic density as estimated by the model, where the magnitude is represented by a color scheme. It shows traffic density is higher in corridor areas rather than through cubicals, which belong to other occupants. This behavior is expected in a workplace where people avoid disturbing co-workers.

Finally, we briefly present path planning as a possible application of a model of motion patterns. Interested readers are referred to [14] for more detailed information. As mentioned above, generally people take alternate routes to avoid entering into others workspaces. If a robot needs to be integrated with humans, such qualities need to be learned. In long term deployment, this could be achieved by observing and learning human motion patterns. Without this knowledge, a robot will plan a path based on the shortest path criteria using the common A^* algorithm as shown in Fig. 10(a). However, it could be much more appropriate to use the knowledge of how humans navigate in the environment to plan the path. This could be achieved by extending the A^* algorithm to accommodate the learned model. More precisely, we can alter the cost function of the A^* algorithm for socially acceptable path planning by taking traffic density into account. The result of such a scenario is shown in Fig. 10(b), where it could be noted that the planned path is significantly longer than the shortest path, however the robot will not cross desk areas disturbing occupants.

V. CONCLUSIONS

In this paper we presented a novel method to learn a SHMM to represent motion patterns in an office like environment. It provides an on-line and unsupervised learning technique. Motion patterns can be described as a joint probability distribution over pose and velocity from which samples can be taken. To avoid the computational complexity of using a sample distribution, a Hidden Markov Model based representation to learn common motion patterns was proposed.

The SHMM obviously has a lower memory footprint than a sample distribution, since we can easily reduce the resolution of the model. This approach is valid as we are interested

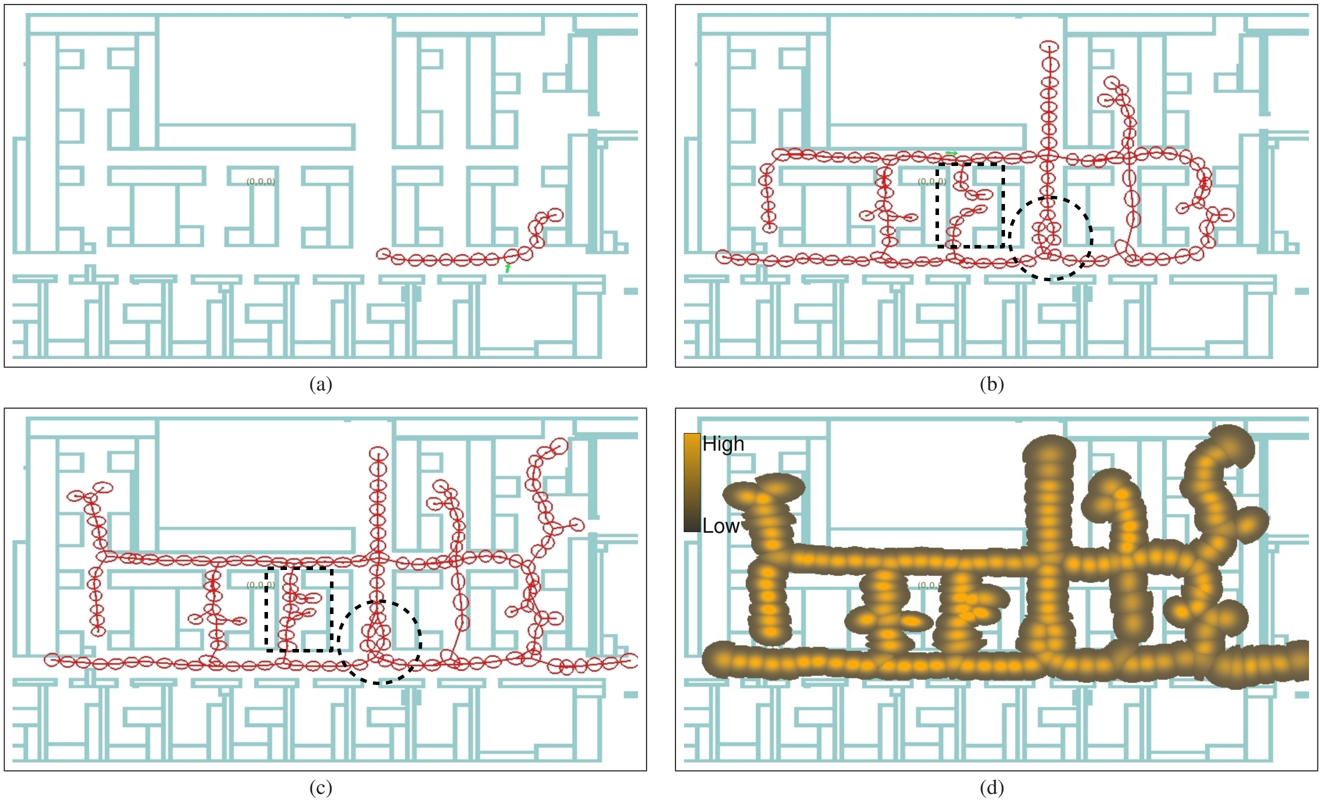

Fig. 8. Model evolution, a) The first observed trajectory in the model. b) The model after the robot observed 7 trajectories. c) The model after observing 25 trajectories. d) The final model after observing more than 60 trajectories. The red ellipses and red lines denote the covariance matrices and state transitions.

in motion patterns rather than detailed trajectories. Further, the SHMM learning approach has the ability to change and adapt the model to accommodate new observations, which is crucial for any mobile robotic deployment.

Finally, the use of such a model for path planning was briefly outlined. Above this, the ability to use the motion pattern model for prediction of future poses of moving people is of great interest in mobile robotics and is part of future publications. Future work will focus on the use of such predictions to improve tracking and interaction with human peers in office spaces.

VI. ACKNOWLEDGMENTS

This work is supported by the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government.

REFERENCES

[1] S. Thrun, M. Bennewitz, W. Burgard, A. Cremers, F. Dellaert, D. Fox, D. Haehnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz, “Minerva: A second generation mobile tour-guide robot,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 1999.

[2] N. Roy, W. Burgard, D. Fox, and S. Thrun, “Coastal navigation – mobile robot navigation with uncertainty in dynamic environments,” in *In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1999, pp. 35–40.

[3] D. Makris and T. Ellis, “Finding paths in video sequence,” in *Proceedings of the British Machine Vision Conference*, 2001, pp. 263–272.

[4] E. Swears, A. Hoogs, and A. G. A. Perera, “Learning motion patterns in surveillance video using hmm clustering,” in *Motion and video Computing, 2008. WMVC 2008. IEEE Workshop on*, 2008, pp. 1–8.

[5] D. Vasquez, C. Fraichard, and C. Laugier, “Incremental learning of statistical motion patterns with growing hidden markov models,” in *Proceedings of The 13th International Symposium of Robotics Research*, November 2007, pp. 348–353.

[6] M. Bennewitz, W. Burgard, and S. Thrun, “Learning motion patterns of persons for mobile service robots,” in *IEEE International Conference on Robotics and Automation (ICRA)*, vol. 4, May 2002, pp. 3601–3606.

[7] 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)*, April 2005, pp. 3948–3953.

[8] S. Sehestedt, S. Kodagoda, A. Alempijevic, and G. Dissanayake, “Efficient learning of motion patterns for robots,” in *The Australasian Conference on Robotics and Automation (ACRA)*, 2009.

[9] L. Rabiner, “A tutorial on hidden markov models and selected applications in speech recognition,” *Proc. IEEE*, pp. 257–286, Feb. 1989.

[10] L. Liao, D. Fox, and H. A. Kautz, “Learning and inferring transportation routines,” in *Proceedings of the Nineteenth National Conference on Artificial Intelligence, Sixteenth Conference on Innovative Applications of Artificial Intelligence*, 2004, pp. 348–353.

[11] B. Fuglede and F. Topsoe, “Jensen-shannon divergence and hilbert space embedding,” in *International Symposium on Information Theory*, June 2004, p. 31 ff.

[12] M. Bolić, P. M. Djurić, and S. Hong, “Resampling algorithms for particle filters: A computational complexity perspective,” *EURASIP Journal on Advances in Signal Processing*, vol. 2004, no. 15, pp. 2267–2277, 2004.

[13] 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.

[14] S. Sehestedt, S. Kodagoda, and G. Dissanayake, “Robot path planning in a social context,” in *IEEE International Conference on Robotics, Automation and Mechatronics (RAM)*, 2010.