

Socially aware path planning for mobile robots

Sarath Kodagoda*, Stephan Sehestedt and Gamini Dissanayake

Centre for Autonomous Systems (CAS), The University of Technology, Sydney, Australia

(Accepted June 2, 2014. First published online: July 1, 2014)

SUMMARY

Human–robot interaction is an emerging area of research where a robot may need to be working in human-populated environments. Human trajectories are generally not random and can belong to gross patterns. Knowledge about these patterns can be learned through observation. In this paper, we address the problem of a robot’s social awareness by learning human motion patterns and integrating them in path planning. The gross motion patterns are learned using a novel Sampled Hidden Markov Model, which allows the integration of partial observations in dynamic model building. This model is used in the modified A^* path planning algorithm to achieve socially aware trajectories. Novelty of the proposed method is that it can be used on a mobile robot for simultaneous online learning and path planning. The experiments carried out in an office environment show that the paths can be planned seamlessly, avoiding personal spaces of occupants.

KEYWORDS: Mobile robots; Navigation; Path planning; Mechatronic systems; Service robots.

1. Introduction

Path planning is a fundamental yet crucial ability that a robot should possess to execute most mobile robotic tasks. In the past, it has been extensively explored with great diversity of solutions.^{1,2} Although there was work undertaken to perform path planning in dynamic environments, non-static objects (such as people) were often treated as purely random occurrences, sometimes using a model to predict those occurrences, in an effort to reduce the problem to a static environment.^{3,4} However, one main focus of current robotic research is the human interaction, where a robot carry out its intended tasks, collaborating with humans. In these types of applications, robots should be equipped with more sophisticated path planning algorithms to deal with all the complex dynamics that are inherent in human-populated environments.

Human-populated spaces usually exhibit common motion patterns, which has been noted in different areas of research.^{5,6} These motion patterns can be a result of a complex combinations of goals, tasks, physiological and social constraints. General observation of people is that they do not always plan the shortest paths. Depending on the task at hand, environmental and social constraints, they may choose to use a longer path. As an example, an office clerk may choose not to walk through cubicals even if it is the shortest path, but opt to take a longer path through a corridor to avoid invading occupants’ workspace. In a social context, these types of behaviors are very constructive. Therefore, if a robot could learn these human behaviors and plan “human-like” paths, it can be considered as a positive step toward a socially acceptable robotics. Being a robot, this knowledge could be exploited by observing and learning dynamics in an environment and planning paths based on the learned models.

Learning could be achieved through the information based on infrastructure-mounted sensors overlooking the region of operation or based on sensors mounted on mobile robots or both. The infrastructure-mounted sensors are mainly utilized in video surveillance applications.^{7,8} However, there is a few literature available on learning models based on infrastructure-based sensors to be used in mobile robots with limited success.

* Corresponding author. E-mail: Sarath.Kodagoda@uts.edu.au

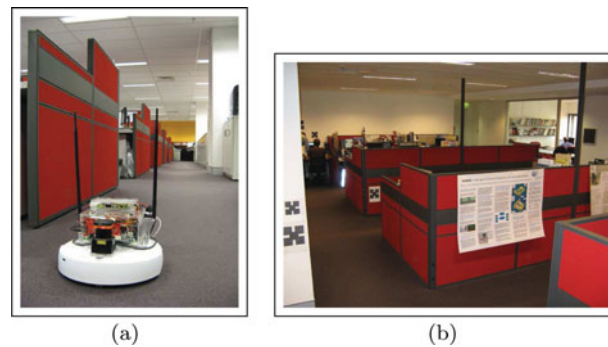


Fig. 1. (a) LISA robot navigating in its environment. (b) The office environment.

Kruse *et al.*⁹ and Govea *et al.*¹⁰ utilized stationary cameras to observe motion patterns in an office space and car park respectively. In Kruse *et al.*,⁹ a statistical representation of motion patterns is proposed. The trajectories are modeled as a Poisson process in consecutive locations with linearization to minimize the complexity. Furthermore, it is proposed to combine similar trajectories, which eventually lead to loss of information. Once the model is learned, it is utilized for pre-planning to minimize the probability of collisions and improved reactive behavior. It results in the robot preferring areas of low traffic density. It has also reported to have improved motion planning with moving obstacles. In Govea *et al.*,¹⁰ a model of trajectories is learned incrementally by dividing the space into Voronoi regions. These regions are used to define the states and state transitions in the proposed Growing Hidden Markov Models to model motion of people or cars in a parking area. Even though the method is promoted as being useful for mobile robotics, no experiments were presented which would support this claim. These methods may not be suitable for most of the mobile robotic applications due to the requirement of infrastructure modifications, higher cost, physical constraints in observations due to large occlusions etc.

Data acquired through infrastructure-based sensor networks can also be found in literature. Bennewitz *et al.*¹¹ use Expectation Maximization (EM) to learn trajectories of individual persons in an office environment. The learning procedure is carried out off-line based on sensor network data, and then fed to a mobile robot to implement reactive behaviors. The algorithm specifically requires complete trajectories between defined resting points, which may not be always available in most human-populated environments. Kanda *et al.*¹² propose to use a sensor network to track people walking in a shopping mall. Local behaviors, such as fast-walk and idle-walk, of people are learned to subsequently form a histogram of local behaviors in each grid cell of a discretized space. Global behaviors of people are then analyzed using state chains of local behaviors. A number of global behaviors are extracted from a large data set, which range from passing through the observed space to window shopping. This model is then used to enable a robot to identify the pattern a person is engaged in and approach to motivate entering the store. Use of sensor networks for data acquisition may require careful positioning of sensors, and it may need infrastructure modifications.

Although these methods show appealing results, the approaches do not exploit the full potential of learned information to be used in Human–Robot Interaction. These methods are in general implemented off-line learning strategies, which have disadvantages of incorporating incomplete data, chronological changes in motion patterns and sudden changes while used in online application. Therefore, in this paper we propose the motion model learning based on Sampled Hidden Markov Model (SHMM) and use it with the A^* and probabilistic road maps (PRM) algorithms to solve the above-mentioned limitations.

This work of socially aware path planning is based on our previous work.¹³ The observer travels through an office-like environment as shown in Fig. 1(a). Although a 30-m laser range finder was used in the experiments, the environmental conditions, such as cubical walls, lead to partial observations (see Fig. 1(b)). Learning is based on sampling from the observation of trajectories, followed by clustering. This model comprises rich knowledge of human motion patterns. It is then exploited in a modified A^* -based algorithm to realize socially aware path planning. The contributions of this paper are: (1) Synthesis of an online motion pattern learning algorithm based on SHMM, which is

capable of utilizing sensory data as and when they arrive requiring no dedicated learning phase, (2) the SHMM-based motion pattern algorithm is used with A^* algorithm and probabilistic roadmaps to achieve socially acceptable navigational paths and (3) experiments on a robotic platform are presented to validate the results.

The remainder of this paper is organized as follows. In Section 2, learning motion patterns based on SHMM are presented. Section 3 discusses the issues of path planning in populated environments. Experiment results based on a robotic platform are presented in Section 4. Section 5 concludes the paper with an indication for future work.

2. Representation and Learning of Motion Patterns

In this section, the proposed methodology for learning motion patterns based on SHMM is presented.

2.1. Learning motion patterns

In a 2-dimensional (2D) environment, a trajectory can be described as a succession of x - y positions with heading θ and linear speed v_l . Consider a segmentation of a 2D space into a regular grid, where the occurrence of motion could be counted in each grid cell. Extending the grid to include all four dimensions, i.e. x , y , θ and v_l , would result in what is known as a flow field or motion histogram.¹⁴ One could normalize values in the histogram to obtain probabilities resulting in a grid representation of joint probability distribution,

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

which expresses the probability of the simultaneous occurrence of x - y - θ and v_l . Knowledge of this distribution constitutes knowing all motion patterns in the environment independent of time. The distribution is very complex and thus requires a significant amount of data to succeed.

Here a probabilistic method is developed to build an estimate of Eq. (1) considering it to be a moving observer with limited field of view. The goal is to incrementally build the belief of D_t (approximation of Eq. (1)) using all sensor readings $z_{0..t}$, all robot poses $\zeta_{0..t}$ and all observations of moving people $\xi_{0..t}$ up until time t ,

$$Bel(D_t) = P(D_t | \xi_t, \zeta_t, z_t, \dots, \xi_0, \zeta_0, z_0). \tag{2}$$

From the above equation, an incremental update rule can be derived using the well-known Bayes theorem as

$$Bel(D_t) = \eta P(\xi_t | D_t, \zeta_t, z_t, \xi_{t-1}, \dots, \xi_0, \zeta_0, z_0) P(D_t | \zeta_t, z_t, \xi_{t-1}, \dots, \xi_0, \zeta_0, z_0), \tag{3}$$

where, $\eta = P(\xi_t | \zeta_t, z_t, \xi_{t-1}, \dots, \xi_0, \zeta_0, z_0)$ is a constant. Since $Bel(D_t)$ is the belief of D at time t , given all past observations, sensor readings and observer poses, it is not an efficient solution without further simplifications. 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,

$$Bel(D_t) = \eta P(\xi_t | D_t, \zeta_t, z_t) \overbrace{P(D_t | \zeta_t, z_t, \xi_{t-1}, \dots, \xi_0, \zeta_0, z_0)}^{\text{prior belief}}. \tag{4}$$

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

$$Bel(D_t) = \eta P(\xi_t | D_t, \zeta_t, z_t) Bel(D_{t-1}). \tag{5}$$

This result allows the update $Bel(D_t)$ to use only the most recent observations of moving people. Due to the intricacies of human spatial behaviors, $Bel(D_t)$ is complex, and is of unknown distribution. Therefore, an adequate representation has to be chosen.

2.2. Hidden Markov model

A Hidden Markov Model (HMM) is a statistical model that represents a system as a directed graphical model. Here we briefly outline HMMs following the notation used by Rabiner.¹⁵ HMMs are defined by N states of a system $S = s^1, s^2, \dots, s^N$, together with the observation symbols $V = v^1, v^2, \dots, v^M$, with M being the number of symbols.

A state transition probability distribution $A = a_{ij}$ is given as

$$a_{ij} = P(q_{t+1} = s^{(j)} | q_t = s^{(i)}), \quad 1 \leq i \leq N \\ 1 \leq j \leq N. \quad (6)$$

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

$$b_{ij} = P(v^{(i)} | s^{(j)}), \quad 1 \leq i \leq M \\ 1 \leq j \leq N. \quad (7)$$

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

$$\pi_i = P(q_1 = s^{(i)}), \quad 1 \leq i \leq N. \quad (8)$$

A large number of variations of HMMs are proposed in literature, and generally HMMs are reported to be working well in different areas of research.^{10,11}

2.3. Sampled hidden Markov model

While a grid-based approach is a possible representation of $Bel(D_t)$, there are a number of shortcomings to this approach. First, it leads to high computational costs and inefficient implementations due to the need of updating the whole grid whenever the belief is updated, irrespective of the field of view of the observation. Moreover, as the grid is ignorant of the environment's structure, a grid representation of motion patterns would require maintaining a belief in regions where even no motion would be possible (e.g. inside walls). Finally, the grid's resolution has to be chosen carefully, but even then it is difficult to guarantee a good resulting approximation.

Considering the inherent shortcomings of grid-based approaches, a sampling algorithm which can predict, weigh and resample to incrementally learn an approximation of Eq. (1) is proposed. The sample-based representation overcomes the grid-based approaches' problems by only generating samples in the areas of interest, i.e. where motion was observed and the number of samples can be controlled by means of resampling or subsampling.

The belief $Bel(D_t)$, as defined in Eq. (5), can be represented as a set of weighted samples,

$$X_t = \langle x_t^{(i)}, \omega_t^{(i)} \rangle, \quad 1 \leq i \leq N, \quad (9)$$

where

$$x_t^{(i)} = [x^{(i)} y^{(i)} \theta^{(i)} v_t^{(i)}]^T \quad (10)$$

and $\omega_t^{(i)}$ being the weight of the i th sample. The belief of D_t is then defined as

$$Bel(D_t) = X_t = \eta \overbrace{P(\xi_t | D_t, \zeta_t, z_t)}^{\text{People Tracking}} Bel(D_{t-1}), \quad (11)$$

where a particle filter is implemented for people tracking.

Consider a person walking along a corridor in the direction of the arrow as shown in Fig. 2(a). The particle filter-based tracker would produce a series of poses along the observed trajectory. Each of these poses would be represented by a cluster of samples (see Fig. 2(b)) which approximates a

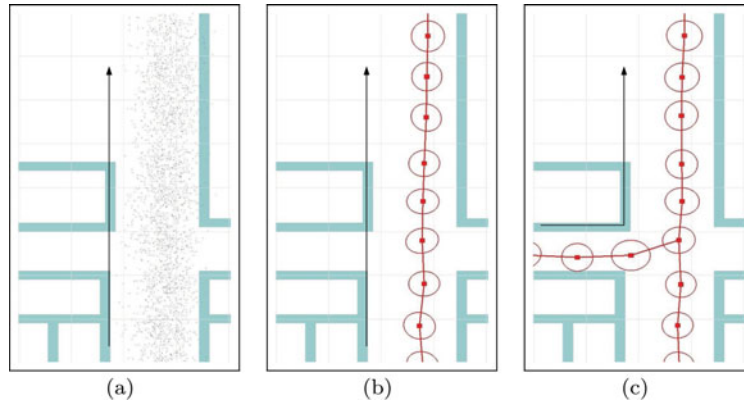


Fig. 2. Learning example: (a) A person was moving in the direction of the arrow with samples from the tracker. (b) The resulting SHMM. (c) After merging a new trajectory from the left. In (b) and (c), the states are represented by means, “.”, covariances by ellipses and transition probabilities by the thickness of the straight lines joining the states.

probability distribution over x, y, θ and v_j . The transitions between states are directly observable,

$$S = s^{(i)} = \begin{bmatrix} \mu^{(i)} \\ \Sigma^{(i)} \end{bmatrix} \quad 1 \leq i \leq N, \tag{12}$$

where $\mu^{(i)}$ and $\Sigma^{(i)}$ are the mean and the covariance of the i th sample cluster and N is the number of states. Whenever another moving object is observed in a region where a model was previously learned, the statistics of states are updated by combining the corresponding sample clusters. Therefore, the time dependency needs to be incorporated into the definition of a state; however, it is omitted in the formulation due to unworthy complexity.

2.3.1. Model adaptation and growth. By observing a moving object, the resulting cluster can be seen as the j th state $s^{(j)-}$ in the path of the object. The superscript “-” means that it is either a new state or may add new information to an already existing state in the model of motion patterns. The decision can be made based on the symmetric Kullback–Leibler divergence (KLD)¹⁶ representing $s^{(i)}$ and $s^{(j)-}$ as probability distributions,

$$\text{KLD}(s^{(i)} \parallel s^{(j)-}) = \text{KLD}_s(s^{(i)} \parallel s^{(j)-}) + \text{KLD}_{s-}(s^{(j)-} \parallel s^{(i)}), \quad \begin{matrix} 1 \leq i \leq N \\ 1 \leq j \leq K, \end{matrix} \tag{13}$$

where K is the number of clusters from the tracked object’s path. The first and the second terms of the equation are computed as

$$\text{KLD}_s(s^{(i)} \parallel s^{(j)-}) = \int_{-\infty}^{\infty} P_{s^{(i)}}(x) \log \frac{P_{s^{(i)}}(x)}{P_{s^{(j)-}}(x)} dx, \quad \begin{matrix} 1 \leq i \leq N \\ 1 \leq j \leq K \end{matrix} \tag{14}$$

and

$$\text{KLD}_{s-}(s^{(j)-} \parallel s^{(i)}) = \int_{-\infty}^{\infty} P_{s^{(j)-}}(x) \log \frac{P_{s^{(j)-}}(x)}{P_{s^{(i)}}(x)} dx, \quad \begin{matrix} 1 \leq i \leq N \\ 1 \leq j \leq K. \end{matrix} \tag{15}$$

If the KLD between state $s^{(i)}$ and cluster $s^{(j)-}$ is below a certain threshold, the sample clusters can be combined and the state statistics can be updated accordingly. The threshold can be computed based on the expected distance between consecutive states, or be chosen as a fixed threshold based

Table I. Personal space zones as defined in ref. [18].

Personal space zone	Distance (m)
1. Intimate	0.0–0.45
2. Personal	0.45–1.2
3. Social – consultative	1.2–3.0
4. Public	3.0+

on experimentation. To avoid growing computational effort with a growing model, the KLD is only computed for clusters which are closely located in the state space. It is to be noted that the robot localization is assumed to be known. Figure 2(c) shows a situation where a new trajectory from the left is joined based on the KLD criteria.

2.3.2. Reducing the dimensionality of the model. Higher dimensions demand higher computational requirements. The dimensionality is proposed to be reduced by making use of the properties of the HMM. When sampling with a fixed frequency is proposed, the distance between the means of successive states encodes the average speed of the target in the area. Thus, a distance measure for speeds or binning replaces the explicit use of v in the model, and hence the state can be represented as (x, y, θ) . Similarly, it is possible to use binning for heading θ , reducing the dimensions of a single state to just two, (x, y) .

3. Path Planning in Populated Environments

Path planning in dynamic environments has been studied before;^{3,17} however, most of the time the context of a human populated space has not been exploited. In psychology and other fields, use of a space by multiple people is an active field of research, which yields many interesting results.^{6,18} It is a fact that the interaction between humans and robots in this context is not extensively studied yet because of lack of existing real-world scenarios. However, with the fast growing field of human–robot interaction, it is perceived that such robots operate in a socially acceptable manner for seamless integration with humans.

The concept of personal space as described by Hall¹⁸ is central to human space sharing. There it is stated that the space needed by a person is more than the body's volume, but there are areas around a person which should only be intruded for a particular interaction (see Table I). The innermost area should only be entered if an intimate relationship exists between the involved persons, whereas middle region is only for personal relationships. The outer circle defines an area which is used for consultive interaction, such as between colleagues. Unexpected intrusion of the personal space may result in discomfort, lower perceived privacy and lower work performance,¹⁹ which is enviable to avoid if possible. From the perspective of robotics, it is desirable to develop path planning algorithms giving due regards to such social issues.

3.1. The A* algorithm

The A* algorithm and its derivatives are a popular solution for the path planning problem.^{2,3,20} It performs the best-first search on a grid which is pre-computed using a collision detector with defined configuration values. More precisely, a configuration space C is computed that contains all static obstacles in the d -dimensional space of the robot. There a C_{free} exists, which contains all collision-free configurations. A path planning algorithm searches a path such that the path lies in C_{free} .

A* is defined by the functions $g(x)$, which is the shortest path from start to goal by Euclidean distance (often called the path-cost function), $h(x)$, which is used as a heuristic estimate of the length of the path, and $f(x)$, which defines the sum of $g(x)$ and $h(x)$. The algorithm searches for a path using a priority queue, where the priority of node x is higher, its $f(x)$ being lower. Hence, it is called the best-first search.

This can be exploited to include prior information about dynamics in the environment. A cost function is used to evaluate the cost of a path with respect to a model of motion patterns as

$$g_D(x) \sim D(x), \quad (16)$$

where D denotes the learned model of motion patterns; $g_D(x)$ returns a low value if node x is in an area of high traffic density, and returns a high value if it is in an area of low traffic density. Consequently, $g(x)$ in the standard A^* algorithm can be replaced by the below given function to calculate $f(x)$,

$$G(x) = g(x) + g_D(x). \quad (17)$$

This cost function is suitable for a robot, which is supposed to prefer a commonly taken path by people. However, in complex scenarios, it is appropriate to have more flexibility in choosing path planning solutions. For example, in some cases it may be better for the robot to choose the shortest path. This flexibility could be introduced by reformulating Eq. (17) as

$$G(x) = g(x) + w * g_D(x), \quad 0 \leq w \leq 1, \quad (18)$$

where the factor w is chosen depending on the current requirements. Setting w to zero means the paths are planned based purely on A^* . If w is set to 1, the robot prefers common paths, whereas any number between 0 and 1 denotes the combination of pure A^* with $g_D(x)$.

3.2. Probabilistic roadmaps

The integration of PRM path planning with the proposed model of motion patterns is straightforward and yields some interesting properties.

Probabilistic roadmap was introduced as a method to overcome the issue of growing complexity in higher dimensions.^{2,21} The basic algorithm first constructs an undirected graph G , the roadmap, to solve the path planning problem. The nodes of G are generated by random sampling and collision checking. Path planning is done by traversing between nodes which are sufficiently close to each other. There are many publications presenting variations to the sampling step and collision checking in order to improve the efficiency. Generally, PRM has been applied in many successful applications.

The main appeal to use PRM here is that once a graph is constructed, path planning can be done with A^* as detailed above. Consequently, when using the model of motion patterns D , at least part of the graph G can be considered known and sampling can be restricted to unexplored areas, thus improving the efficiency. The more complete D is, the less sampling needs to be done. This highlights the appeal of PRM when using with SHMM. If there is an already learned SHMM, the PRM graph can be considered known or mostly known and planning can be directly done on the SHMM graph.

4. Experimental Results

4.1. Experimental setup

Experiments were carried using an in-house developed Lightweight Integrated Social Autobot (LISA), which was realized using an iRobot Create platform (see Fig. 1(a)). The robot carried an Intel D510MO small-scale computer and a Hokuyo UTM-30LX laser range finder enabling it to localize and navigate in the environment while detecting²² and tracking people using a simple particle filter. The software development environment was Player/Stage²³ and all the algorithms were implemented in C++ within the Orca software framework.²⁴ Figure 3 shows a Simultaneous Localization and Mapping (SLAM)-generated map of the environment where desk areas and corridors are marked appropriately. The LISA robot is shown as a red circle, and the red outline illustrates the observed laser reading. Being a small robot, it has a significantly limited field of view due to the presence of furniture. The map spans approximately 32 m × 20 m. It is important to note the complexity of the environment with a large amount of clutter, semi-static objects such as trash cans and chairs and transparent objects such as glass walls.



Fig. 3. The map used in experiments. Desk areas, hallways and common areas are marked as “D”, “H” and “C” respectively. LISA’s pose is shown by a red circle and the observed laser reading is shown as a red outline. Note the limited observability LISA in the environment.

4.2. Learning motion patterns

In this section, SHMM learning is presented with the robot LISA in the aforementioned office environment. Ten different subjects were included in this experiment and no environment modifications were done. The limited observability at most times means that the robot has to explore the environment to build a model of motion patterns. Furthermore, in order to observe longer trajectories, the robot has to follow people, hence it has to be a mobile observer. For more information about the implementation issues such as updating the transition probability matrix, the readers are referred to Sehestedt *et al.*²⁵

A series of plots in Fig. 4 show the evolution of an SHMM. Figure 4(a) shows the robot following a person, where the person is represented by a yellow cylinder and the trajectory is shown as an orange line. The robot is shown as a red circle, whereas the red outline indicates the observed reading of the forward looking laser sensor. The observed trajectory exhibits a typical human motion and accordingly it is represented in the initial SHMM as shown in Fig. 4(b).

Figure 4(c) shows the model after more than 70 observed trajectories while the robot was on the move. The trajectories were successfully joined and compactly represented. The final representation, including more than 80 trajectories, is shown in Fig. 4(d) as a unimodal Gaussian distribution. It could be noted that trajectories are positioned correctly on free spaces rather than through obstacles. Further, compared with grid-based representations of motion patterns, a greater efficiency is achieved as the belief has to be maintained only in the relevant areas (with human motion) of interest rather than over the entire space.

4.3. Model adaptation

Another important aspect of the evolution of an SHMM is adaptation to changes in the environment. Consider the situation in Fig. 5(a), where people usually walk along a hallway in an almost straight line. The learned model after observing five similar trajectories is illustrated in Fig. 5(b). Then an obstacle is placed on a commonly taken path to partially block it, which leads to a change in person’s path as visible in Fig. 5(c). After five observations of changed trajectories, the SHMM’s states shifted accordingly as in Fig. 5(d), thus effectively adapting to the change in behavior.

For the next experiment, consider the same situation as in Fig. 5(b), however, with a larger obstacle blocking the common path as shown in Fig. 6(a). Naturally, person’s trajectory has to change drastically to avoid an obstacle. As a result, it is not sufficient to merely shift the locations of the existing states to accommodate the change, instead the model is extended with a new part as shown in Fig. 6(b). Initially, the transition \overline{ac} has a lower probability than the transition \overline{ab} as indicated by the thickness of the red state transition lines. However, with more observations the transition probability of \overline{ac} became larger than that of \overline{ab} as can be seen in Fig. 6(d).

4.4. Socially aware path planning

In the previous section, the robot has learned human motion patterns through SHMM, which captures the navigational behaviors of people in a particular environment. Having such a knowledge, the robot can then behave in a similar manner obeying general human social rules. For example, the robot is supposed to avoid office desk spaces of other people, i.e. a high value for w_g in Eq. (18) in path planning. Importance of avoiding the desk areas can be seen in Fig. 7, where the personal space area

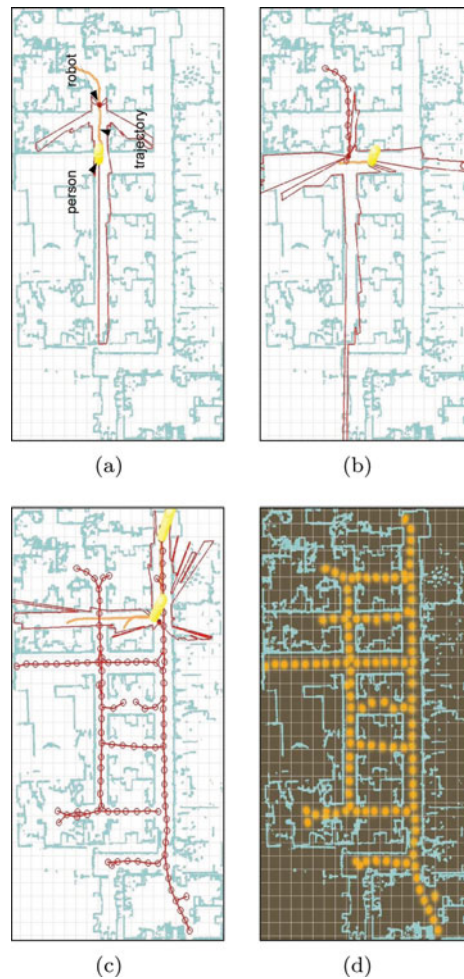


Fig. 4. Model learning with real-world data. (a) Robot follows a person; (b) initializes the SHMM with the observed trajectory in (a); (c) SHMM after observing more than 70 trajectories; (d) learned model is represented as a unimodal Gaussian distribution.

is overlaid on a part of the map. This shows that if a robot moves through one of the desk areas, it may intrude the personal space of the occupier, which is not desirable.

Figure 8 shows the A^* -cost map derived from the SHMM shown in Fig. 4(d). This cost map illustrates the cost for traversal, which is calculated using the spatial distance of a grid cell to the SHMM as well as the observed traffic density in the area. The spatial distance as a factor is useful to incorporate closeness to previously seen trajectories to path planning without a need to stay exactly on such a trajectory. Combination of two factors represents rich information about human path planning and is used for socially aware path planning in the following experiments.

Figure 9 presents some path planning results. Figure 9(a) shows a more traditional planning result based on the shortest path criterion, where among a set of equally long paths the planner has no preference. Given equal probability for choosing any of these paths, in this example there is a 2/3 probability of the robot passing a desk area. Utilizing the proposed method, the planned path became the one shown in Fig. 9(b). Hence, the robot was able to deliberately avoid desk areas, which could have been someone's personal space.

As illustrated in Fig. 9(c), the robot would even accept a longer path to avoid disturbing office workers at their desks. However, if the detour for a socially aware path compared with the shortest path is too long, the robot may choose the shortest path as shown in Fig. 9(d). This stems from the fact that the accumulated cost, including the higher cost for traversal in some areas as illustrated in Fig. 8, is lowest for the shortest path. This could be argued as mimicking human behavior for navigation.

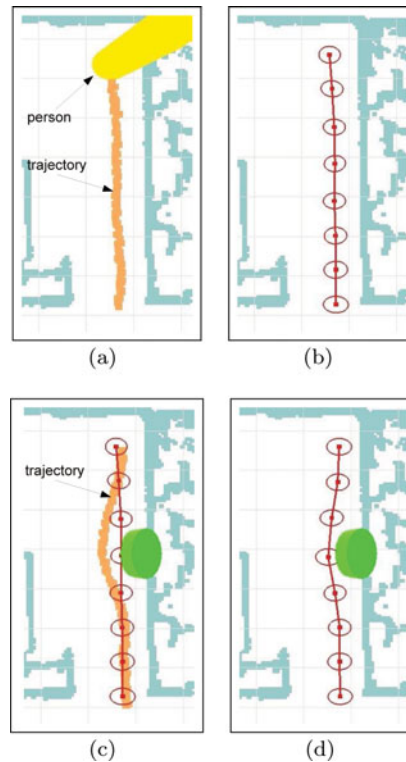


Fig. 5. An SHMM is adapting to a slightly changed environment. (a) The initial common trajectory. (b) The SHMM representation of a common trajectory. (c) An obstacle causes a slight change in trajectory. (d) SHMM adapting to a new situation.

These experiments prove that the experimental data include valuable information about human path planning which can be successfully exploited for socially aware path planning.

4.5. Simultaneous path planning and model learning

The last experiment is designed to demonstrate the seamless integration of path planning and model learning. In Fig. 10(a) LISA plans a path around a desk area based on the information at hand. Once it reaches near the goal (Fig. 10(b)), it detects a walking person and starts tracking while moving toward the goal (Fig. 10(c)). Once the person leaves robot's field of view (i.e. the track is terminated), new information is added to the model of motion patterns as shown in Fig. 10(d).

5. Conclusions and the Future work

In this paper, a method was presented to allow a robot to plan paths considering social behaviors. This philosophy is important as it will lead robots to behave more like humans, reinforcing human–robot interactions. For this purpose, we proposed a motion pattern learning algorithm to enhance the robot path planning algorithm. The learning algorithm was performed online with data received from on-board sensors without involving any infrastructure-mounted sensors. Further, there is no dedicated learning phase utilized where all acquired information is immediately incorporated in the model. Therefore, the robot can be deployed efficiently in many environments. The model is capable of adapting to new information and thus it is suitable for lifelong learning in a changing world. The representation of motion as an SHMM is more memory-efficient than grid-based approaches, as the model only represents motion in the areas of interest. As the model is not fully connected, the transition matrix can be replaced by more compact data structures.

Path planning was achieved using A^* algorithm and PRMs, which integrate coherently with the motion pattern models. A weight is set to control influence on path planning, leading to socially acceptable paths. Experimental results show that the robot can plan paths with regard to a model of motion patterns, avoiding certain areas where people prefer minimal interactions.

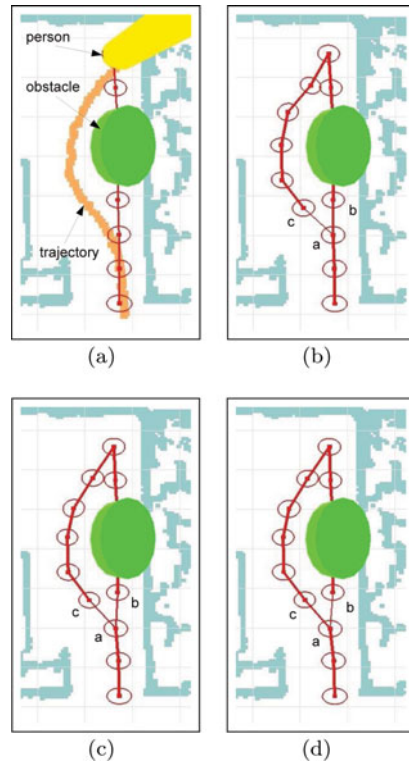


Fig. 6. An SHMM adapting to a large change. The initial trajectory and SHMM are the same as in Fig. 5. (a) A large obstacle causes a drastic change in trajectory. (b) Large change in trajectory leads to the learning of a new branch of SHMM. Initially, the transition \overline{ac} is estimated to be less likely than \overline{ab} as indicated by the thickness of the transition lines. (c) and (d) Observing the changed behavior repeatedly leads to an increase of \overline{ab} transition probability than that of \overline{ac} .




Fig. 7. Personal spaces overlaid on a part of the map.




Fig. 8. The A^* -cost map for socially aware path planning, illustrating the cost for traversal in a 2D grid. This cost map is also indicative for the cost of the PRM graph. Note that the cost map needs to be updated whenever the SHMM is updated.

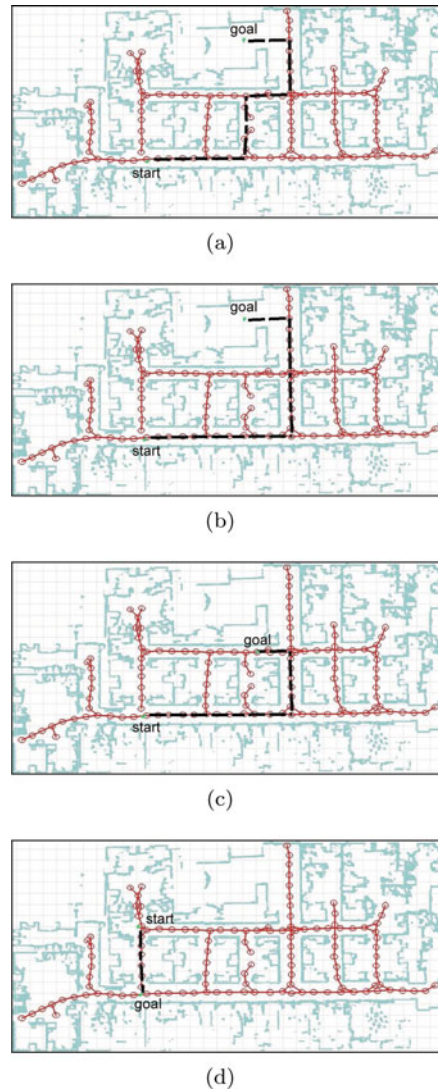


Fig. 9. Socially aware path planning based on real data.

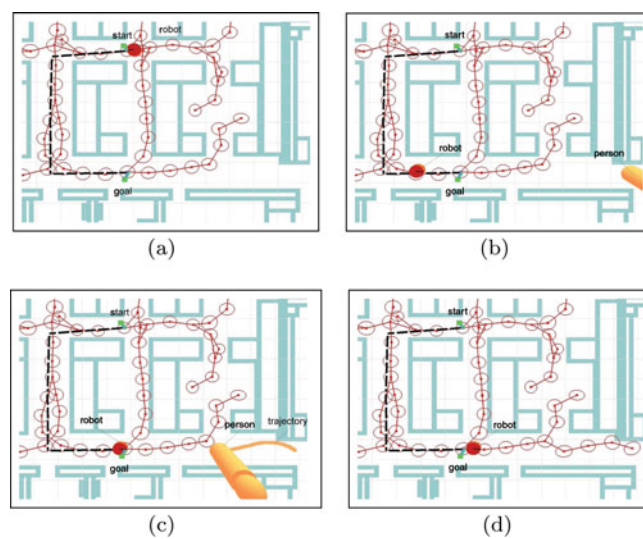


Fig. 10. (a) A path generated using basic PRM. (b) The robot reaches the defined goal. (c) The robot detects a walking person and starts tracking. (d) A newly observed track has been added to the model.

In this study, it was assumed that there were no reactive actions taken by the people to regain their personal spaces. This could be an interesting area for future research. Further, the future research could also focus on the management of the states in SHMM, integration of the SHMM-based motion prediction with other probabilistic data association methods, e.g. sample-based Joint Probabilistic Data Association²⁶ and learning of common locations of interest, e.g. kitchen, printer, which could further improve the long-term prediction capabilities.

Acknowledgments

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

References

1. J. Latombe, *Robot Motion Planning*, Eighth printing (Kluwer, Boston, MA, 1991).
2. S. M. LaValle and M. S. Branicky, "On the relationship between classical grid search and probabilistic roadmaps," *Int. J. Robot. Res.* **23**(7–8), 673–692 (2004).
3. M. Likhachev, D. Ferguson, G. Gordon, A. T. Stentz and S. Thrun, "Anytime Dynamic a*: An Anytime, Replanning Algorithm," *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, Monterey, California, USA (Jun. 5–10, 2005) pp. 262–271.
4. J. V. D. Berg, D. Ferguson and J. Kuffner, "Anytime Path Planning and Replanning in Dynamic Environments," *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Orlando, Florida, USA (May 15–19, 2006) pp. 2366–2371.
5. D. Dean, *Museum Exhibition: Theory and Practice* (Routledge, Oxford, UK, 1994).
6. I. Altman, A. Rapoport and J. Wohlwill, *Environment and Culture* (Springer, New York, NY, 1980).
7. D. Makris and T. Ellis, "Finding Paths in Video Sequence," *Proceedings of the British Machine Vision Conference*, Manchester, UK (Sep. 10–13, 2001) pp. 263–272.
8. L. Liao, D. J. Patterson, D. Fox and H. A. Kautz, "Learning and inferring transportation routines," *Artif. Intell.* **171**(5–6), 311–331 (2007).
9. E. Kruse, R. Gutsche and F. Wahl, "Acquisition of Statistical Motion Patterns in Dynamic Environments and Their Application to Mobile Robot Motion Planning," *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '97)*, vol. 2, Grenoble, France (Sep. 11, 1997) pp. 712–717.
10. D. A. V. Govea, T. Fraichard and C. Laugier, "Incremental Learning of Statistical Motion Patterns with Growing Hidden Markov Models," *Proceedings of the 13th International Symposium of Robotics Research*, Hiroshima, Japan (Nov. 26–29, 2007) pp. 75–86.
11. M. Bennewitz, W. Burgard and S. Thrun, "Adapting Navigation Strategies Using Motions Patterns of People," *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA '03)*, vol. 2, Taipei, Taiwan (Sep. 14–19, 2003), pp. 2000–2005.
12. T. Kanda, D. Glas, M. Shiomi and N. Hagita, "Abstracting people's trajectories for social robots to proactively approach customers," *IEEE Trans. Robot.* **25**(6), 1382–1396 (Dec. 2009).
13. S. Sehestedt, S. Kodagoda and G. Dissanayake, "Models of Motion Patterns for Mobile Robotic Systems," *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Taipei, Taiwan (Oct. 18–22, 2010) pp. 4127–4132.
14. A. Lookingbill, D. Lieb, D. Stavens, S. Thrun, "Learning Activity-Based Ground Models from a Moving Helicopter Platform," In: *Proceedings of the 2005 IEEE International Conference on Robotics and Automation (ICRA 2005)*, Barcelona, Spain (Apr. 18–22, 2005) pp. 3948–3953.
15. L. Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," *Proc. IEEE* **77**(2), 257–286 (Feb. 1989).
16. S. Kullback and R. Leibler, "On information and sufficiency," *Ann. Math. Stat.* **22**(1), 79–86 (1951).
17. A. Bruce and G. Gordon, "Better Motion Prediction for People-Tracking," *Proceedings of the International Conference on Robotics and Automation (ICRA)*, Barcelona, Spain (Apr. 18–22, 2005) pp. 1418–1423.
18. E. T. Hall, *The Hidden Dimension – Man's Use of Space in Public and Private* (Bodley Head, London, 1969).
19. C. B. Danielsson, "Differences in perception of noise and privacy in different office types," *J. Acoust. Soc. Am.* **123**(5), 2970 (2008).
20. T. Allen, A. Hill, J. Underwood and S. Scheduling, "Dynamic Path Planning with Multi-Agent Data Fusion – The Parallel Hierarchical Replanner," *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA '09)*, Karlsruhe, Germany (May 6–10, 2013) pp. 3245–3250.
21. R. Bohlin and L. Kavraki, "Path Planning Using Lazy Prm," *Proceedings of the ICRA '00. IEEE International Conference on Robotics and Automation*, vol. 1, San Francisco, CA, USA (Apr. 24–28, 2000) pp. 521–528.
22. Z. Zainudin, S. Kodagoda and G. Dissanayake, "Torso Detection and Tracking Using a 2D Laser Range Finder," *The Australasian Conference on Robotics and Automation (ACRA 2010)*, Brisbane, Australia (Dec. 1–3, 2010).

23. P. Gerkey, R. T. Vaughan and A. Howard, "The player project," available at: <http://playerstage.sourceforge.net/>, August 2010 (online).
24. D. Brugali, A. Brooks, A. Cowley, C. Côté, A. Domínguez-Brito, D. Létourneau, F. Michaud and C. Schlegel, "Trends in Component-Based Robotics," **In: *Software Engineering for Experimental Robotics***, Springer Tracts in Advanced Robotics Series, vol. 30 (D. Brugali, ed.) (Springer, Berlin, Germany, 2007) pp. 231–251.
25. S. Sehestedt, S. Kodagoda, A. Alempijevic and G. Dissanayake, "Efficient Learning of Motion Patterns for Robots," *The Australasian Conference on Robotics and Automation (ACRA 2009)*, Sydney, Australia (Dec. 2–4, 2009).
26. D. Schulz, W. Burgard, D. Fox and A. Cremers, "People tracking with mobile robots using sample-based joint probabilistic data association filters," *Int. J. Robot. Res.* **22**(2) (2003).