

© 2010 IEEE. Reprinted, with permission, from Sarath Kodagoda , Models of motion patterns for mobile robotic systems, Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on, and Oct. 2010. This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Technology, Sydney's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it

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 behaviours 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. This 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 learnt using a sampling routine which outcome is used to construct an incrementally growing Sampled Hidden Markov Model. This model has a number of interesting properties and can be of use in many applications. For example, the learnt knowledge can be used to predict motion and infer social rules and 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 as unpredictability of human motion may require sudden changes to planned tasks. Thus far, a common approach is to employ a method to minimise 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 allow segmentation of any sensor observations so that sensor data that is detrimental to the operation of tasks such as localisation 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 behaviour by observing a scene. Here, one of the fundamental underlying assumptions is that the observer is stationary. The problem’s complexity is further reduced by observing complete trajectories from start to end. 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 these assumptions usually do not hold thus requiring different strategies for dealing with this issue. The greatest 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 handled.

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. This work uses stationary sensors and it is necessary that complete trajectories are observed. Furthermore, the learning algorithm operates off-line and it assumes that motion always happens between some specified points (e.g. an object travels from a start point A to an end point B). Vasquez et al. [5] propose Growing Hidden Markov Models (GHMM) to incrementally learn motion patterns in an area. This allows for on-line learning, however this work also requires that the observer is stationary, which limits its applicability in mobile robotics.

This paper presents and discusses a novel approach to learning typical motion patterns in an environment of a robot based on Sampled Hidden Markov Models (SHMM). 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.

The remainder of this publication is organised as follows. Section II briefly outlines a sampling procedure to learn a probability distribution of motion patterns. In section III we propose SHMMs represent common motion patterns in an environment which can be learnt on-line, without supervision, on a mobile robot. Furthermore, 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 $x - y - \theta$ location and velocity v . Discretising the state space into a spatial grid followed by building a motion histogram [7] and then normalising the values of the grid cells would result in an approximation of the joint probability distribution

$$P(x, y, \theta, v) \quad (1)$$

which represents the probability of the simultaneous occurrence of $x - y - \theta$ and v . Knowledge of this distribution constitutes knowing all motion patterns in the environment independent of time. The distribution is very complex and thus require a significant amount of data to succeed. Therefore, In [8] we proposed a sampling algorithm to incrementally learn an approximation of Eq. 1. Here we extend

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

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 localisation and object tracking observes a person's trajectory. Tracking algorithms commonly represent each piece of a trajectory as probability distribution from which it is possible to take samples. In Fig. 1(a) a person (green rectangle) walks from the left to the right while being tracked. The samples are taken from the prediction of the tracking algorithm and are weighed according to the observation. In the figure a 2D projection of the samples is shown along with the 95% confidence ellipses in x and y (green ellipses). Fig. 1(b) shows the sample set after more people moved along a similar trajectory.




Fig. 1. 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 learning Sampled Hidden Markov Models (SHMM) using the sampling algorithm outlined above. The main consideration is to derive a model which can be learnt and utilised by a mobile robot to improve its operation in a populated space. Even though our sample based representation is already more efficient than a 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 $S = s^1, s^2, \dots, s^N$, observation symbols $V = v^1, v^2, \dots, v^M$ with M being the number of symbols and state transition probability distribution $A = a_{ij}$, which is given as

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

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

$$b_{ij} = P(v^{(i)} | s^{(j)}), 1 \leq i \leq M, 1 \leq j \leq N \quad (3)$$

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

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

Most HMM frameworks highly depend on prior knowledge of the topology of the model and learning is performed through previously obtained data sets. There is no easy way to update these models over time [9]. Thus, these implementations are not suitable for the application at hand. There are numerous, usually application dependent, derivatives of HMMs reported in the literature and 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, comparatively 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 platform.

B. Deriving a Hidden Markov Model

From the sampling algorithm in Section II a particle cloud is obtained (as shown in Fig. 2(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 person's 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 in Section II a vector of M clusters of weighted samples is obtained which describes an observed trajectory

$$C = [c^{(0)} \ c^{(1)} \ \dots \ c^{(M)}] \quad (5)$$

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

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

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. 2(a) a 2D projection of SHMM states can be seen as the red covariance ellipses in x and y . This figure also shows the learnt 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

$$a_{ij} = \begin{cases} \frac{N^{(ij)}}{P(s^{(j)}|s^{(i)})} & 1 \leq i \leq N \\ & 1 \leq j \leq N \end{cases} \quad (7)$$

where N^{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 learnt transitions in this example are 1.

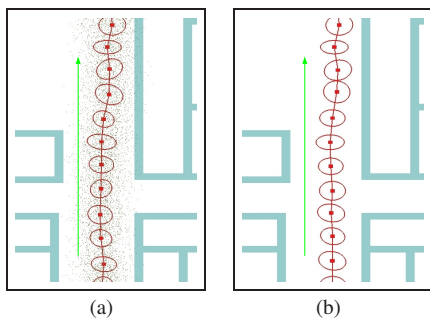


Fig. 2. a/b) A motion pattern is learnt. People are walking from the bottom to the top of the image along the green arrow.

2) *Updating The Model:* When observing another trajectory, a series of sample clusters is produced and data association has to be done to determine whether the trajectory or part thereof is already contained in the model. To do this data association, the symmetrised Kullback-Leibler distance (KLD) [11] is used. It calculates the distance of clusters in C to states in the model. The KL-distance is commonly used in the literature for this purpose, nevertheless it is to be noted that other distance measures such as the Mahalanobis distance could be used as well. The symmetrised KL-distance is defined as follows

$$KLD_{sym}(s^{(i)}|c^{(j)}) = KLD(s^{(i)}|c^{(j)}) + KLD(c^{(j)}|s^{(i)}), \quad 1 \leq i \leq N \quad (8) \\ 1 \leq j \leq M$$

where $KLD_{sym}(s^{(i)}|c^{(j)})$ denotes the symmetrised 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 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 as in a normal particle filter with systematic resampling [12]. Finally, the transition probabilities are updated as

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

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.

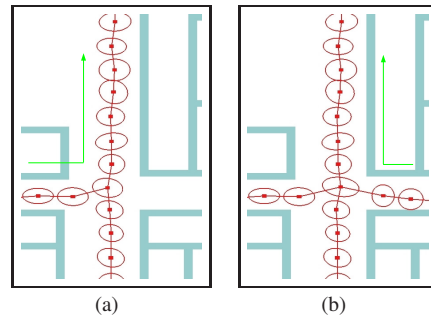


Fig. 3. a) A person is tracked coming from the left and then walking along the previously learnt path. b) Same situation with a person coming from the right. The robot (not visible) changed its location during observations. (real data)

To update the state transitions the knowledge about the sequence of newly added and associated states can be exploited. Where a transition is already known the values can be updated by incrementing $N^{(ij)}$ and updating the transition probabilities accordingly.

Fig. 2(b) shows the updated model after a second person was observed moving along a similar trajectory as the first person (as indicated by the green arrow). Towards the top of the figure the robot lost track of the person and it can be seen that the two top most states were not updated. Since the two trajectories were very similar no new states and transitions had to be added.

In contrast Fig. 3(a) shows the SHMM after another person was observed coming from the left, again following the trajectory indicated by the green 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 learnt. The situation is similar in Fig. 3(b) where a person was tracked coming from the right and the model is updated accordingly.

C. Using the SHMM On a Mobile Robot

The above definition of the SHMM to model common motion in an environment did not take all aspects into consideration which may be of importance for a mobile robot. Most importantly computational complexity and limited observability have to be considered.

1) *Complexity Considerations:* Firstly, by applying the above without further consideration a very detailed model would be learnt. I.e. with a rate of observations of 10Hz, there would be a state every 10cm for objects moving at 1m/s. Our interest is a model of commonly occurring motion patterns rather than detailed trajectories. Therefore, we discard some of the observations to obtain a distance of approximately 0.7m between states for people moving

at a normal walking pace. Still, for learning we take all observations into account.

Handling 4-dimensional Gaussian distributions can be a demanding task, especially under the computational constraints of a mobile robot. Therefore we propose to exploit the structure of the SHMM to reduce the dimensionality to just 2.

Firstly, there are two possible ways to exclude v from the state estimate. It is possible to either sample with a fixed frequency or do binning. When using a fixed frequency with which clusters of samples are produced, the distance between clusters refers to the speed of the observed object thus implicitly representing v . For binning, speed domains may be chosen and for each domain a distance between successive clusters is defined. Our current implementation uses the latter method since it is less vulnerable to timing inaccuracies.

Considering the 2 dimensional structure of an SHMM as in Fig. 3(b) it can be seen that the expected orientation can be derived from the relative location of successive states. Moreover, when there are multiple transitions a probability for the matching headings can be obtained using the transition probabilities. Hence, we can drop the explicit use of θ without loss of information and obtain a 2-dimensional description of a state.

2) *Limited Observability*: In most cases a mobile robot will not be able to see all or most parts of its environment at all times. More precisely, when it comes to calculating the transition probabilities of the SHMM, it cannot be guaranteed that all parts of the environment are observed for the same duration of time during robot operation. Firstly, this can lead to poor estimates for the transition probabilities. Secondly, the value $N^{(ij)}$ cannot be used as a measure of traffic density, because the relation between the values of different transitions is not known, due to the possibly largely different time frames during which parts of the environment were in the robot's field of view.

To overcome this, we add the overall time of observation of a part of the model to the state's transitions

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

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. Consequently, the transition probabilities are computed as

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

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



Fig. 4. a) The IRobot Create in its environment. b) The office space.

IV. EXPERIMENTAL RESULTS

All Experiments were conducted using our IRobot Create platform, which carries a Hokuyo UTM-30LX laser and a small size notebook (eeePC) (see Fig. 4(a)) for localisation and people tracking. The environment is an open office space of approximately 20x25m as shown in Fig. 4(b) and Fig. 8(a). We first present experiments concerning model learning. To conclude this section we briefly present possible applications of a model of motion patterns as we propose it. In the following figures the robots pose 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.

Firstly, Fig. 5 shows a typical human trajectory modeled as an SHMM. It can be seen, as noted in literature [13], that human trajectories are usually not straight but curved. Also when walking around a corner, a smooth curve rather than a 90 degree turn on the spot can be observed. However, such detail may not be visible in the model any, when more observations are added. This can be observed in Fig. 8(c) where in some sections the model describes almost straight lines (usually in places where many observations were made).




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

The second experiment illustrates the adaptability of the proposed model of motion patterns. In Fig. 6 the robot observed people walking from the bottom to the top of the image. Fig. 6(b) shows the result after 3 trajectories were perceived. Then an obstacle was put close to the learnt path, so that people would have to alter their trajectories slightly. Fig. 6(b) shows how the model slowly adapts to the change until it converges after a while (Fig. 6(c)).

The third experiment shows a similar situation, however this time with a larger obstacle blocking the normal path, forcing people to alter their trajectory more than in the previous experiment. Due to the substantial difference in the observations, it can be seen that the new trajectory is added




Fig. 6. People are movint 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.

to the model in Fig. 7. Above this it can be seen that initially the transition from A to C has a lower probability than the transition from A to B, as indicated by the thickness of the lines in Fig. 7(b). With more observations the transition from A to C gets a higher probability as shown in 7(c). Also due to the change of trajectories the mean of some states shifted slightly.




Fig. 7. People are movint from the bottom to the top while a robot learns the patterns. b) An obstacle is introduced. c) The model converged adapted according to the new information.

The fourth experiment demonstrates the learning of a large model of motion patterns. More than 60 trajectories were observed from various locations in the office environment. The sequence of figures (Fig. 8(a) - 8(c)) show the evolution of the learnt model from a single trajectory to complex motion patterns. Although the model is learnt 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 green rectangle). Once a trajectory has been observed, the model becomes continuous with observed state transitions (see Fig. 8(c)). Another interesting observation can be made inside the yellow 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 sharp corners due to the phenomenon described in the first experiment.

Fig. 8(d) shows the learnt Gaussian distributions of trajectories with which the states and transitions are represented. As the structure (map) of the environment is not taken into consideration, there are some apparent overlaps of




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




Fig. 10. a) A path generated using basic A^* . b) Path generated with regard to motion pattern model, which is significantly longer than the shortest path.

the distributions with obstacles, such as walls. With more observations this effect would be reduced. Lastly, Fig. 9 shows the traffic density in the model. The expectation is that there are more people walking in corridors than in desk areas, which is confirmed in our data.

Finally, we briefly present path planning as a possible application of a model of motion patterns. The details of this are within the scope of another publication. Consider a robot operating in the above office space as a co-worker, i.e. in long term deployment. When people plan a path they do not only plan a shortest path but take social requirements into account as well. One of those is to not disturb co-workers unnecessarily. Fig. 10(a) shows a planned path using common A^* path planning, hence the shortest path is used as the sole criteria. Using the above learnt model, we extended A^* to take the model into account. More precisely, we altered the cost function to use the information about traffic density which can even lead to the robot planning a significantly longer path in order to not cross a desk area as long as the trade-off is not too high (see Fig. 10(b)).

V. CONCLUSIONS

In this publication we presented a novel method to learn a SHMM to represent motion patterns on-line and unsupervised with a mobile robot. 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 the use of a Hidden Markov Model to represent common motion patterns in an environment was proposed.




Fig. 8. a-d) The green arrow, red ellipses, and red lines denote the robots pose, covariance matrices and state transitions respectively. 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 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 in patterns rather than detailed trajectories. Furthermore, the ability to change and adapt the model to accommodate new information was presented.

Finally, the use of such a model for path planning and motion prediction was briefly outlined. In particular the ability to use the motion pattern model for prediction of future poses of moving people is of great interest in mobile robotics. In future work we will exploit this predictive potential for improving tracking and interaction with human peers in the office space. Also the structure of the model can be analysed to identify intersections.

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 *Proc. of the IEEE International Conference on Robotics and Automation (ICRA '99)*, 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*, 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 *Robotics and Automation, 2002. Proceedings. ICRA '02. IEEE International Conference on*, vol. 4, May 2002, pp. 3601 – 3606.
- [7] D. S. D. T. S. Lookingbill, A. Lieb, "Learning activity-based ground models from a moving helicopter platform," in *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, April 2004, 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, December, 2009, (ACRA 2009)*, 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 *Information Theory, 2004. ISIT 2004. Proceedings. International Symposium on*, 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. Archavaleta, 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.