Learning User Models of Mobility-Related Activities Through Instrumented Walking Aids

Jared Glover
School of Computer Science
Carnegie Mellon University

Sebastian Thrun
Computer Science Department
Stanford University

Judith T. Matthews
School of Nursing
University of Pittsburgh

Abstract— **We present a robotic walking aid capable of learning models of users' walking-related activities. Our walker is instrumented to provide guidance to elderly people when navigating their environments; however, such guidance is difficult to provide without knowing what activity a person is engaged in (e.g., where a person wants to go). The main contribution of this paper is an algorithm for learning models of users of the walker. These models are defined at multiple levels of abstractions, and learned from actual usage data using statistical techniques. We demonstrate that our approach succeeds in determining the specific activity in which a user engages when using the walker. One of our proto-type walkers was tested in an assisted living facility near Pittsburgh, PA; a more recent model was extensively evaluated in a university environment.**

I. Introduction

We present a robotic walker for elderly people designed to provide guidance to people who are cognitively or mentally frail and otherwise in danger of getting lost. To assist such people in their daily walking-related activities, it is beneficial for the walker to acquire a model of people's daily routines. Our walker does just this: by passively monitoring people's walking activities, it develops a hierarchical model of people's daily walking routines.

Our walkers extend commercial walking aids, as shown in Figure 1. Both proto-types are equipped with a laser-based navigation system for localization relative to a learned environment map, a display for providing directions to its users, a touch-based interface for receiving commands, and an active drive mechanism equipped with a clutch for switching between active and passive mode. The guidance provided by the walker is similar to car-based GPS systems, in that it informs individual users where to go when attempting to navigate to a target destination [10].

A key ability of our walker is that it learns models of people's motion behaviors. These models are acquired when the device is used with and without providing guidance. The model is defined at multiple levels of abstraction: It includes a representation of principled activities, topological locations through which a person may navigate, and low-level metric locations. A hierarchical hybrid semi Markov model ties together these multiple models into a single coherent mathematical framework. The parameters of the model are learned in a separate teach-in phase, in which a person labels specific activities (e.g., a caregiver). When used for every-

(a) Early prototype (b) Current light-weight walker

Fig. 1. Two robotic walkers developed on top of a commercial walking aid. Both walkers provide navigational guidance and can, though a clutch, be controlled so as to park themselves.

day navigational assistance, our learned model is capable of identifying individual walking-related activities with high reliability. We conjecture that the ability to learn such models and recognize individual activities just from the way it is used is an essential precondition to build truly effective robotic walking aids for the elderly.

Experimental results illustrate that a highly accurate model is learned after only a few days of using the walker. In particular, we have found 100% accuracy in classification of activities when tested on independently collected data—for the duration of an entire testing day.

II. Prior Work

The idea of building robotic walking aids is not new. Most existing robotic devices are active aids—meaning that they share control over motion with the user—and are aimed at obstacle avoidance and path navigation. There exist a number of wheelchair systems [14], [17], [19], [23] as well as several walker- and cane-based devices [5], [13], [9], [21] targeted at blind and elderly people. A technology with some similarities to ours is the walker-based Guido system. Guido evolved from Lacey and MacNamara's PAM-

AID, and was designed to facilitate independent exercise for the visually impaired elderly. It provides power-assisted wall or corridor following [9]. Dubowksy et al's PAMM (Personal Aid for Mobility and Monitoring, distinct from PAM-AID) project focuses on health monitoring and navigation for users in an eldercare facility, and most recently has adopted a custom-made holonomic walker frame as its physical form [6], [25]. Wasson and Gunderson's walkers rely on the user's motive force to propel their devices and steer the front wheel to avoid immediate obstacles [30], [29]. A similar device by Morris et al [21] also provides guidance and force feedback through a haptic interface. All four of these walkers are designed to exert some corrective motor-driven force, although passive modes are available. Our overall approach is similar to [6], [10], [25] in physical shape and appearance, in that it is based on a light-weight off-the-shelf walker frame. The ability to provide guidance is similar in functionality to the one [10], [21]. However, none of these systems learns and analyzes the motion of its users. This paper fills this important gap: our walker is unique in its ability to learn a user model.

Outside the realm of robotic walkers, the idea of learning models of people's motion is not new. Most notably, Bennewitz et al [2], [3] have developed techniques for learning models of people's motion, as observed from a nearby mobile robot. Others have learned behavioral models of people from camera images [1], [7], [11]. The activity of discrete activities is also related to the rich literature of plan recognition [12]. The work here is related, in that it acquires statistical models of behavior. However, it applies these techniques to a new and important domain. Further, our approach integrates learning of behaviors at multiple levels of abstraction, and it ties these together when analyzing high-level activities.

The specific mathematical models proposed here are hierarchical and mixed discrete-continuous. Within the realm of discrete statistical models, a more general class of hierarchical models were proposed in [22], [8], and learning algorithms were presented in [27]. The work here places an instance of this more general mathematical model in the context of a specific application; further, it extends it by a continuous component, as previously proposed for non-hierarchical models in [16].

III. Learning Models of Users

A. *Hierarchical State Space*

Our approach models activities at three levels:

1. The **metric location** of a person operating the walker is comprised of her x-y-location along with her heading direction θ. The location vector at time t is denoted α_t. Determining α_t for an instrumented walker is essentially a metric localization problem, for which a number of effective algorithms exist [4], [15]. In our system, the location α_t is obtained by running the Carmen software package [20].

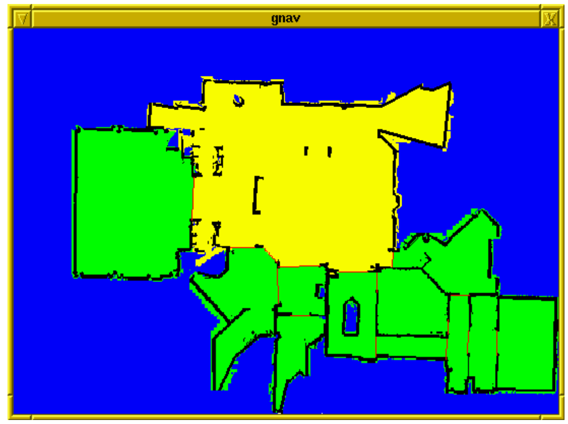

Fig. 2. Topological decomposition of a large foyer environment in the Longwood assisted living facility near Pittsburgh, PA..

2. The **topological location** of a person is determined based on a manually partitioned environment map into topological regions. Regions correspond to rooms, corridors, foyers, and so on. Each of these regions is given a unique identifier. The topological location at time t is denoted β_t. The topological location is a function of the metric location: $\beta_t = g(\alpha_t)$. Since we obtain accurate metric coordinates from our metric localizer, we trivially obtain topological locations as well. Figure 2 depicts a topological decomposition of the environment. While this decomposition was specified manually, algorithms exist for finding similar decompositions automatically [28].

3. The **logical activity** in which a person is engaged forms the most abstract level of our hierarchy. We distinguish two types of activities: Activities carried out in a single location (e.g., a person eating lunch), and activities that involves motion between multiple locations (e.g., walking from the dining hall back to one's room). Each activity is given a unique identifier. The logical activity at time t will be denoted γ_t. In the training phase, we assume the activity is provided (e.g., a caregiver manually labels the data sequence). During everyday operation, the activity is *not* directly observable; thus, we need a statistical framework for estimating activity from sequences of locations.

Clearly, the state at each level changes over time. However, it does so at vastly different time scales. Changes at the metric location level occur continuously, and are reported back at a sample rate of ten Hertz. At the the topological level, changes occur much less frequently: It may take more than a minute for frail elderly people to move from one topological region to another. At the activity level, the change is even slower: An activity can easily persist for half an hour.

To accommodate these vastly different time scales, our approach utilizes different time indices for the different levels. At the lowest level, we use the regular fixed time interval provided by the Carmen software; time will be denoted by t. At the topological level, we will use the time index k. The variable k is incremented whenever the topological location changes. Finally, at the activity level we will use the time index s. The value of s is incremented whenever the activity changes. Both more abstract time indices are variable and depend on a person's actions. Markov chains in which states transition at variable rates are known as semi-Markov chains [18], [26].

The set $B = \{\beta_k, t[k]\}$ denotes the sequence of topological events; here $t[k]$ is the time at which a person's topological location changes. $C = \{\gamma_s, t[s]\}$ shall be the sequence of activities. Again, $t[s]$ models the time at which such a change occurs. We note that it is straightforward to extract the *duration* of an event. For example, the duration of an event in B is given by $\delta_k = t[k+1] - t[k]$.

B. The Hierarchical Probabilistic Semi Markov Model

Our generative probabilistic model—which forms the basis for the inference of activities from data—is defined through four conditional probability distributions that characterize the evolution of state over time. The first two of these distribution operate at the topological time resolution k, whereas the other two are defined for the activity level time s.

- $p(\beta' \mid \beta, \gamma)$ is the the transition probability between topological locations, conditioned on the activity γ. This probabilistic function defines state transitions at the topological level.

- $p(\delta \mid \beta, \gamma)$ is the distribution over durations spent in topological regions β, conditioned on the activity γ. Here δ is a continuous variable. Notice that this distribution is defined over a continuous domain.

- $p(\gamma' \mid \gamma)$ measures the transition probability for activities, modeled at the activity level.

- $p(f(t[s]) \mid \gamma)$ is a time-of-day distribution for activities: It measures the time of day at which an activity γ may be initiated. Here $f(t[s])$ is a function that extracts the time-of-day from a time stamp t by removing the date information. For example, f("11:45:22 on 7/12/2003") ="11:45:22".

Under this model, the probability of the data sequences B, C is then given by the following product:

$$p(B,C) = \prod_k p(\beta_k \mid \beta_{k-1}, \gamma_{k-1})\, p(\delta_k \mid \beta_k, \gamma_k) \cdot \prod_s p(\gamma_s \mid \gamma_{s-1})\, p(f(t[s]) \mid \gamma_s) \quad (1)$$

Clearly, the probabilistic model has been designed carefully so as to model the essentials of activities of elderly people using our walker. For example, our model ignores the specific metric trajectory defined by the variables α; those are only used to calculate the topological region β. The reason for being oblivious to the specific trajectory is its dependence on a great number of factors, such as other people that might block the way. Our specific choice of temporal models—the time a person stays at a single topological location and the time-of-day an activity is initiated, are highly informative: The former allows us to identify activities in which a person stays in the same single topological location for extended periods of time (e.g, watching television). The latter helps us identify activities that occur at regularly scheduled times, such as eating lunch.

C. Learning The Model

The first two probabilities are defined over discrete spaces. Hence, we use a Laplacian estimator for estimating these transition probabilities:

$$p(\beta' \mid \beta, \gamma) = \frac{\sum_k I(\beta_k = \beta' \wedge \beta_{k-1} = \beta \wedge \gamma_{k-1} = \gamma) + c}{\sum_k I(\beta_{k-1} = \beta \wedge \gamma_{k-1} = \gamma) + c|\beta|} \quad (2)$$

Here I is the indicator function which is 1 if its argument is true, and 0 otherwise. The parameter c is the parameter of a Dirichlet prior: It can be thought of as a "pseudo"-observation that prevents transition probabilities of zero (a common technique in the literature on speech recognition). For $c = 0$, this expression becomes the standard maximum likelihood estimator.

Similarly, for the activities γ we have

$$p(\gamma' \mid \gamma) = \frac{\sum_s I(\gamma_s = \gamma; \wedge \gamma_{s-1} = \gamma) + c}{\sum_s I(\gamma_{s-1} = \gamma) + c|\gamma|} \quad (3)$$

The remaining probability distributions are defined over continuous values, but conditioned on discrete variables. Our approach represents these distributions by conditional Gaussian distributions:

$$p(\delta \mid \beta, \gamma) \sim \mathcal{N}(\mu_{\beta,\gamma}, \sigma^2_{\beta,\gamma}) \quad (4)$$

$$p(f(t) \mid \gamma) \sim \mathcal{N}(\nu_\gamma, \tau^2_\gamma) \quad (5)$$

where $\mathcal{N}(\mu, \sigma^2)$ denoted a Gaussian with mean μ and variance σ^2. The mean and variance are obtained using the standard estimation equations:

$$\mu_{\beta,\gamma} = \frac{\sum_k \delta_k \; I(\beta_k = \beta \wedge \gamma_k = \gamma)}{\sum_k I(\beta_k = \beta \wedge \gamma_k = \gamma)} \quad (6)$$

$$\sigma^2_{\beta,\gamma} = \frac{\sum_k [\delta_k - \mu_{\beta,\gamma}]^2 \; I(\beta_k = \beta \wedge \gamma_k = \gamma)}{\sum_k I(\beta_k = \beta \wedge \gamma_k = \gamma)} \quad (7)$$

and

$$\nu_\gamma = \frac{\sum_s f(t_s) \; I(\gamma_s = \gamma)}{\sum_s I(\gamma_s = \gamma)} \quad (8)$$

$$\tau^2_\gamma = \frac{\sum_s [f(t_s) - \nu_\gamma]^2 \; I(\gamma_s = \gamma)}{\sum_s I(\gamma_s = \gamma)} \quad (9)$$

These estimators generate the maximum likelihood Gaussians.

D. Inferring Activities

During everyday use, we cannot observe the activities γ. We are only given the set B of topological transitions, and the times at which γ changes (e.g., detected by a person engaging or disengaging from the walker). The problem of inferring the activities γ from data is then a semi-HMM, short for semi hidden Markov model. Inference for this model can then be carried out using any of the standard HMM inference algorithms, such as the Baum Welch algorithm [24] and its hierarchical extensions [22].

With our walker, we are interested in inferring the present activity of a person in real time. This is achieved by the Bayes filter, an algorithm equivalent to the forward pass in Baum Welch. The Bayes filter calculates, for any time t, the probability that the person's activity is γ_t given the present and past data. If we denote the data up to time t by $B[0;t]$, we seek to estimate $p(\gamma_t \mid B[0;t])$. This expression nicely decomposes, thanks to our choice of the hierarchical model. First, we note that if we define $s*$ as the time index of the most recent activity change, we obtain:

$$\begin{aligned} & p(\gamma_t \mid B[0;t]) \\ &= \sum_{\gamma_{s*}} p(\gamma_t \mid \gamma_{s*}, B[0;t]) \; p(\gamma_{s*} \mid B[0;t]) \\ &= \sum_{\gamma_{s*}} p(\gamma_t \mid \gamma_{s*}, B[s^*;t]) \; p(\gamma_{s*} \mid B[0;s^*]) \quad (10) \end{aligned}$$

Here we split the data B into two parts: $B[0;s^*]$ and $B[s^*;t]$. The set $B[0;s^*]$ contains all items collected before the time at which s^* occurred (this time is denoted $t[s^*]$). The remaining data, gathered in the time interval from $t[s^*]$ through t, is denoted $B[s^*;t]$. The transformation above exploits the fact that the hidden variable γ is the only hidden state in the model—every other state variable is observable. Thus, γ renders the past and future conditionally independent—which is the defining property of Markov chains.

In other words, whenever an activity changes, it suffices to memorize the posterior distribution $p(\gamma_s \mid B[0;s])$ over the activity at that time. Data gathered before that activity change carries no further information relative to the problem of estimating the current activity. This important characteristic of our approach (and Markov chains in general) is documented by the fact that (10) is indeed a recursion.

Unfortunately, activities change slowly. However, a similar Markov property can be exploited for the estimates between activity changes.

$$p(\gamma_t \mid \gamma_{s*}, B[s^*;t]) \;\propto\; p(f(t[s^*]) \mid \gamma) \prod_{\beta_k \in B[s^*;t]} p(\beta_k \mid \beta_k - 1, \gamma) \; p(\delta_k \mid \beta_k, \gamma) \quad (11)$$

This again lends itself nicely to a recursive implementation: While no activity change occurs, the posterior probability of each activity γ is simply updated in proportion to the transition probabilities $p(\beta_k \mid \beta_k - 1, \gamma)$ and the duration probabilities $p(\delta_k \mid \beta_k, \gamma)$.

The resulting algorithm is depicted in Table I. Notice that it is extremely simple: Whenever a state change is observed, the corresponding probability is multiplied into the posterior state estimate. once a posterior estimate of the activity has been obtained, it is straightforward to calculate the likelihood of the data sequence from Equation (1).

Initially, set $\pi(\gamma) = \text{uniform}$ *for all activities* γ.
When activity change detected at time t, *use* $\pi(\gamma') = p(t \mid \gamma') \sum_\gamma p(\gamma' \mid \gamma)\pi(\gamma)$ *as the new estimate (after normalization).*
When the topological location changes from β *to* β' *after being in* β *for a duration of* δ, *multiply* $\pi(\gamma)$ *by* $p(\beta' \mid \beta, \gamma) \cdot p(\delta \mid \beta, \gamma)$ *and normalize.*

TABLE I

ALGORITHM FOR CALCULATING POSTERIORS π OVER ACTIVITIES γ.

IV. RESULTS

We conducted a number of experiments to evaluate the ability of our approach to learn good predictive models of its users. The model learning results were achieved on data

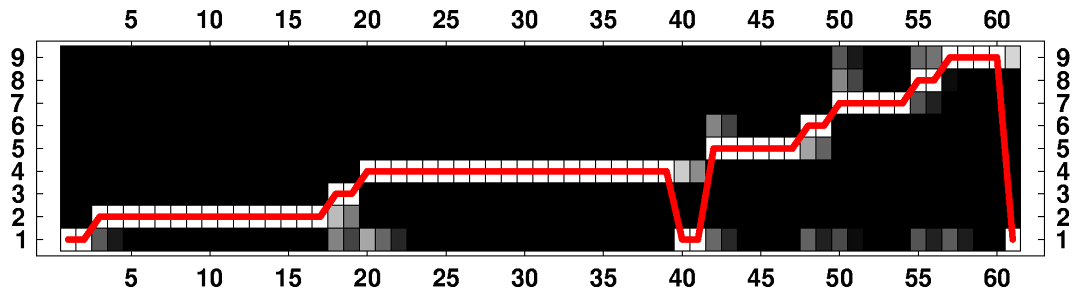

Fig. 3. Predicted activity using our learned model plotted as log-likelihoods, and the actual activity of a person during an entire day. Each time step on the horizontal axis corresponds to a change of the topological location or the activity, and each row corresponds to one of nine different activities. The predictions are remarkably accurate!

collected over a four-day period with an individual user (a student). Figure 4 shows the testing environment, which covers three different floor levels in two different buildings connected by a walkway and two elevators. All results involve genuine motion. For learning, the guidance system was switched off to avoid the obvious bias asserted by the active guidance system. Within those four days, we collected more than 60,000 position data, from which we derived a total of 213 topological state transitions. The map was subdivided into 86 locations. It spanned three different buildings, and within these buildings a total of three different floors, which were accessed through three different sets of elevators. One of the days was withheld from the data to serve as independent testing data; all other data was used for training.

We found that our model predicted people's activity with 100% accuracy, for a total of 61 activities and topological location changes in the testing data. This result is illustrated in Figure 3. Shown there is a sequence of 61 probability distributions over 9 possible activities. Each distribution is plotted as log-likelihood: the brighter an activity, the more likely it is. The red line in this diagram depicts ground truth: clearly, the prediction of activities is remarkably accurate. This illustrates that the features chosen in our model are well-suited for modeling user activities.

Components of the learned model are visualized Figures 5 through 7. Figure 5 shows two examples of topological transition tables for the conditional probability distribution $p(\beta' \mid \beta, \gamma)$. This distribution measures the probability that a person enters region β' from β in activity γ. As should be apparent from this graph, there is a huge diversity of transition functions. For the activity "at lunch," the person remains at a single location (the dining hall), whereas for the activity "returning from lunch" she traverses a number of regions in mostly fixed order.

Figure 6 shows the transition table between activities, that is, the learned probability distribution $p(\gamma' \mid \gamma)$. Again, most activities occur in some sort of sequence, though not all. This remarkably deterministic behavior is a key reason for the high predictive accuracy of our approach. Finally, Figure 7 shows the distribution for the time of day at which an activity is usually carried out. Here we find specific time dependence for a number of activities. This should come as little surprise, since certain activities (such as lunch-related activities) occur at about the same time every day.

Our guidance activities were rather informal, and are mostly documented in [10]. We essentially tested the walker with a number of elderly people, who by and large showed excitement for this new concept. An informal lab evaluation showed that pointing to the next topological region leads to more intuitive guidance than pointing in the direction of the final target location. In a previous related system [21], we found that the guidance can effectively deliver people at locations that they might otherwise be unable to find.

V. Conclusion

We have presented a robotic walker designed to provide guidance to people, and that is able to learn models of people's walking activities. Our approach to learning this model is a hierarchical Markov model that operates at three different levels: A metric motion level at which location is described by metric coordinates, a topological motion level which uses topological regions as its basic element, and an activity level, at which a person's walking activities are logically subdivided into broader categories.

Our model is trained from labeled data. In particular, our approach learned transition probabilities for the two upper levels, and duration and time-of-day distributions. Once learned, it uses Bayesian filtering to determine the specific activity in which a person engages. We find after only a few testing days that our system predicts activities with 100% accuracy on an independent testing day.

While these results are encouraging, more needs to be done to turn this walker into a profitable guidance system. Most importantly, we plan to utilize the learned models in our guidance system, in the hope of providing the right guidance at the right time even if a person fails to specify the target location. This should now easily be possible, given our ability to determine the target location (a function of the

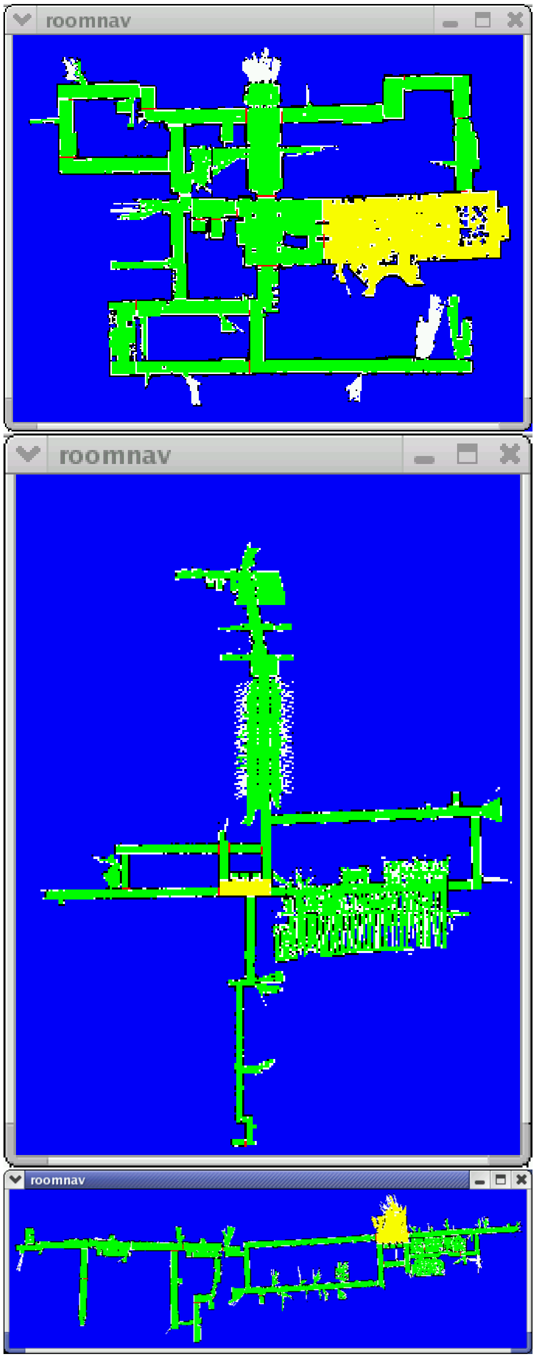

Fig. 4. These three maps together describe the environment in which the walker is being operated. Each corresponds to a different floor, connected by three different sets of elevators. The total distance spanned by these maps is several hundred meters.

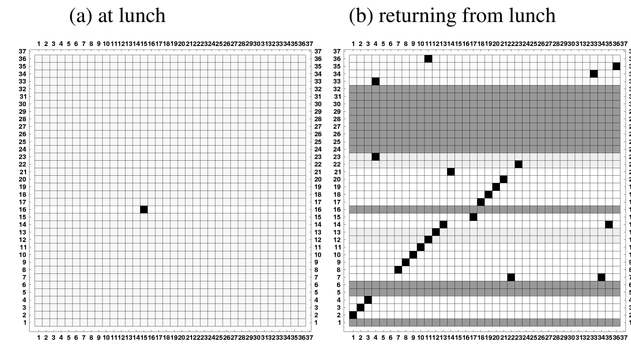

Fig. 5. Two samples of the topological location transition probability $p(\delta \mid \beta, \gamma)$, for the activity "at_lunch" and "returning_from_lunch." The former activity takes place at a single location, whereas the latter involves a long walk back through a number of topological regions.

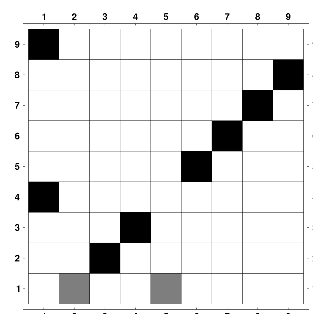

Fig. 6. The activity transition probability table $p(\gamma' \mid \gamma)$ learned from data. Some of the activities tend to occur in sequence.

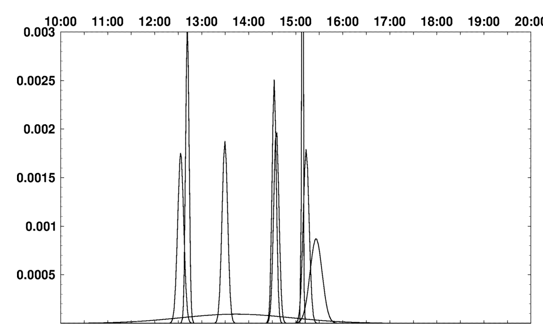

Fig. 7. The Gaussians modeling the time-of-day probability $p(f(t[s]) \mid \gamma)$, for the nine different activities in our model. Some of these activities are remarkably time-specific, whereas others are not.

activity). On the mathematical side, we plan to employ techniques that can automatically segment time series, so as to improve our ability to detect activity change.

Despite these limitations, this paper presents the somewhat surprising result that walking activities can successfully be modeled using relatively little training data, and an appropriately equipped robotic walker.

ACKNOWLEDGMENT

We acknowledge the countless contributions of the members of the class "Assistive Robotic Technology in Nursing and Health Care" (David Holstius, Michael Manojlovich, Keirsten Montgomery, Aaron Powers, Jack Wu, and Sara Kiesler), which took place in the Spring of 2003 at Carnegie Mellon University and the University of Pittsburgh. This research has been sponsored by the National Science Foundation under their ITR Program, which is gratefully acknowledged.

REFERENCES

[1] T. Balch, Z. Khan, and M. Veloso. Automatically tracking and analyzing the behavior of live insect colonies. In *Proceedings of the Autonomous Agents Conference*, Montreal, CA, 2001.

[2] M. Bennewitz, W. Burgard, and S. Thrun. Using EM to learn motion behaviors of persons with mobile robots. In *Proceedings of the Conference on Intelligent Robots and Systems (IROS)*, Lausanne, Switzerland, 2002.

[3] M. Bennewitz, W. Burgard, and S. Thrun. Adapting navigation strategies using motion patterns of people. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2003.

[4] J. Borenstein, B. Everett, and L. Feng. *Navigating Mobile Robots: Systems and Techniques*. A. K. Peters, Ltd., Wellesley, MA, 1996.

[5] J. Borenstein and I. Ulrich. The guidecane: A computerized travel aid for the active guidance of blind pedestrians. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1283–1288, 1997.

[6] S. Dubowsky, F. Genot, and S. Godding. PAMM: A robotic aid to the elderly for mobility assistance and monitoring: A "helping-hand" for the elderly. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA, 2002. ICRA.

[7] A. C. Fang and N. S. Pollard. Efficient synthesis of physically valid human motion. In *ACM Transactions on Graphics, SIGGRAPH Proceedings*, 2003.

[8] S. Fine, Y. Singer, and N. Tishby. The hierarchical hidden markov model: Analysis and applications. *Machine Learning*, 32(1):41–62, 1998.

[9] Lacey G., S. MacNamara, and K. Dawson-Howe. Personal adaptive mobility aid for the frail and elderly blind. In V. Mittal, H. Yanco, J. Aronis, and R. Simpson, editors, *ecture Notes in Artificial Intelligence 1458: Assistive Technology and Artificial Intelligence*, pages 211–220. Springer, New York, 1998.

[10] J. Glover, D. Holstius, M. Manojlovich, K. Montgomery, A. Powers, J. Wu, S. Kiesler, J. Matthews, and S. Thrun. A robotically-augmented walker for older adults. Technical Report CMU-CS-03-170, Carnegie Mellon University, Computer Science Department, Pittsburgh, PA, 2003.

[11] N. Johnson and D.C. Hogg. Learning the distribution of object trajectories for event recognition. In D. Pycock, editor, *Proceedings of the British Machine Vision Conference*, pages 583–592. BMVA, 1995.

[12] H. Kautz and J. Allen. Generalized plan recognition. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, Philadelphia, PA, 1986. IJCAI.

[13] G. Lacey and K.M. Dawson-Howe. The application of robotics to a mobility aid for the elderly blind. *Robotics and Autonomous Systems*, 23:245–252, 1998.

[14] A Lankenau and T. Röfer. A versatile and safe mobility assistant. *IEEE Robotics and Automation Magazine*, 9:29–37, 2001.

[15] J. Leonard, J.D. Tardós, S. Thrun, and H. Choset, editors. *Workshop Notes of the ICRA Workshop on Concurrent Mapping and Localization for Autonomous Mobile Robots (W4)*. ICRA Conference, Washington, DC, 2002.

[16] U. Lerner, B. Moses, M. Scott, S. McIlraith, and D. Koller. Monitoring a complex physical system using a hybrid dynamic bayes net. In *Proceedings of the 18th Annual Conference on Uncertainty in AI (UAI)*, pages 301–310, 2002.

[17] S. P. Levine, D. Bell, L. Jaros, R. Simpson, Y. Koren, and J. Borenstein. The navchair assistive wheelchair navigation system. *IEEE Transactions on Rehabilitation Engineering*, 7(4), 1999.

[18] S. Mahadevan and N. Khaleeli. Robust mobile robot navigation using partially-observable semi-Markov decision processes. internal report, 1999.

[19] M. Mazo, , J. C. Garcia, F. J. Rodriguez, J. Urena, J. L. Lazaro, and F. Espinosa. Integral system for assisted mobility. *IEEE Robotics and Automation Magazine*, 129(1-4):1–15, 2000.

[20] M. Montemerlo, N. Roy, and S. Thrun. Perspectives on standardization in mobile robot programming: The carnegie mellon navigation (CARMEN) toolkit. In *Proceedings of the Conference on Intelligent Robots and Systems (IROS)*, 2003.

[21] A. Morris, R. Donamukkala, A. Kapuria, A. Steinfeld, J. Matthews, J. Dunbar-Jacobs, and S. Thrun. A robotic walker that provides guidance. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2003.

[22] K. Murphy and M. Paskin. Linear time inference in hierarchical HMMs. In *Advances in Neural Information Processing Systems (NIPS)*. MIT Press, 2002.

[23] E. Prassler, J. Scholz, and P. Fiorini. A robotic wheelchair for crowded public environments. *IEEE Robotics and Automation Magazine*, 8:38–45, 2001.

[24] L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. In *Proceedings of the IEEE*. IEEE, 1989. IEEE Log Number 8825949.

[25] M. Spenko. Design and analysis of the smartwalker, a mobility aid for the elderly. Master's thesis, Department of Mechanical Engineering, MIT, Cambridge, MA, 2001.

[26] R.S. Sutton, D. Precup, and S. Singh. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112:181–211, 1999.

[27] G. Theocharous, K. Rohanimanesh, and S. Mahadevan. Learning hierarchical partially observed markov decision process models for robot navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2001.

[28] S. Thrun. Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, 99(1):21–71, 1998.

[29] G. Wasson and J. Gunderson. Variable autonomy in a shared control pedestrian mobility aid for the elderly. In *Proceedings of the IJCAI'01 Workshop on Autonomy, Delegaiton, and Control*, 2001.

[30] G. Wasson, J. Gunderson, S. Graves, and R. Felder. An assistive robotic agent for pedestrian mobility. In *Proceedings of Agents*, pages 169–173, 2000.