

# The use of hidden semi-Markov models in clinical diagnosis maze tasks

Einat Marhasev<sup>a</sup>, Meirav Hadad<sup>b</sup>, Gal A. Kaminka<sup>b</sup> and Uri Feintuch<sup>a,c</sup>

<sup>a</sup>*Caesarea Rothschild Institute for Interdisciplinary Applications of Computer Science, University of Haifa, Haifa, Israel*

<sup>b</sup>*The MAVERICK Group, Department of Computer Science, Bar Ilan University, Ramat-Gan, Israel*

<sup>c</sup>*School of Occupational Therapy, Faculty of Medicine, Hadassah-Hebrew University Medical Center, Jerusalem, Israel*

**Abstract.** In this paper, we investigate the use of hidden semi-Markov models (HSMMs) in analyzing data of human activities, a task commonly referred to as activity recognition. In particular, we use the models to recognize normal and abnormal two-dimensional joystick-generated movements of a cursor, controlled by human users in a computerized clinical maze task. This task – as many other activity recognition tasks – places a lot of emphasis on the duration of states. To model the impact of these durations, we present an extension of HSMMs, called Non-Stationary Hidden Semi Markov Models (NSHSMMs). We compare the performance of HMMs, HSMMs and NSHSMMs in recognizing normal and abnormal activities in the data, revealing the advantages of each method under different conditions. We report the results of applying these methods in analyzing real-world data, from 75 subjects executing clinical diagnosis maze-navigation tasks. For relatively simple activity recognition tasks, both HSMMs and NSHSMMs easily and significantly outperform HMMs. Moreover, the results show that HSMM and NSHSMM successfully differentiate between human subject behaviors. However, in some tasks the NSHSMMs outperform the HSMMs and allow significantly more accurate recognition. These results suggest that semi-Markov models, which explicitly account for durations of activities, may be useful in clinical settings for the evaluation and assessment of patients suffering from various cognitive and mental deficits.

## 1. Introduction

Activity recognition is the process of processing an observed behavioral data stream, tracking and mapping it to a model that allows analysis, recognition, and prediction. The data-stream is generated by observing the actions observable features of an agent (e.g., its position at any time), and the objective of the analysis is typically to infer hidden or unobservable state of the agent, allowing classification of its activity (e.g., recognizing that it is moving along an abnormal trajectory). It is an important task in surveillance, assistive technology, and user-modeling [1,4,12].

Various methods for matching data to models are in use for activity recognition. A popular family of such methods rely on Hidden Markov Models (HMMs) and variants for recognition. Recently, it has become clear that in many cases, the recognition task has a strong temporal component. Thus HMM variations which take duration explicitly into account have become more popular, e.g., Hidden Semi-Markov Models (HSMMs) [12] and Switching Hidden Semi-Markov Models [4]. These allow explicit modeling of the duration of a particular recognized state (i.e., the length of time the observed agent is expected to be in this state). For instance, HSMMs allow modeling the fact that an observed airline passenger may spend a long time at the security check area, and much less time in other areas connecting the check-in area to the flight boarding gate.

These models, however, do not account for changes in *transition* probabilities based on the duration an agent has spent in a given state. For instance, an airline passenger spending a short time at the check-in counter is more likely to have gone through the check-in process without glitches, and thus more likely to head towards the gate. On the other hand, a passenger spending a long time at the check-in area is more likely to have encountered problems, and thus the likelihood of the passenger heading towards a supervisor is greater. Existing models encounter difficulty modeling such distinctions, and are therefore insufficient for the complex set of recognition problems found in real-world applications.

In this paper, we investigate the use of Markov models that explicitly model the dependency of transition probabilities on state duration. We present an augmentation of the standard semi Markov models, called *Non-stationary Hidden Semi Markov Models* (NSHSMMs) which enable more accurate recognition of duration-based behavioral patterns. Just as HSMMs expand HMMs to include state duration (to more accurately model realistic observed behaviors), so do NSHSMMs expand HSMMs to model the transition probabilities' dependence on state duration. Variants of NSHSMMs have been previously applied in speech recognition and handwritten character recognition [25,22], but their use in activity recognition is novel.

Indeed, we contrast the performance of the different models in two clinical diagnosis tasks, with 75 subjects. Specifically, we use the models with real world data by administering the maze tasks to healthy young adults who were instructed to solve computerized mazes using a joystick. Such tasks are often employed for clinical assessment of cognitive impairments. Being able to detect and classify patterns of movement may provide relevant information regarding the nature of a patient's impairment or help assess his or her progress [23,26]. Up until now, such patterns were only recognizable to trained experts. To our best knowledge, this is the first attempt at using activity recognition methods in this type of task. The results show that both HSMMs and NSHSMMs are successful in recognition of behavioral patterns which were not revealed by the conventional HMMs. The NSHSMM model, described in this paper, outperforms the HSMM in one of the tasks, while the HSMM outperforms the NSHSMM in the other.

This paper is organized as follows. The next section presents related work in activity recognition. Section 3 presents the NSHSMM model. Section 4 illustrates the use of the models on synthetic data, for clarity. Section 5 reports on results of recognition in the clinical diagnosis task. Section 6 concludes.

## 2. Related work

Hidden Markov Models (HMMs) and their many variants have been successfully applied to a number of scientific and engineering problems. Rabiner reviews the theoretical facets of HMMs [21]. He provides a formal definition of the elements of HMMs and a discussion of the three fundamental problems of HMMs, including methods for solving these problems. Rabiner also reviews in this work the types of HMMs that have been studied and shows an application of HMMs in speech recognition.

Although HMMs have been studied and used in several areas of applications, their use in analyzing behavioral data for activity recognition is fairly recent. Oliver et al. [18] developed a representation based on layered HMM (LHMM) that offers real-time interpretations of human activity in an office environment. The layers in LHMM are designed to portray different levels of abstraction operating at different temporal granularity of observation sequences in HMMs. In this work, the characteristics of the LHMM are used to distinguish between activity sets taking place in an office domain. LHMMs are compared with standard HMMs in terms of scalability, space needed for learning, robustness to changes, and avoiding over-fitting.

Han and Veloso [11] introduce an HMM-based recognition algorithm that performs recognition of robot behavior in a soccer game. This algorithm recognizes multiple behaviors, where each behavior is represented by a single HMM, called a Behavior HMM (BHMM). In order to recognize a single execution of a behavior, the BHMM must be instantiated around the time when the behavior begins executing. Therefore, Han and Veloso's algorithm performs probabilistic segmentation and the BHMMs are created periodically at regular intervals to capture the starting time of the behaviors.

To handle the hierarchical formation of complex activities, Bui et al. [1] used another variant, Abstract Hidden Markov Model (AHMM), for vision-based activity recognition. The AHMM is a hierarchical probabilistic model for representing a hierarchical abstraction of an agent's activity at changing levels of detail in uncertain and noisy domains. Yin et al. used AHMMs to track users moving in a wireless environment [27].

As Rabiner points out [21], in conventional HMMs the duration of states is not modeled explicitly. Instead, an exponential state duration distribution is implicitly created when self-transitions (transitions from a node to itself) are used to allow a system to remain in its current state. This has been shown to be problematic in recognizing real-world activities. Hidden Semi-Markov Models have been developed to allow explicit modeling of duration probabilities. The following works are examples for the use of HSMMs in various fields.

In conventional hidden semi-Markov models, the probability of state duration is established by a distribution function [8,15]. Levinson [15] suggested constructing a continuous variable-duration hidden Markov model (CVDHMM) by substituting the probability distribution duration in HMM with continuous probability density functions. Levinson used the gamma density function for state duration probability and described the estimation procedure for this model's parameters with a proof of the technique's appropriateness. The model's performance was evaluated using simulated data for automatic speech recognition. Levinson's work has been studied and referred to in many speech recognition papers as well as papers which focus on probabilistic models.

Yu and Kobayashi [28] proposed a hidden semi-Markov model that allows for missing data and multiple observation sequences. They applied it to mobility tracking in wireless networks. They provided a categorization of observation patterns, and introduced estimation algorithms for missing observations and for two observation sequences. In order to evaluate the models and algorithms, Yu and Kobayashi used simulation experiments of mobility tracking in wireless internet services provisioning. Their work underlines the importance of addressing the possible missing patterns of observations in trying to achieve better estimation procedures for HSMMs in applications where these types of observation sequences are potential input to the models.

Hongeng and Nevatia [12] applied HSMMs for recognizing events in a video stream. The probabilities of *primitive* events were inferred by using Bayesian networks, while *composite* events, which are complex events composed of a sequence of *primitive* events, were recognized using a hidden semi Markov finite state model (essentially, HSMM). Hongeng and Nevatia compared the use of standard HMMs to that of HSMMs in event segmentation tasks. They reported that in HSMM the event segmentation in noisy video sequences is more reliable than in HMMs, and by assuming uniform or normal duration distribution, Hongeng and Nevatia composed a less complex algorithm for inference using HSMMs.

Natarajan and Nevatia [17] later combining the work on HSMMs with work on coupled hidden Markov models, to form Coupled Hidden Semi-Markov Models (CHSMMs). These allow both modeling durations of states (as in HSMMs), as well as modeling of interactions between otherwise independent observation streams (e.g., when two independently-tracked humans meet). Natarajan and Nevatia show that the computational complexity of reasoning using CHSMMs is higher than either HSMMs or coupled

HMMs, and provide a number of efficient algorithms for such reasoning. In contrast to Natarajan and Nevatia's work, this paper does not focus on efficient algorithms, but on contrasting the use of the representational power of HMMs, HSMMs, and NSHSMMs for clinical diagnosis task. This paper also does not consider coupled models.

Duong et al. [4] introduced switching hidden semi-Markov models (S-HSMM), which are essentially two-layered extensions of the HSMM, for high-level activity recognition and abnormality detection. Their work included a formal definition of S-HSMMs, parameter estimation methods and experimental validation for the innovative model. They applied S-HSMMs to the tasks of learning and making online classifications in series of home-related activities, and to the task of abnormality detection in an elderly care domain. Duong et al. also demonstrated the S-HSMM's superiority over the flat HSMM in these tasks.

These previous works concentrate on HSMMs in which the duration distribution does not affect the state transition probabilities. In contrast, Non-stationary Hidden Semi-Markov Models (NSHSMMs) explicitly model not only state duration likelihood, but also the dependence of different transition probabilities on the duration.

NSHSMMs have been applied in speech technology and in handwritten character recognition [22]. In Vaseghi's [25] time-dependent state transition HMM, the state transitions are obtained from the state durations' cumulative density distribution functions and are formulated on the basis of a specific type of HMM, a left-right HMM, commonly used in speech modeling. Vaseghi discusses the re-estimation and recognition algorithms for this model and presents experimental results for recognition isolated spoken words with a data set of spoken English alphabet. The results show an increase in recognition accuracy using the proposed modified HMM.

Sin and Kim [22] used NSHSMMs to recognize on-line handwritten characters, and showed that they outperformed both standard HMM and HSMM. Furthermore, Sin and Kim deduced that NSHSMMs provide high recognition accuracy in online handwritten character recognition tasks. The NSHSMM equations we present later on are new, compared to Sin and Kim's model. In their work, Sin and Kim used NSHSMMs on data which has different characteristics than the type of data observed in the clinical diagnosis task we discuss in this paper. In particular, recognition of handwritten characters should take into account the history of all possible events from the beginning of the observation sequence. In contrast, for the clinical diagnosis task, we are only interested in the previous state for the probability calculations (i.e., pure Markovian models). The resulting model (presented below) is therefore simpler than that of Sin and Kim's handwritten character recognition model.

Our work also contrasts with other probabilistic approaches to activity recognition. Kaminka et al. [13] have investigated overhearing by plan recognition in cooperative agent teams. They have utilized a recognition model that is reminiscent of hierarchical non-hidden semi-Markov models. The system used average duration of plan-step as an input to a Gamma distribution, to estimate the duration of unobserved plan-steps. Kaminka et al. report that the use of this model produced very poor results in accurately predicting the joint states of multi-agent teams, due to the unmodeled variance of actual durations. They use this to motivate their work on multi-agent heuristics which compensate for the technique's drawbacks and boost recognition accuracy.

Several probabilistic inference models have been applied to activity recognition (e.g., [9,20,2]). These works do not model state duration in their computations but highlight the necessity for methods which optimize the probabilistic recognition models for goals and activities. For instance, Fox et al. [9] survey Bayesian filter techniques and implementations while focusing on location estimation tasks. By representing uncertainty at different levels of abstraction, these techniques are important for situations

in which location sensors provide imperfect measurements or errors, circumstances common to any real world application. Like Markov models, Bayesian filters rely on the Markov assumption. They do not use statistics on state duration in their computations.

Various works in recognition suggest using statistical models in order to deal with uncertainty in recognition tasks. These include parsing of probabilistic state-dependent grammars (PSDG) [20], and the Bayesian-networks models by Charniak and Goldman [2]. In both of these, there is no explicit treatment of the importance of the amount of time spent in a certain state. Some recent work [10] does consider state duration in Bayesian networks but it is not applied to activity recognition with hidden states as in our applications.

Our own paper in the MOO (Modeling Others from Observations) workshop [16] introduced an earlier version of the NSHSMM model, and a subset of the examples appearing in Section 4, used to illustrate the use of the model. In contrast, this paper discusses the model and associated methods in full detail, and reports on its use in clinical diagnosis tasks.

### 3. Non-stationary recognition models

We begin with a short overview of familiar models (HMMs and HSMMs), before providing a full discussion of the NSHSMM model and its use in activity recognition.

#### 3.1. HMMs and HSMMs

An HMM (Hidden Markov Model) is defined in [21] as:

1. A set of  $N$  states  $S = \{S_1, S_2, \dots, S_N\}$ . A state at time  $t$  is denoted as  $q_t$ .
2. The state transition probability distribution  $A = \{a_{ij}\}$ , where  $\{a_{ij}\}$  is the probability to transition from state  $S_i$  to state  $S_j$ .  $a_{ij} = P[q_{t+1} = S_j | q_t = S_i] \quad 1 \leq i, j \leq N$ .
3. The initial state probability distribution  $\pi = \{\pi_i\}$  where  $\pi_i = P[q_1 = S_i] \quad 1 \leq i \leq N$ .
4. A set of  $M$  observation symbols  $V = \{v_1, v_2, \dots, v_M\}$ .
5. The observation symbol probability distribution in state  $j$ ,  $B = \{b_j(k)\}$ , where  $b_j(k) = P[v_k \text{ at } t | q_t = S_j], \quad 1 \leq j \leq N \quad 1 \leq k \leq M$ .

The parameters of HMM are represented by  $\lambda = (A, B, \pi)$ . Given a sequence of observations  $O = O_1 O_2 \dots O_T$  and  $\lambda$ , the probability of the observing sequence can be evaluated as:

$$P(O|\lambda) = \sum_{i=1}^N \alpha_T(i) \quad (1)$$

for any  $t \in [1, T]$ , the forward probability  $\alpha$  is defined by:

$$\alpha_t(j) = \sum_{i=1}^N \alpha_{t-1}(i) a_{ij} b_j(O_t) \quad (2)$$

As noted in previous work [15], in the standard HMM, the transition probability distribution  $a_{ij}$  is constant in time. At each time step, the model transfers to the next state according to the transition probabilities and the symbol observed. To model state duration, a self-transition (a transition from the state to itself) is used, with transition probability  $a_{ii}$  (Fig. 1). Thus the HMM standard model does not

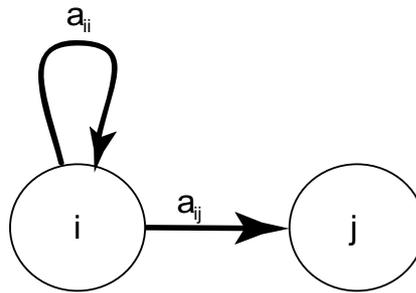


Fig. 1. HMM with exponential state duration.

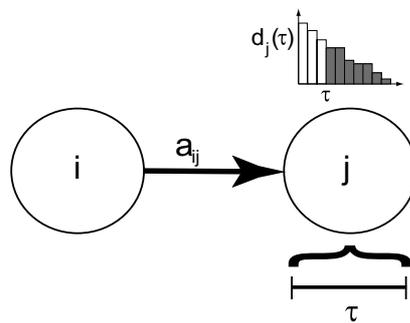


Fig. 2. HSMM with explicit state duration.

differentiate between a self-transition and a transition to another state when calculating the probability of an observation sequence, causing an exponential probability distribution for remaining in the current state. The likelihood of remaining in state  $i$  for  $t$  time-steps, is  $a_{ii}^t$ .

Since it means that there is no accounting for the amount of time spent in a certain state, for many applications, this exponential state duration probability distribution is inappropriate. For this reason, the HSMM (Hidden Semi-Markov Model) was devised to broaden the HMM by eliminating self-transitions (see Fig. 2), and instead, modeling the state duration explicitly.

As in the HMM model,  $A$  is the state transition probability distribution  $A = \{a_{ij}\}$ , where  $\{a_{ij}\}$  is the probability to transition from state  $S_i$  to state  $S_j$ .  $1 \leq i, j \leq N$ ,  $i \neq j$ . However, the self-transition  $a_{ii}$  is replaced with duration probability density function  $d_j(\tau)$  which denotes the probability of staying at least of length duration  $\tau$  in state  $S_j$ ,  $1 \leq \tau \leq D_j$ . The rest of the HMM definitions apply also for the HSMM.

In particular, the probability of observing a sequence, in HSMM, is computed in the same manner as for HMMs Eq. (1), but the forward probability  $\alpha$  is defined differently. Specifically, for any  $t \in [1, T]$  and  $\tau \leq t$ , the forward probability  $\alpha$  is defined by:

$$\alpha_t(j) = \sum_{\tau=1}^{D_j} \sum_{\substack{i=1 \\ i \neq j}}^N \alpha_{t-\tau}(i) a_{ij} d_j(\tau) \prod_{\theta=1}^{\tau} b_j(O_{t-\tau+\theta}) \tag{3}$$

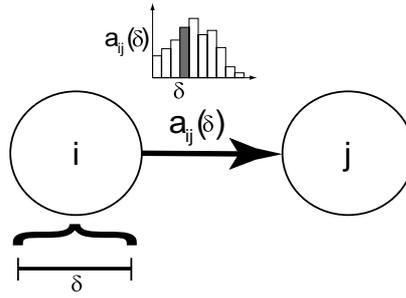


Fig. 3. NSHSMM with state transition probability distribution.

### 3.2. Non-stationary hidden semi-markov model

We consider an expanded model which includes not only the modifications of HSMM to HMM but a refinement to the state transition probability to include time dependency.

Formally, we define the state transition probability distribution of an NSHSMM as follows.  $A = \{a_{ij}(\delta)\}$ , where  $\{a_{ij}(\delta)\}$  is the transition probability from state  $S_i$  to state  $S_j$  after having remained in state  $S_i$  for a duration of length exactly  $\delta - 1$  (a pictorial illustration is given in Fig. 3). This definition of the transition probability reflects the non-stationary property we wish to capture. By assigning the transition probability to be the exact probability of spending a specific duration before the transition to a certain state, we take into account the concept that different durations in a state affect the choice of where to transition to. As explained by the airline passenger example (introduction section), the NSHSMM takes into account not only the explicit duration of a state but also the choice of where to transition to according to that exact duration. In doing so, the NSHSMM can improve the accuracy of the recognition abilities of the model.

Taking in consideration the time-dependent transition probability, we define the forward probability  $\alpha$  of NSHSMM as:

$$\alpha_t(j) = \sum_{\tau=1}^{D_j} \sum_{\substack{i=1 \\ i \neq j}}^N \sum_{\delta=1}^{D_i} \alpha_{t-\tau-\delta+1}(i) a_{ij}(\delta) d_j(\tau) \prod_{\theta=1}^{\tau} b_j(O_{t-\tau+\theta}) \quad (4)$$

We initiate the computation for each state by the probabilities to begin at that state and remain there until  $D_i$ . We denote this initiation as  $\alpha_t^\tau(j)$ :

$$\alpha_t^\tau(j) = \pi_j d_j(\tau - 1) \prod_{\theta=1}^{\tau} b_j(O_{t-\tau+\theta}) \quad (5)$$

where  $j = 1 \dots N$ ,  $t = 1 \dots D_j$  and  $D_j$  denotes the maximum length duration being in state  $S_j$ . Thus, given Eq. (4) above, for any  $t \in [1, T]$  the likelihood function  $P(O|\lambda)$  of NSHSMM can be evaluated by Eq. (1).

It is important to notice the difference in meaning between  $a_{ij}(\delta)$  and  $d_j(\tau)$ . In Eq. (4),  $D_i$  refers to the maximum length that an agent may spend in state  $S_i$  ( $i$  is the state an agent transitions from) and  $D_j$  refers to the maximum length that an agent may spend in state  $S_j$  ( $j$  is the state an agent transitions into). It is clear from this representation that an agent therefore cannot spend more than  $D_i + D_j$  time-steps in these states. As previously described,  $a_{ij}(\delta)$  denotes the transition probability from state  $S_i$  to state  $S_j$

after having remained in  $S_i$  for exactly  $(\delta - 1)$  time-steps and  $\delta$  is bounded by  $D_i$ . In contrast,  $d_j(\tau)$  denotes the probability of remaining at state  $S_i$  for at least  $\tau$  time-steps and  $\tau$  is bounded by  $D_j$ .

Examining the computation involved in the calculation of the likelihood function  $P(O|\lambda)$  of NSHSMM shows that  $\alpha_t(j)$  requires an order of  $O(ND^2 + T)$  calculations, where  $D = \text{argmax}_{D_i}$  and therefore the computation of  $P(O|\lambda)$  requires  $O(N^2TD^2)$ .

In the following sections, we compare the performance of each model in recognizing normal and abnormal behavior when  $d_j(\tau)$  and  $a_{ij}(\delta)$  are discrete probability distribution functions, estimated from the data. In the case of discrete functions, the space involved is in the order of  $O(ND^2 + TN)$ , where  $O(TN)$  is the space needed for the  $\alpha$  values table. The algorithms for estimation of the probability density functions are discussed in the following section.

### 3.3. Parameter estimation

This section describes the estimation of the models' parameters appropriate for each model's parametric probability characteristics. In all three models, the parameter estimates are drawn from the training set given to the models, i.e., the observation sequences available for training. The observations, transitions and the remaining time in each state are counted and the probability measures,  $A$ ,  $B$  and  $\pi$  are then estimated based on the processed data. We refer to each such observation sequence in the training set as an example.

In all three models, the estimates of the  $\pi$  and  $B$  parameters are done in the same manner. The initial state probability in state  $S_i$ , as shown in [21], is estimated by

$$\pi_i = \text{the number of examples in state } S_i \text{ at time step } t = 1$$

With the following constraint kept, by normalization:

$$\sum_{i=1}^N \pi_i = 1 \tag{6}$$

The observation symbol probability for symbol  $v_k$  and state  $S_j$  is estimated by [21].

$$b_j(k) = \frac{\text{the number of examples where } v_k \text{ was observed while in state } S_i}{\text{the number of examples that go through state } S_i}$$

With the following constraints kept (using normalization):

$$\sum_{i=1}^M b_j(k) = 1, 1 \leq j \leq N \tag{7}$$

Previous sections discuss in detail the differences between HMM, HSMM and NSHSMM in their state duration probability density functions and in their transition probability density functions. The models account for these differences in the parameter estimation procedures.

In standard HMM, the transition probability  $a_{ij}$  is constant in time and is estimated in [21] by

$$a_{ij} = \frac{\text{the number of examples where transitions take place from } S_i \text{ to } S_j}{\text{the number of examples where transitions from } S_i \text{ (to any other state) occur}} \tag{8}$$

With the following constraint kept (by normalization):

$$\sum_{i=1}^N a_{ij} = 1, 1 \leq i \leq N$$

In HSMM, the state duration differs from that in HMM. HSMM does not have a self transition  $a_{ii}$  in state  $S_i$  (which is calculated in the same manner as a transition  $a_{ij}$  from one state  $S_i$  to another  $S_j$ ), and instead, it has a duration probability density function  $d_i(\tau)$ .  $d_i(\tau)$  is estimated by

$$d_i(\tau) = \sum_{j=\tau}^{D_i} \frac{\text{the number of examples with a } j\text{-long stays in } S_i}{\text{the number of examples with a stay in } S_i} \quad (9)$$

The summation Eq. (9) represents the probability of remaining in state  $S_i$  for *at least*  $\tau$  time steps. Thus, the state duration probability will begin at maximum and gradually reduce for durations where there is evidence of shorter remaining periods.

Because the self-transition  $a_{ii}$  does not exist in HSMM, the constraints over parameter  $A$  were tuned accordingly (see Eq. (10)).

$$\sum_{\substack{j=1 \\ j \neq i}}^N a_{ij} = 1 - d_i(\tau), 1 \leq i \leq N, 1 \leq \tau \leq D_j \quad (10)$$

NSHSMM is a refinement of the transition probability density distribution  $A$  in HSMM, to include a time dependency. We define  $e_{ij}^\delta$  as the number of transitions from  $S_i$  to  $S_j$  after remaining in  $S_i$  exactly  $(\delta - 1)$  time steps. The transition probability  $a_{ij}(\delta)$  from state  $S_i$  to state  $S_j$  in NSHSMM is then estimated by

$$a_{ij}(\delta) = \frac{e_{ij}^\delta}{\sum_{k=1}^N e_{ik}^\delta} \quad (11)$$

Equation (12) presents the constraints on the transition probability distribution  $A$  in NSHSMM. These constraints are the reason for the alterations made to adjust for the time-dependent changes in NSHSMM.

$$\sum_{\substack{j=1 \\ j \neq i}}^N a_{ij}(\delta) = 1 - d_i(\delta), 1 \leq i \leq N, 1 \leq \delta \leq D_i \quad (12)$$

Given the equality constraints over the parameters  $A$ ,  $B$  and  $\pi$  presented in this section, for any  $t \in [1, T]$ , the likelihood function  $P(O|\lambda)$  can be evaluated according to each of the three model by Eq. (1), as discussed in preceding subsections.

#### 4. Illustration using synthetic data

In this section, our goal is to illustrate in depth the differences in capabilities and performance between the three models, HMM, HSMM and NSHSMM. To do this, we use simulated data. We present two small-scale experiments, carefully designed to bring out the characteristics of each model.

To test and compare the three models for the task of recognizing agents' activities, we utilized a simulator which simulates up to dozens of passengers moving about a large area according to pre-defined paths (with noise in the movement). This simulator is part of an industry effort in surveillance. It includes a window which shows the position of dynamic objects (e.g. agents) in a  $XY$ -grid, and simulates the (noisy) movements of these objects. The simulator enables the choice of the path and the speed of an object's movement on the grid as well as the ability to determine its deviation from the predefined path.

In our experiments, the agents can move within a square-shaped  $200 \times 200$  grid. Each grid cell represents a unique single position defined by two coordinates,  $X$  and  $Y$ , both ranging from 0 to 199. An agent's movements, which are represented by  $(x, y)$  positions in the grid followed by the specific time that the movements were observed, are recorded into a file. A series of detected movements with their appropriate occurrence times represent a single path. Thus, each file includes a path which the agent has taken. We denote the time marking each movement in the recorded files as *time steps*. In these experiments, a single second is divided into five time steps.

We generated output files to build training sets for our models and then we tested the models' ability to detect paths which are consistent with the training sets. In these experiments, the training sets for our models were groups of 40 paths which the models received. Using the training set, the models learned all paths in the training set, and calculated the probabilities associated with the model, e.g., transition probabilities from one state to another, state-duration probabilities, etc.

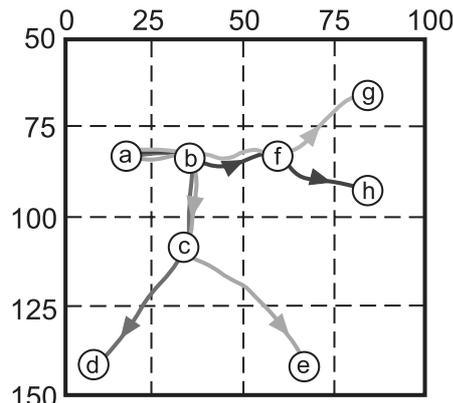
The goal of these experiments is to illustrate the differences between the models. In the general case,  $(x, y)$  positions are observed with uncertainty (i.e., hidden). We have compared models that relate equally to the hidden property, as can be seen in the likelihood function equations. In the experiments, the observation probability of a state is identical in all the models. For this reason, the observation probability of each state is 1. This means that the ratio of probabilities between the models remains (with lower resulting probabilities). Hence, we overlooked the hidden part in these experiments, but addressed the theoretical grounding as it is important in real-world applications.

For the following experiments, we used a 64-block tessellation of  $25 \times 25$  size blocks for the simulated data. Each block represents a single state in the model. An observation, which is a reading of a certain position in the grid, is associated with a state if it's position in the grid is within the state's scope. We chose a  $25 \times 25$  block tessellation since this block size distinguishes areas in which the agents linger from areas where agents pass. For instance, in the airline passenger example, the boarding area is an area where a passenger is prone to be stalled and therefore it is a lingering state, while an airport corridor is an area where a passenger is more likely to pass without making any stops.

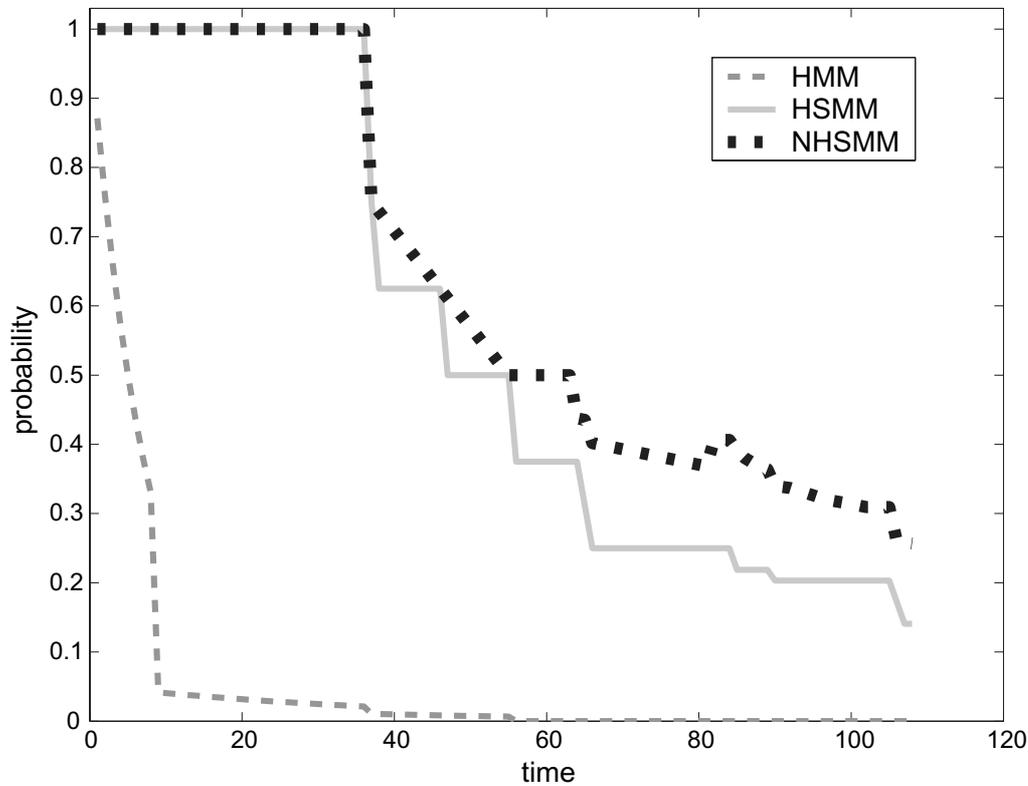
In the experiments discussed in the following sections, we used the cross-validation technique to estimate the accuracy of the models' ability to classify a path according to it's resemblance to the paths in the training set.

For the first experiment we conducted, we used simulated data that is of a non-stationary nature. The goal of this experiment was to test the behaviors of HSMM and NSHSMM in this type of data and compare their ability to estimate sequences of observation which are compatible with the non-stationary group of paths provided to the models as a training set. In this experiment, the non-stationary training set was devised of four types of behaviors on the grid. Each type differs from the other by its motion rate and by the designated positions along its path.

Figure 4(a) shows a sketch of the types of the four paths. All four paths start at the same position, denoted  $a$ , and move together on the grid. When they all reach the point marked  $b$  (coordinate  $(60, 60)$  on the  $XY$ -grid), the four paths separate into two different directions, based on the amount of time spent in the state encompassing  $b$ . The two slower paths keep moving straight towards point  $f$  after having stayed



(a) A sketch of the four paths comprising the training set for this experiment. Movement is left-to-right.



(b) A graph showing the resulting probability values for each model from testing the four types of paths of agents moving on the grid according to the paths described over a period.

Fig. 4. First experiment: Settings and Results (averaged over 30-some different trials).

a longer period of time in the area of *b*, and the two faster paths turn to the right, towards *c*. Further along, another split occurs between each of the two pairs of paths. In *c*, the trajectory turning right (towards *d*) remained a shorter period of time than the trajectory turning left (towards *e*). Similarly, in *f*,

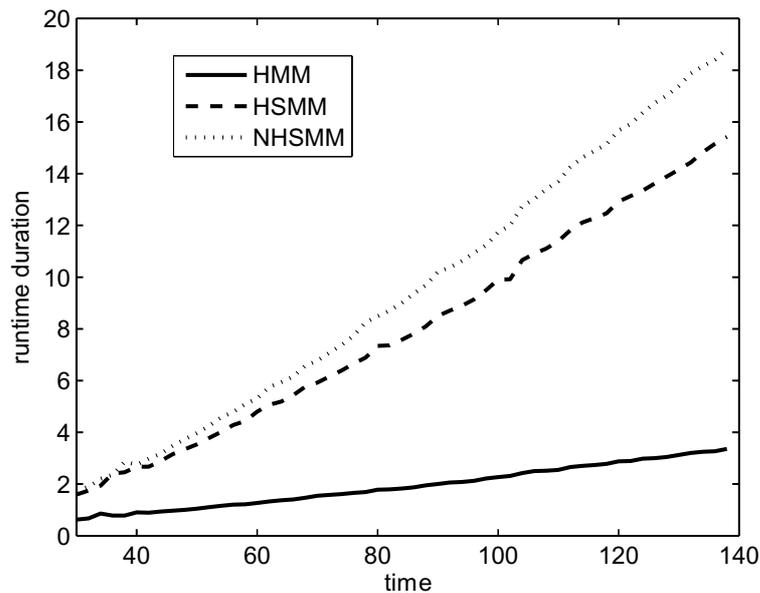


Fig. 5. First experiment: A graph showing the run-time of each model (in seconds) over observation time steps.

the trajectory turning right towards  $h$  remained in  $f$  less time than the other trajectory (heading towards  $g$ ).

Figure 4(b) shows the results of comparing the ability of the three models to estimate the probability of observation sequences that are compatible with all four types of behaviors comprising the non-stationary data introduced above. We see an exponential decline in the graph of the HMM. For the HSMM and NSHSMM, the graph first shows very high probabilities for both models. These are the high probabilities received in the first phase of the training set, where all four paths are moving together on the grid. At time step 42 there is a clear drop in both graphs and further along the graphs appear to gradually decline with a few more noticed drops. These drops in the probabilities of the tested observation sequences are explained as before, by the splitting of the paths causing a descent in the probability of transitioning to each state along each path.

However, from time step 42 where the first split occurs, the graphs of HSMM and NSHSMM no longer converge. The NSHSMM graph surpasses that of the HSMM throughout the time and significantly outperforms it (paired t-test,  $p$  value is  $2.0909 \times 10^{-21}$  i.e., much less than 0.001). This can be explained by the difference in the transition probability functions of the two models. While the HSMM has a constant transition probability, the NSHSMM holds an explicit transition probability distribution from state to state allowing the NSHSMM to model the non-stationary behavior of the data. These results show the ability of the NSHSMM to capture the non-stationary nature of the given data and its superiority over HSMM in cases where the time spent in a certain state has implications on where to proceed to in the tested environment.

Figure 5 refers to the first experiment and compares the run-time for calculating the likelihood function over consecutive observations by each of the three models. As can be seen, there is a substantial increase in computation time when comparing HMM with both HSMM and NSHSMM. However, between HSMM and NSHSMM, the additional computation in NSHSMM appears minor.

In order to better highlight the difference between HSMM and NSHSMMs, we compared their ability to detect non-stationary abnormal behavior. Our goal was to examine, in this second experiment, how

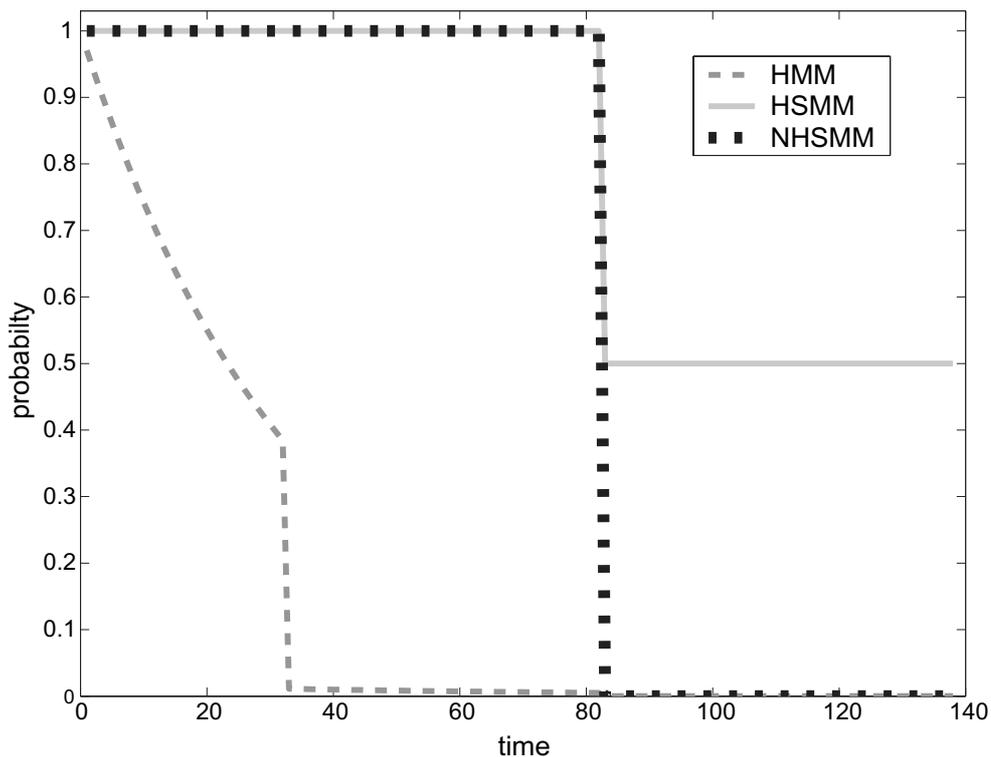


Fig. 6. Second experiment: Settings and results. A graph showing the resulting probability values for each model for their ability to detect non-stationary abnormal behavior. Results are averaged over 30-some different trials.

well the models detect an observation sequence which is not compatible with the training set in terms of duration, but is compatible in terms of the states in the models that it visits.

The training set for this experiment was comprised of two types of behaviors over a period of 140 time steps. Both types of behaviors start at the same point in the  $XY$ -grid and continue to the next state, where they spend a different amount of time and each turns to opposite directions. In other words, the path spending a short period of time at the splitting state turns right and the path which lingers at the splitting state, turns left. Thus, non-stationary abnormal behavior in this experiment is an observation sequence that remains a short number of time steps at the splitting state but turns left (instead of right).

Figure 6 shows the likelihood probability graphs of the non-stationary abnormal behavior for each of the three models. After examining the resulting graphs of HMM and HSMM in this experiment, it becomes very clear that these two models do not detect the non-stationarity of the abnormal behavior and refer to this observation sequence as matching the training set. Due to long remaining periods in certain states, the HMM graph decreases exponentially over time. This type of behavior in HMM was seen in prior experiments using stationary and non-stationary data. The low probability seen the HMM graphs is due to its exponential probabilities, and the graph in this experiment is similar to graphs of HMM probability likelihood for observations matching the training set, making HMM inappropriate for detecting the non-stationary of the abnormal behavior.

Figure 6 also shows an abrupt drop in the graphs of both HSMM and NSHSMM is at time step 84, which according to the two type paths composing the training set, is the point where the two paths split, with one path making a turn to a certain direction and the other path remaining at that point a while

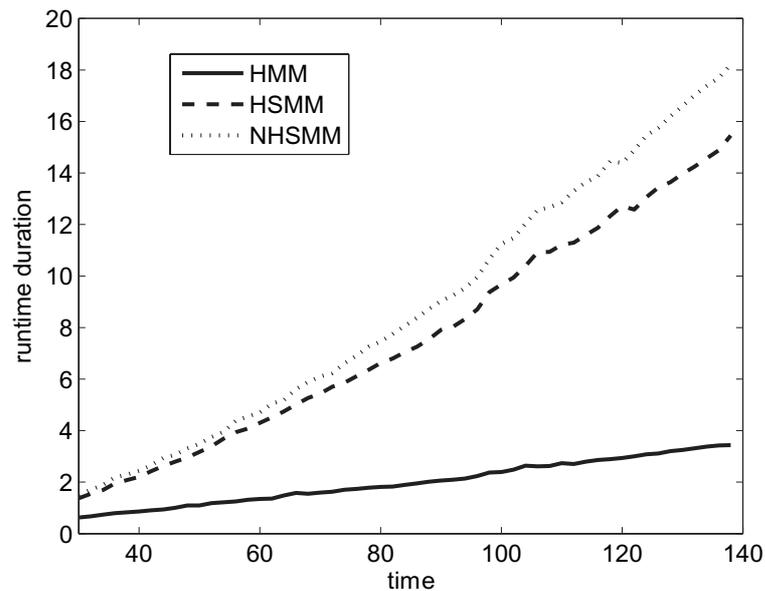


Fig. 7. Second experiment: A graph showing the run-time of each model (in seconds) over observation time-steps.

longer before turning to the opposite direction. For the HSMM graph, this split results in a drop in the probability values from a probability of almost 1, when both paths advance jointly, to almost 0.5 at the time the splitting occurs and remains constant around 0.5. These high probabilities indicate that HSMM regards the abnormal path as a path matching the training set and does not recognize the non-stationary characteristic of the tested data.

This is not the case in NSHSMM. The NSHSMM graph in Fig. 6 clearly shows an abrupt fall at time step 84 to very low values. This is due to the unexpected turn to the wrong direction. While the transition probability in NSHSMM predicted a turn to the right after spending a short time in the splitting state, the path followed a turn to the left, which resulted in a very low transition probability in the explicit state transition probability function. This experiment shows this model's ability to detect abnormal behavior in which the transition from one state to another does not match the duration probability in the prior state.

A run-time comparison for the second experiment is presented in Fig. 7. This figure depicts the resulting graphs of each model's run-time for calculating the likelihood function over observations. Again, there is a substantial increase in computation time when comparing HMM with both HSMM and NSHSMM. However, between HSMM and NSHSMM, the additional computation in NSHSMM appears minor. Note the similarity to Fig. 5.

## 5. Recognition in a cognitive task

We now turn to experiments using the different models in recognizing real-world human behavior, as part of a clinical diagnosis task. Our goal is to evaluate the utility of the different models to recognizing normal and abnormal behavior.

For this task, we apply the models to data gathered from human subjects, in two computerized maze navigation tasks. Maze tasks produce data that has many time-dependent patterns of human behavior

and movements. Being able to detect and classify such patterns may assist, for instance, in clinical diagnostic processes, by achieving better understanding of a patient's nature of impairment or providing help in assessing the patient's progress.

Maze tasks involve both lower and higher cognitive and executive functions as well as psycho-motor skills, and it offers a variety of research and intervention methods (e.g., [24]). Furthermore, since maze tasks are not verbal they may be administered to a wide range of healthy and impaired populations (e.g. [19]). Maze tests have been used, for example, to classify personality disorders [23], and for assessing patients' improvement when undergoing treatment [26]. Some of these behavioral patterns may be recognized by the trained eye according to simple classification rules [19]. Implementation of automated activity recognition methods, however, may identify more subtle patterns and help the clinician in the diagnostic process [6]. Similarly, analyzing patterns of behaviors in haptic mazes may help us, for example, in exploring the navigation strategies of people who are blind [14]. These insights are relevant, for example, for the design of haptic maps for the blind, where they learn and traverse novel environments, prior to approaching them in the real world [5].

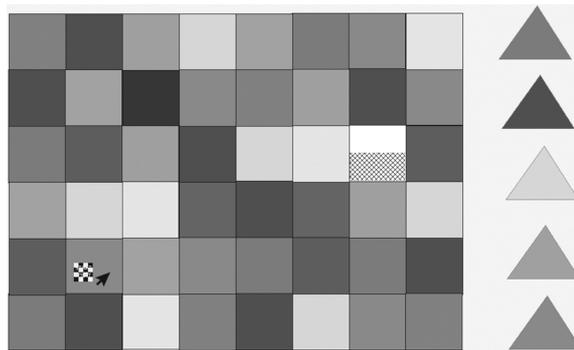
We used two types of maze tasks, each one requiring different faculties and skills, thus producing different types of data and allowing us to explore the abilities of the activity recognition models in two different settings [7]. We employed a computerized color maze (Section 5.1) and a computerized haptic maze (Section 5.2), where the participants control a cursor using a joystick. The mazes were implemented using the "Sense The Surface" software [7] which allows the researcher/clinician to design multi-sensory simulations. The interface used to navigate within the mazes was Microsoft's Sidewinder force-feedback joystick, an off-the-shelf product. During the trials, the systems sampled the cursor's position at 5 Hz and saved the  $X$ ,  $Y$  coordinates at each sample time  $t$  for offline analysis by our software.

We recruited 75 healthy volunteer subjects (48 subjects for the color maze experiments and 27 for the haptic maze experiments). They were 18–30 years of age, with right hand dominance and normal or corrected to normal vision. None of the participants suffered from any upper limb motor or sensory deficit, as well as diagnosed attention or cognitive disorders. The participants were university students from various departments. All of them had previous experience with computers but not necessarily with gaming interfaces. None, however, encountered difficulties in operating the system or completing the task. The sample was a convenience volunteer sample. The groups are of different sizes as we unfortunately had to discard some of the data of the haptic maze group due to technical reasons, resulting from variations in the hardware used to collect the data.

### 5.1. Task 1: Color maze

In the color maze task, participants were instructed to solve a maze by moving the cursor to the white square (the goal square in the maze). They were allowed to move the cursor only on squares of five predefined colors (See Fig. 8). The experiments were conducted individually. Each participant's movements were observed and manually separated into pattern groups (this can also be done with existing clustering methods). The software sampled the cursor's position and saved the data as described above.

Similar to the  $XY$ -grid in the simulated data, the scope of every square in the maze represents a state in the models. However, due to small movement errors, the cursor may move to adjacent cells when nearing a border, and thus observations of the cursor being in an adjacent cell can also be evidence of the original cell's state, i.e., border areas overlap, and states do have a hidden property. In particular, a border observation corresponds to all squares (i.e., states) sharing the border, each with some probability. The tessellation in the maze case is  $97 \times 97$  non-overlapping blocks, a size which offers a cognitive



(a) A screen-shot of a color maze. Subject has to move cursor to the white square. Cursor can be moved only on squares which are of five predefined colors (shown on right in triangles).



(b) Illustration of a subject solving the color maze.

Fig. 8. Color maze.

challenge compatible with non-computerized mazes. Here also, the data collected for these experiments is in the form of  $XY$ -positions following their time stamp. Out of the 75 volunteers, we recruited 48 healthy volunteer subjects for the color maze experiments (a typical number of subjects in this type of cognitive experiments). Each subject's trial in solving the maze is referred to as *a path*.

#### 5.1.1. Results.

Our interest was in finding patterns in the paths, building training sets according to the patterns found, and evaluating the models' ability to detect paths matching their own pattern group while recognizing their incompatibility with other pattern groups. We found several patterns within the collected data from the color maze trials, built training sets for each pattern accordingly and tested the models with paths from the collected data which are similar to the pattern groups we found (see Fig. 9 for examples of the patterns). Our experiments showed that HMM was not able to ascribe a path to its appropriate pattern group or detect its incompatibility with the other pattern groups. However, both HSMM and NSHSMM succeeded in these tasks. Hence, the following experiment results focuses only on HSMM and NSHSMM.

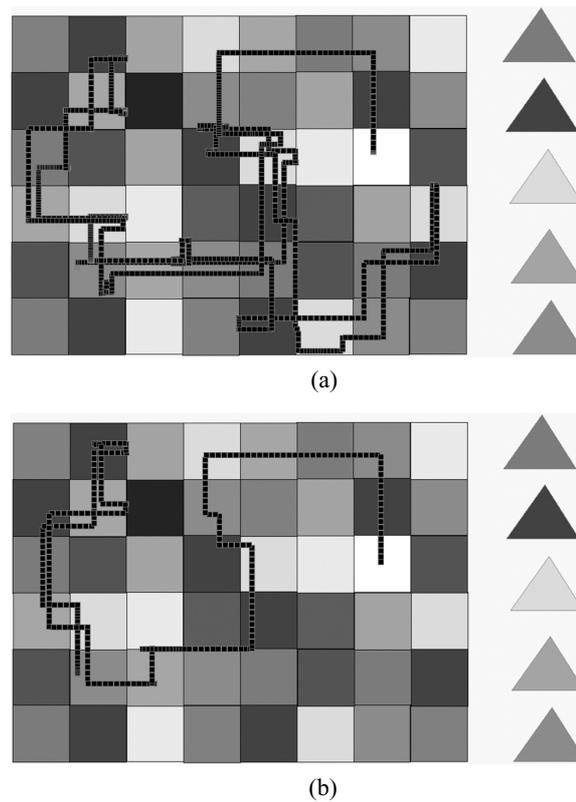


Fig. 9. Examples of patterns in the color maze.

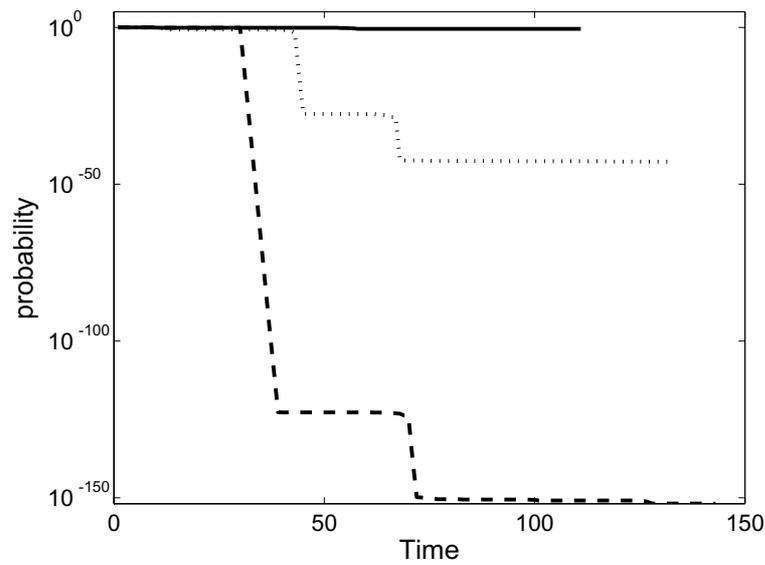
Our experiment included testing the recognition ability of each model in detecting paths matching their pattern group, paths of different patterns than those of the training set, and paths of the same pattern as the training set, but with different lengths (i.e., the length duration of the paths) than of the average length in the training set.

Figures 10(a) and 10(b) show the probability graphs for HSMM and NSHSMM. The upper graph in both models is the resulting recognition probability of a path similar to the tested pattern group, matching the training set. The two other graphs are the recognition probabilities of paths which are not consistent with the pattern group in the training set.

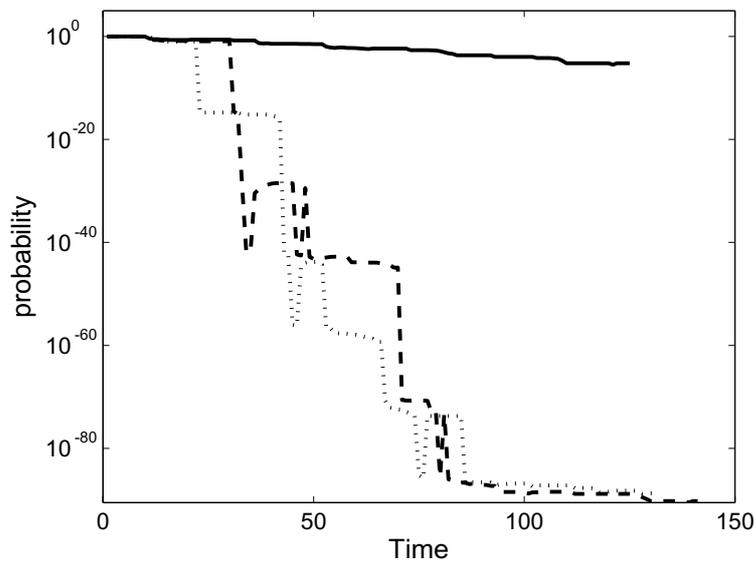
These results clearly show that both models succeeded in recognizing patterns as well as detecting incompatibilities with patterns. Yet, although both HSMM and NSHSMM recognized these patterns, the results typically show that the HSMM outperforms the NSHSMM. These results may be improved by using continuous probability density function, estimated through more examples. Furthermore, the fact that these results indicate that the NSHSMM model may need many more examples than HSMM to work well, is an issue of importance to practitioners.

## 5.2. Task 2: Haptic maze

A second maze task involved using a force-feedback joystick in order to provide haptic feedback to the subjects (i.e., information related to the sense of touch). In this task, the participants had to solve a maze without seeing it (see Fig. 11 for an example of a haptic maze). The only information they received



(a) HSM.

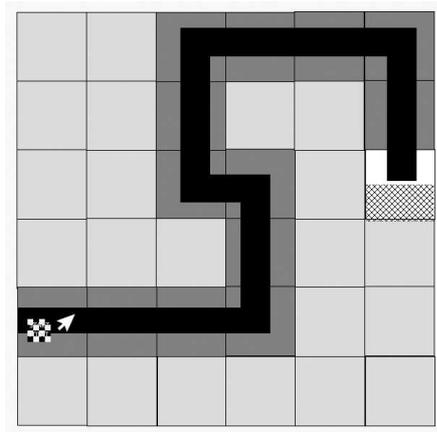


(b) NSHSM.

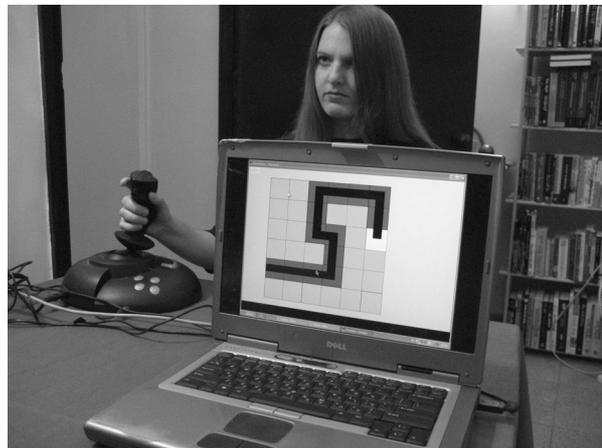
Fig. 10. Activity recognition results on clinical color maze-navigation task for HSM and NSHSM. The black line depicts the probability of a path similar to the training set. The dashed lines and the dotted lines depict the probability of paths which are not consistent with the patterns in the training set.

was a vibration delivered via the joystick. One type of vibration meant “correct route”, while another implied hitting a maze wall. Solving such a maze may require different types of strategy and employ different brain mechanisms.

Investigating the way haptic information is processed has implication for design of applications for people who are blind [5,14]. The ability to recognize behaviors that are distinctive to these populations are vital for appropriate design of such applications. It is important, for example, to detect the difference



(a) A screen-shot of a haptic maze. Subject has to move cursor to the white square without looking at the screen. Cursor can be moved only along the black path and receive force-feedback indicating whether the cursor is on the predefined path.

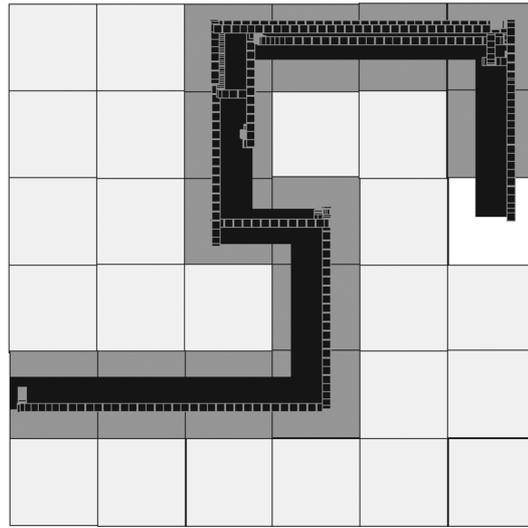


(b) Illustration of a subject solving the haptic maze. The subject is facing the back side of the screen and therefore can only rely on the force feedback from the joystick.

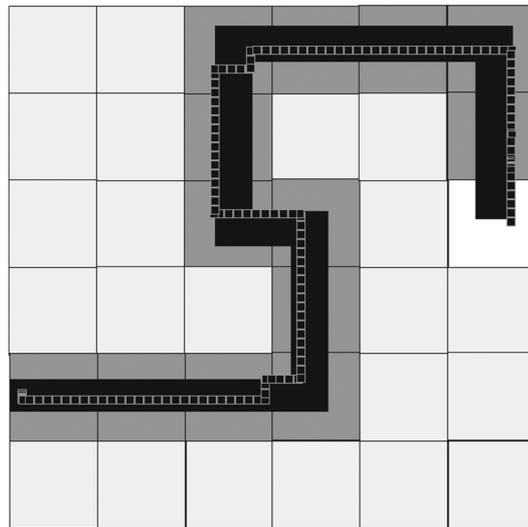
Fig. 11. Haptic maze.

between spatial processing performed by people who are congenitally blind as opposed to people who became blind in a later stage in life. Such distinctions are relevant as it is hypothesized that they navigate in different ways and employ different strategies to find their way. Additionally, such maze tasks may add to understanding the way the brain processes information provided by the haptic channel and may prove itself relevant for cognitive rehabilitation for post-stroke patients where haptic learning may replace damaged visual learning abilities (e.g. [3]).

The emphasis of the experiments with the haptic maze data is on the ability to capture the characteristic behavior from the tested paths in relation to particular behavior patterns drawn from the collected data. These behavior patterns portray certain characteristic behaviors that simulate the progress of blind people in an unfamiliar environment.



(a)



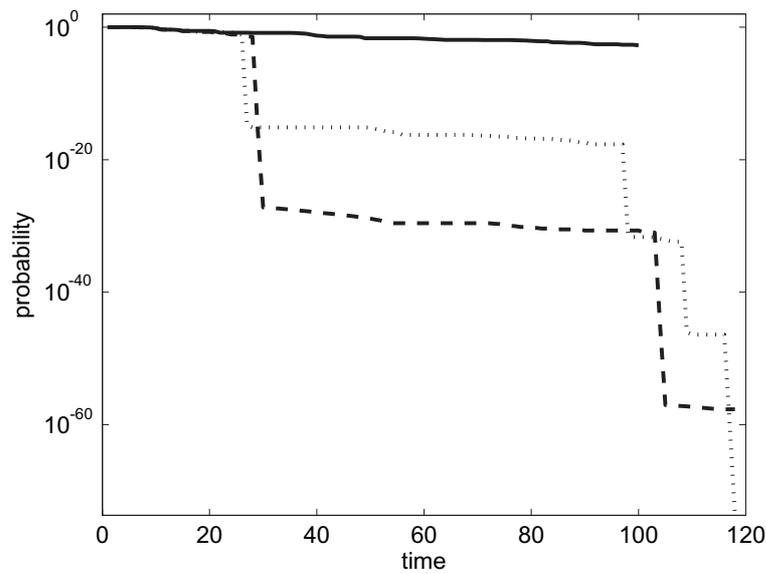
(b)

Fig. 12. Examples of patterns in the color maze.

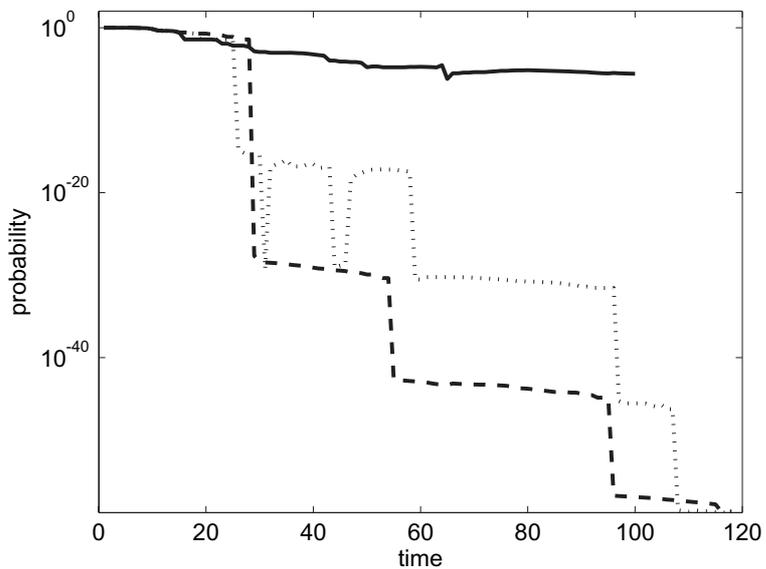
The tessellation in the haptic maze case is the same as that of the color maze. This tessellation was chosen for the similar reasons. Also, we later discuss some experiments where we compare patterns from both haptic maze and color maze data. The identical division of both mazes into blocks (i.e, states) allows us to make this comparison.

### 5.2.1. Results

Similarly to the color maze experiments, the data collected from the haptic maze trials was used to build training sets according to typical pattern groups found in the data. The goal of these experiments was to evaluate the models' ability to detect paths which resemble their pattern group as well as detect



(a) HSM.



(b) NSHSM.

Fig. 13. First experiment: Activity Recognition results on haptic maze for HSM and NSHSM. The black smooth graphs depicts the probability of a path similar to the training set. The dashed graphs and the dotted graphs depict the probability of paths which are not consistent with the patterns in the training set.

paths which are not consistent with a training set, i.e., belong to different pattern groups. Figure 12 shows examples of different patterns observed in the haptic experiments.

Our experiments showed that an HMM is inappropriate for this task, since it was not able to recognize a path belonging to its pattern group nor detect its incompatibility with the other pattern groups. This means that HMM failed cross validation. We therefore focus on the two models, HSM and NSHSM, and present the results of our experiments for these two models.

Out of the possible pattern groups obtained from the haptic maze data, we have chosen to present results for a pattern group that has characteristic blind-like behavior. By blind-like behavior we refer to this typically observed by people who are blind upon making a navigation error in a novel environment. It may include pausing in place, or getting stuck moving between two obstacles, until finding a way out. This pattern group contains paths in which the participants were observed to similarly linger at specific areas in the haptic maze which were more difficult to go through without viewing the maze (e.g., corners and turns). In these areas the participants received force-feedback from the joystick. However, since they were not allowed to look at the screen, they had to figure out how to extricate themselves without knowing their precise location in the maze or its solution. This resulted in creating paths with lingering states in certain areas on the haptic maze which are more impediment-prone. We have chosen to demonstrate the models' recognition abilities on a pattern group that consisted of paths with two principal lingering areas in the haptic maze.

The graphs in Fig. 13 show the resulting probabilities of HSMM and NSHSMM respectively. In both Fig. 13(a) and 13(b), the black smooth graphs represent the probability of a path that resembles the pattern group comprising the training set to belong to this group over a period, and the two remaining types of graphs, the dotted graph and the dashed graph represent the probability of two types of paths from other pattern groups to match the pattern in training set over a period. Since the solution to the haptic maze is the same as the color maze, the participants were in fact asked to solve the same graph. Therefore, in addition to testing the models' ability to detect paths from pattern groups taken from the haptic maze's pattern groups, we have tested these abilities in detecting paths taken from the color maze's pattern groups. The dashed graphs in Fig. 13(a) and 13(b) depicts a path from a pattern group taken from the collected data of color maze trails and the dotted graphs depicts a path from a different pattern group within the collected data from the haptic maze trails.

The graphs clearly show (the black smooth lines in Fig. 13(a) and 13(b) that both HSMM and NSHSMM were able to detect paths matching the training set. The graphs show that both models were able to ascribe the appropriate probabilities according to the given training set and are able to recognize a path matching the pattern of that training set's group. The two other graphs in Fig. 13(a) and 13(b) (dotted lines and dashed lines in Fig. 13(a) and 13(b) that represent the probability of paths of two different pattern groups from that of the training set to match it, show that both models, the HSMM and NSHSMM, were able to detect the abnormality and incomparability of the two tested patterns.

The graphs in Fig. 14 depict the probability of two paths that belong to pattern groups other than that of the training set, to resemble the training set's pattern group over a period. The black graphs depict the resulting probabilities of HSMM while the gray graphs depict the resulting probabilities of NSHSMM. The chosen paths in this experiment were picked from different pattern groups within the collected data from the haptic maze trails.

As can be seen in the graphs from Fig. 14, the probabilities of the black graphs remained high over a period while the two gray graphs show a drop in the probabilities in both graphs. This means, that the HSMM was not able to detect the abnormality of the two paths taken from pattern groups other than that of the training set's pattern group, while the NSHSMM succeeded in both of these tasks. The results in these experiments show that NSHSMM is superior to HSMM in recognizing abnormal behavior.

### 5.3. Discussion

In this section we applied our algorithms to real world data, gathered by behavioral experiments conducted with healthy volunteer subjects. The participants were asked to solve one of two types

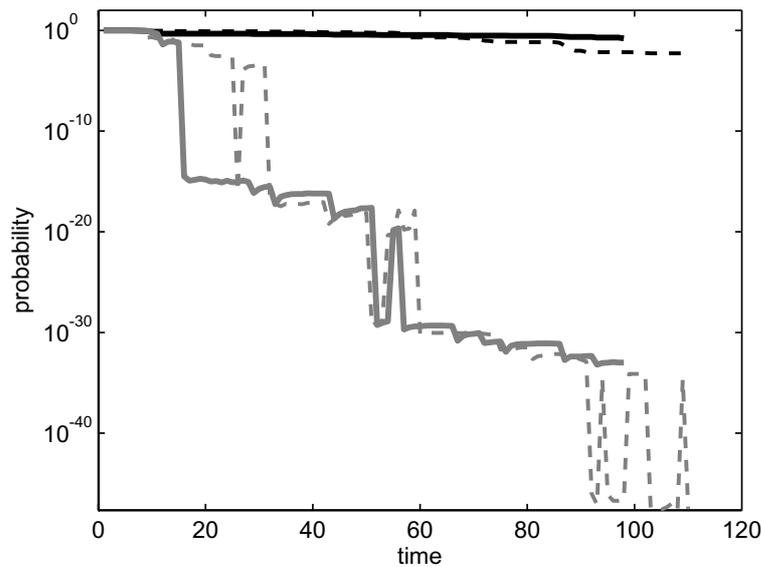


Fig. 14. Second experiment: Activity recognition results on haptic maze. The graphs show results of two paths not resembling the training set pattern. The black graphs depict the probabilities of HSM and the gray graphs depict the probabilities of NSHSMM.

of mazes, one color-based which required visual learning and one non-visual which required haptic learning. Maze tasks were chosen as they may be applied in clinical settings. Seventy five participants were recruited, divided into two groups and solved the mazes using an off-the-shelf force-feedback joystick and the specially designed “Sense the surface” software.

In both experiments, the color maze and the haptic maze, we were able to identify different patterns of cursor movement within the mazes. After training on such distinguishable patterns we tested the ability of the HMM, HSM and NSHSMM approaches to recognize unfamiliar patterns and classify them as belonging (or not) to the familiar patterns. This is done as one follows pattern activities over time and compares their probabilities of belonging to the training set pattern.

When this comparison was performed for the color maze it was evident that the HMM was not able to detect whether a pattern belonged or not to the training set. However, both the HSM and the NSHSMM were able to do so. As demonstrated in Fig. 10, the difference in likelihood between the similar pattern and the different pattern at  $t = 100$  was at least forty orders of magnitude, and typically much more.

In the next phase we applied the same method to the haptic maze data. There, too, the HMM failed in classification of patterns as similar or different than the training set. The HSM and NSHSMM showed again their ability to detect similarity as even in the worst case, after approximately 100 time steps, the probability difference between similar and different activity patterns was over ten orders of magnitude.

Being able to detect and classify patterns of movement may provide relevant information regarding the nature of a patient’s impairment or help to assess his or her progress. We believe that such data analyses may serve as a proof of concept, indicating the possible gains expected when applying to relevant clinical populations. Put together, it appears that the NSHSMM and HSM models and their potential of classification of subtle behaviors, may eventually serve as a valuable tool in the clinician’s toolbox.

## 6. Conclusions

The use of Markov model variants in activity recognition is quickly gaining popularity. Markov models provide relatively efficient and familiar methods for recognition. However, up until now, the use of such models in activity recognition has been hindered by lack of treatment of the effect of state durations on transition probabilities.

In this paper, we investigate hidden semi-Markov models in which state durations are explicitly modeled. We present an extension of semi-Markov models in activity recognition, allowing for transition probabilities conditioned on the duration of stay in a state: Non-stationary Hidden Semi-Markov Model (NSHSMM). These allow more accurate recognition of duration-based behavioral patterns.

We evaluate NSHSMMs in two application areas, using movement data of simulated airport passengers, and real world data from human subjects engaged in a complex maze navigation task. In both applications, we compare the performance of NSHSMMs to the results obtained with stationary HSMM and with HMM in recognition of normal and abnormal activities. The results elucidate the strengths and weaknesses of each method in these types of settings.

The evaluations we have conducted reveal that both semi Markov models are successful in recognition of behavioral patterns which are not recognized by the standard HMMs. Although HMMs and variants cannot handle specific conditions, this fact is actually disregarded by much previous works in activity recognition that utilize HMMs. In addition, there are important cases in which NSHSMMs outperform the HSMMs, and allow significantly more accurate recognition.

The results also point at an innovative application of clinical diagnosis, in which semi Markov models may be applied. Experiments with real world data show each of the two models, HSMMs and NSHSMMs, under certain conditions, have some advantages over the other method. These initial results suggest that semi Markov models may be useful in clinical settings for assessment of patients suffering from various deficits, and may assist in clinical practice. We hope that our analysis would help to shed light on the applicability of different models.

## Acknowledgements

This research was supported in part by IRST Trento-CRI Haifa, and by Israel Science Foundation Grants #1357/07 & #1242/04. We thank Martin Golumbic and the anonymous reviewers for useful comments. Thanks to K. Ushi for her support and patient proof-reading.

## References

- [1] H. H. Bui, S. Venkatesh and G. West, Policy recognition in the abstract hidden Markov model. *Journal of Artificial Intelligence Research* **17** (2002), 451–499.
- [2] E. Charniak and R. Goldman, A Bayesian model of plan recognition, *Artificial Intelligence Journal* (64) (1993), 53–79.
- [3] B.B. Connor, A.M. Wing, G.W. Humphreys, R.M. Bracewell and D.A. Harvey, Errorless learning using haptic guidance: research in cognitive rehabilitation following stroke, In *ICDVRAT 2004*, pages 77–83, 2002.
- [4] T.V. Duong, H.H. Bui, D.Q. Phung and S. Venkatesh, Activity recognition and abnormality detection with the switching hidden semi-Markov model, In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR05)*, 2005.
- [5] U. Feintuch, J. Haj and P.L. Weiss, Low-cost haptic-aural virtual environments for way-finding by children with congenital blindness: feasibility study. In *the 5th IEEE International Workshop on Virtual Rehabilitation*, pages 78–81, 2006.
- [6] U. Feintuch, L. Manevitz, E. Mednikovand, D. Rand, A. Dvorkin, R. Kizony, M. Shahar and P.L. Weiss, Integrating Artificial Intelligence and Virtual Reality in the diagnostic process – Feasibility study. In *the 11th Annual Cybertherapy Conference.*, 2006.

- [7] U. Feintuch, D. Rand, R. Kizony and P. Weiss, Development of a “low-end” multimodal feedback program for motor, cognitive and sensory rehabilitation, *Cyberpsychology & Behavior* **7**(3) (2004), 281–282.
- [8] J. Ferguson, Variable duration models for speech. In *Symposium on the Application of Hidden Markov Models to Text and Speech*, pages 143–179, 1980.
- [9] D. Fox, J. Hightower, L. Liao, D. Schulz and G. Borriello, Bayesian filtering for location estimation, *IEEE Pervasive Computing*, 2003.
- [10] K. Gopalratnam, H.A. Kautz and D.S. Weld, Extending continuous time bayesian networks, In *AAAI*, pages 981–986, 2005.
- [11] K. Han and M. Veloso, Automated robot behavior recognition applied to robotic soccer, In *IJCAI-99 Workshop on Team Behaviors and Plan Recognition.*, 1999.
- [12] S. Hongeng and R. Nevatia, Large-scale event detection using semi-HMMs, In *9th IEEE International Conference on Computer Vision (ICCV’03)*, 2003.
- [13] G.A. Kaminka, D.V. Pynadath and M. Tambe, Monitoring teams by overhearing: A multi-agent plan recognition approach, *Journal of Artificial Intelligence Research* **17** (2002).
- [14] O. Lahav and D. Mioduser, Multisensory virtual environment for supporting blind persons’ acquisition of spatial cognitive mapping, orientation, and mobility skills, In *ICDVRAT 2004*, pages 213–220, 2002.
- [15] S. Levinson, Continuously variable duration hidden Markov models for automatic speech recognition, *Computation Speech Language* (45) (1986), 1–29.
- [16] E. Marhasev, M. Hadad and G.A. Kaminka, Non-stationary hidden semi markov models in activity recognition, In *Proceedings of the AAAI Workshop on Modeling Others from Observations (MOO-06)*, 2006.
- [17] P. Natarajan and R. Nevatia, Coupled hidden semi markov models for activity recognition, In *IEEE Workshop on Motion and Video Computing*, 2007.
- [18] N. Oliver, E. Horvitz and A. Garg, Layered representation for human activity recognition, In *Fourth IEEE International Conference on Multimodal Interfaces (ICMI02)*, pages 3–8, 2002.
- [19] S.D. Porteus, *Porteus Maze Test, Fifty Years’ Application*, Pacific Books, Palo Alto, California, 1973.
- [20] D. Pynadath and M. Wellman, Probabilistic state-dependent grammars for plan recognition. In *the Conference on Uncertainty in Artificial Intelligence*, pages 507–514, 2000.
- [21] L.R. Rabiner, A tutorial on hidden Markov models and selected applications in speech recognition. In *IEEE*, volume 2, pages 257–286, 1989.
- [22] B. Sin and J. Kim, Nonstationary hidden Markov models. *Signal Processing* **46** (1995), 31–46.
- [23] M.C. Stevens, R.F. Kaplan and V.M. Hesselbrock, Executive cognitive functioning in the development of antisocial personality disorder, *Addictive Behaviors* **28** (2003), 285–300.
- [24] K.A. Sullivan and S.C. Bowden, Which tests do neurophysiologists use? *Journal of Clinical Psychology* **55** (1997), 657–661.
- [25] S.V. Vaseghi, State duration modelling in hidden Markov models, *Signal Processing* **41** (1995), 31–41.
- [26] J.H.G. Williams, J.T. O’Brien and S. Cullum, The time course of response to electroconvulsive therapy in elderly depressed subjects, *International Journal of Geriatric Psychiatry* **12** (1997), 563–566.
- [27] J. Yin, X. Chai and Q. Yang, High-level goal recognition in a wireless LAN, In *American Association for Artificial Intelligence*, 2004.
- [28] S.Z. Yu and H. Kobayashi, A hidden semi-Markov model with missing data and multiple observation sequences for mobility tracking, *Signal Processing* **83**(2) (2003), 235–250.