



Published in final edited form as:

J Ambient Intell Humaniz Comput. 2010 March 1; 1(1): 57–63. doi:10.1007/s12652-009-0007-1.

Recognizing independent and joint activities among multiple residents in smart environments

Geetika Singla, Diane J. Cook, and Maureen Schmitter-Edgecombe

School of Electrical Engineering and Computer Science, Washington State University, Box 642752, Pullman, WA 99164-2752, USA, cook@eecs.wsu.edu

Abstract

The pervasive sensing technologies found in smart homes offer unprecedented opportunities for providing health monitoring and assistance to individuals experiencing difficulties living independently at home. A primary challenge that needs to be tackled to meet this need is the ability to recognize and track functional activities that people perform in their own homes and everyday settings. In this paper, we look at approaches to perform real-time recognition of Activities of Daily Living. We enhance other related research efforts to develop approaches that are effective when activities are interrupted and interleaved. To evaluate the accuracy of our recognition algorithms we assess them using real data collected from participants performing activities in our on-campus smart apartment testbed.

Keywords

Smart environments; Passive sensors; Activity recognition; Multiple residents; Parallel activities

1 Introduction

A convergence of technologies in machine learning and pervasive computing has caused interest in the development of smart environments to emerge and assist with valuable functions such as remote health monitoring and intervention. An estimated 9% of adults age 65+ and 50% of adults age 85+ need assistance with everyday activities, and the resulting cost for governments and families is daunting. The resulting need for development of such technologies is underscored by the aging of the population, the cost of formal health care, and the importance that individuals place on remaining independent in their own homes. When surveyed about assistive technologies, family caregivers of Alzheimer's patients ranked activity identification, functional assessment, medication monitoring and tracking at the top of their list of needs (Rialle et al. 2008).

To function independently at home, individuals need to be able to complete both basic (e.g., eating, dressing) and more complex (e.g., food preparation, medication management, telephone use) Activities of Daily Living (ADLs) (Wadley et al. 2007). Smart environments can play an assistive role in this context. Our long-term goal is to design smart environment technologies that monitor the functional well-being of residents and provide assistance to help them live independent lives in their own homes. The goal of this project is to design an algorithm that labels the activity that an inhabitant is performing in a smart environment based on the sensor data that is collected by the environment during the activity. In the current study, our goal is

to design and test probabilistic modeling methods that can recognize activities that are performed by multiple residents in a single smart environment. To test our approach, we collect sensor data in our smart apartment testbed while participants perform activities. Some of these activities are performed by the two residents independently and in parallel, while other partial or full activities are performed in cooperation by the two residents. We use this collected data to assess the recognition accuracy of our algorithms.

There is a growing interest in designing smart environments that reason about residents (Cook and Das 2004; Doctor et al. 2005), provide health assistance (Mihailidis et al. 2004), and perform activity recognition (Philipose et al. 2004; Sanchez et al. 2008; Wren and Munguia-Tapia 2006). However, several challenges need to be addressed before smart environment technologies can be deployed for health monitoring. These include the design of activity recognition algorithms that generalize over multiple individuals and that operate robustly in real-world situations where activities are interrupted. This technology, if accurate, can be used to track activities that people perform in their everyday settings and to remotely and automatically assess their functional well-being.

2 Data collection

To validate our algorithms, we test them in a smart apartment testbed located on the WSU campus. The testbed is equipped with motion and temperature sensors as well as analog sensors that monitor water and stove burner use (see Fig. 1). The motion sensors are located on the ceiling approximately 1 m apart and are focused to provide 1 m location resolution for the resident. Voice over IP (VOIP) technology captures phone usage and we use contact switch sensors to monitor usage of the phone book, a cooking pot, and the medicine container. Sensor data is captured using a customized sensor network and stored in a SQL database.

To provide physical training data for our algorithms, we recruited 40 volunteer participants to perform a series of activities in the smart apartment. The smart space was occupied by two volunteers at the same time which performed the assigned tasks concurrently, making it a multi-resident environment. The collected sensor events were manually labeled with the activity ID and the person ID. The experimenter made careful notes on the timing of each step within the activity to ensure the accuracy of the labels.

For this study, we selected 15 ADLs that are often listed in clinical questionnaires (Reisberg et al. 2001). Observed difficulties in these areas can help to identify individuals who may be having trouble functioning independently at home (Schmitter-Edgecombe et al. 2008). These activities are listed below by activity name and by activity type (individual, individual requesting assistance, or cooperative).

Person A

1. *Filling medication dispenser (individual)*: For this task, the participant works at the kitchen counter to fill a medication dispenser with medicine stored in bottles.
2. *Moving furniture (cooperative)*: When Person A is requested for help by Person B, (s)he goes to the living room to assist Person B with moving furniture. The participant returns to the medication dispenser task after helping Person B.
3. *Watering plants (individual)*: The participant waters plants in the living room using the watering can located in the hallway closet.
4. *Playing checkers (cooperative)*: The participant brings a checkers game to the dining table and plays the game with Person B.

5. *Preparing dinner (individual)*: The participant sets out ingredients for dinner on the kitchen counter using the ingredients located in the kitchen cupboard.
6. *Reading magazine (individual)*: The participant reads a magazine while sitting in the living room. When Person B asks for help, Person A goes to Person B to help locate and dial a phone number. After helping Person B, Person A returns to the living room and continues reading.
7. *Gathering and packing picnic food (individual)*: The participant gathers five appropriate items from the kitchen cupboard and packs them in a picnic basket. (S)he helps Person B to find dishes when asked for help. After the packing is done, the participant brings the picnic basket to the front door.

Person B

1. *Hanging up clothes (individual)*: The participant hangs up clothes that are laid out on the living room couch, using the closet located in the hallway.
2. *Moving furniture (cooperative)*: The participant moves the couch to the other side of the living room. (S)he requests help from Person A in moving the couch. The person then (with or without the help of Person A) moves the coffee table to the other side of the living room as well.
3. *Reading magazine (individual)*: The participant sits on the couch and reads the magazine located on the coffee table.
4. *Sweeping floor (individual)*: The participant fetches the broom and the dust pan from the kitchen closet and uses them to sweep the kitchen floor.
5. *Playing checkers (cooperative)*: The participant joins Person A in playing checkers at the dining room table.
6. *Setting the table (individual)*: The participant sets the dining room table using dishes located in the kitchen cabinet.
7. *Paying bills (cooperative)*: The participant retrieves a check, pen, and envelope from the cabinet under the television. (S)he then tries to look up a number for a utility company in the phone book but later asks Person A for help in finding and dialing the number. After being helped, the participant listens to the recording to find out a bill balance and address for the company. (S)he fills out a check to pay the bill, puts the check in the envelope, addresses the envelope accordingly and places it in the outgoing mail slot.
8. *Gathering and packing picnic supplies (cooperative)*: The participant retrieves a Frisbee and picnic basket from the hallway closet and dishes from the kitchen cabinet and then packs the picnic basket with these items. The participant requests help from Person A to locate the dishes to pack.

As can be noted from the activity descriptions, for some of the activities (Moving furniture, playing checkers, paying bills, and packing picnic supplies) the residents cooperate to jointly accomplish the task. The remaining activities are performed independently and in parallel. Figure 2 shows two participant residents, Person A and Person B, performing two of these tasks (Filling medication dispenser and Hanging up clothes) and indicates the sensor readings that are generated during these activities. The average activity times and number of sensor events generated for each activity are shown in Table 1. We anticipate that the shorter events (e.g., Person A/Activity 2—Moving furniture, Person B/Activity 2, Reading magazine) will be more difficult to recognize using any technique.

Our goal from this study is to identify the activities that are performed in a smart environment. While activity recognition algorithms have been explored in the literature (Cook and Schmitter-Edgecombe 2009; Liao et al. 2005; Maurer et al. 2006; Philipose et al. 2004), they typically focus on one resident performing a single activity at a time. Recognizing activities in multi-resident settings is much more challenging because the sensor readings for an activity are not contiguous. Instead, they jump back and forth between the activities that the residents are performing in parallel.

3 Probabilistic models

Because activities are performed in a real-world environment, there exists a great deal of variation in the manner in which the activity is performed. This variation is increased dramatically when the model used to recognize the activity needs to generalize over more than one possible resident. Researchers (Maurer et al. 2006; Liao et al. 2005) frequently exploit probabilistic models to recognize activities and detect anomalies in support of individuals living at home with special needs.

We investigate the use of probabilistic models for this task because of its ability to represent and reason about variations in the way an activity may be performed. Specifically, we design a hidden Markov model to determine an activity that most likely corresponds to an observed sequence of sensor events. A Markov Model (MM) is a statistical model of a dynamic system. A MM models the system using a finite set of states, each of which is associated with a multidimensional probability distribution over a set of parameters. The parameters for the model are the possible sensor event values described above. The system is assumed to be a Markov process, so the current state depends on a finite history of previous states (in our case, the current state depends only on the previous state). Transitions between states are governed by transition probabilities.

For any given state a set of observations can be generated according to the associated probability distribution. We could generate one Markov model for each activity that we are learning (Cook and Schmitter-Edgecombe 2009). However, this approach would ignore the relationship between different activities performed in parallel as well as sequentially. In addition, a Markov chain would not always be appropriate because we will not be able to separate the sensor event sequence into non-overlapping subsequences that correspond to each individual activity.

A hidden Markov model (HMM) is a statistical model in which the underlying model is a stochastic process that is not observable (i.e., hidden) and is assumed to be a Markov process which can be observed through another set of stochastic processes that produce the sequence of observed symbols. A HMM assigns probability values over a potentially infinite number of sequences. Because the probabilities values must sum to one, the distribution described by the HMM is constrained. This means that the increase in probability values of one sequence is directly related to the decrease in probability values for another sequence.

In the case of a Markov chain, all states are observable states and are directly visible to the observer. Thus, the only other parameter in addition to the prior probabilities of the states and the distribution of feature values for each state is the state transition probabilities. In the case of a hidden Markov model, there are hidden states which are not directly visible, and the observable states (or the variables) influence the hidden states. Each state is associated with a probability distribution over the possible output tokens. Transitions from any one state to another are governed by transition probabilities as in the Markov chain. Thus, in a particular state an outcome can be generated according to the associated probability distribution.

HMMs are known to perform very well in cases where temporal patterns need to be recognized which aligns with our requirement in recognizing interleaved activities. The conditional probability distribution of any hidden state depends only on the value of the preceding hidden state. The value of an observable state depends only on the value of the current hidden state. The observable variable at time t , namely x_t , depends only on the hidden variable y_t at that time. We can specify an HMM using three probability distributions: the distribution over initial states $\Pi = \{\pi_k\}$, the state transition probability distribution $A = \{a_{kl}\}$, with $a_{kl} = p(y_t = l | y_{t-1} = k)$ representing the probability of transitioning from state k to state l ; and the observation distribution $B = \{b_{il}\}$, with $b_{il} = p(x_t = i | y_t = l)$ indicating the probability that the state l would generate observation $x_t = i$. These distributions are estimated based on the relative frequencies of visited states and state transitions observed in the training data.

Given a set of training data our algorithm uses the sensor values as parameters of the hidden Markov model. Given an input sequence of sensor event observations, our goal is to find the most likely sequence of hidden states, or activities, which could have generated the observed event sequence. We use the Viterbi algorithm (Viterbi 1967) to identify this sequence of hidden states following the calculation in Eq. 1

$$\arg \max_{x_1, \dots, x_t} P(y_1, \dots, y_t, y_{t+1} | x_{1:t+1}). \quad (1)$$

In our implementation of a hidden Markov model, we treat every activity as a hidden state. As a result, our HMM includes 15 hidden states, each of which denotes one of the 15 modeled activities. Next, every sensor is treated as an observable state in the model due to the fact that every sensor which is used is observable in our dataset. Figure 3 shows a portion of the generated HMM for the multi-resident activities.

The challenge here is to identify the sequence of activities (i.e., the sequence of visited hidden states) that corresponds to a sequence of sensor events (i.e., the observable states). For this, we calculate based on the collected data the prior probability (i.e., the start probability) of every state which represents the belief about which state the HMM is in when the first sensor event is seen. For a state (activity) a , this is calculated as the ratio of instances for which the activity label is a . We also calculate the transition probability which represents the change of the state in the underlying Markov model. For any two states a and b , the probability of transitioning from state a to state b is calculated as the ratio of instances having activity label a followed by activity label b , to the total number of instances. The transition probability signifies the likelihood of transitioning from a given state to any other state in the model and captures the temporal relationship between the states. Lastly, the emission probability represents the likelihood of observing a particular sensor event for a given activity. This is calculated by finding the frequency of every sensor event as observed for each activity.

As an initial evaluation of our hidden Markov model, we put all of the sensor data for the 15 activities into one dataset and evaluate the labeling accuracy of the HMM using threefold cross validation. The HMM recognizes both the person and the activity with an average accuracy of 60.60%. This accuracy is higher than the expected random-guess accuracy of 7.00%, but still needs to be improved before the smart environment can reliably track and react to daily activities that the residents perform.

Figure 4 shows the accuracy of the HMM broken down by activity. As the figure shows, the activities which take more time and generate more sensor events (e.g., Read magazine A, 94.38% accuracy) tend to be recognized with greater accuracy. The activities which are very

quick (e.g., Set table B, 21.21% accuracy) do not generate enough sensor events to be distinguished from other activities and thus yield lower recognition results.

4 Separating models for residents

In our HMM implementation, a single model is implemented for both residents. The model thus represents transitions between activities performed by one person. However, it also represents transitions between residents and transitions between different activities performed by different residents. In some cases the residents are performing unrelated activities. As a result, there is not a strong underlying pattern for the model to use.

We hypothesize that multi-resident activities can be better recognized when one model is learned for each resident. To validate our hypothesis we generate one HMM for each resident. Each of the models contains one hidden node for each activity and observable nodes for the sensor values. In this case, we assume that we know the person ID for each sensor event. However, the sensor data is still collected from the combine multiple-resident testbed where the residents are performing activities in parallel. The average accuracy of the new model is 73.15%. The resulting accuracy for each activity is shown in Fig. 5.

The goal of this project was to design an algorithmic approach to recognize activities performed in a real-time, complex, smart environment. Our experimental results indicate that it is possible to distinguish between activities that are performed in a smart home even when multiple residents are present.

In all of the experiments, the accuracy level varied widely by activity. This highlights the fact that smart environment algorithms need to not only perform automated activity recognition, but they need to base subsequent responses on the recognition accuracy that is expected for a particular activity. If the home intends to report to the caregiver changes in the performance activity, such changes should only be reported for activities where recognizing and assessing the completeness of the activity can be accomplished with consistent success.

5 Conclusions

While our study revealed that Markov models are effective tools for recognizing activities, there are even more complex monitoring situations that we need to consider before the technology can be deployed. In particular, we need to design algorithms that perform accurate activity recognition and tracking for environments that house multiple residents. In addition, we need to design method for detecting errors in activity performance and for determining the criticality of detected errors.

In our data collection, an experimenter informed the participants of each activity to perform. In more realistic settings, such labeled training data will not be readily available and we need to research effective mechanisms for training our models without relying upon excessive input from the user. We hypothesize that ADL recognition and assessment can be performed in such situations and our future studies will evaluate the ADL recognition and assessment algorithms in actual homes of volunteer participants.

Closely related to the problem of activity recognition for multiple residents is the challenge of recognizing activities when they are interleaved. Our earlier study (Singla et al. 2009) provides evidence that HMMs can be used to recognize interleaved activities. However, we have not yet considered situations in which multiple residents are performing activities separately, interleaved, in parallel, and in cooperation. Our future work can consider these situations as well.

In this work, we described an approach to recognizing activities performed by smart home residents. In particular, we designed and assessed algorithms that built probabilistic models of activities and used them to recognize activities in complex situations where multiple residents are performing activities in parallel in the same environment. Ultimately, we want to use our algorithm design as a component of a complete system that performs functional assessment of adults in their everyday environments. This type of automated assessment also provides a mechanism for evaluating the effectiveness of alternative health interventions. We believe these activity profiling technologies are valuable for providing automated health monitoring and assistance in an individual's everyday environments.

References

- Cook, DJ.; Das, SK., editors. Smart environments: technology, protocols, and applications. 1st. Wiley; New York: 2004.
- Cook DJ, Schmitter-Edgecombe M. Assessing the quality of activities in a smart environment. *Methods Inf Med*. 2009 in press.
- Doctor F, Hagrais H, Callaghan V. A fuzzy embedded agent-based approach for realizing ambient intelligence in intelligent inhabited environments. *IEEE Trans Syst Man Cybern Part A* 2005;35(1): 55–56.
- Liao, L.; Fox, D.; Kautz, H. Location-based activity recognition using relational Markov networks. *Proceedings of the international joint conference on artificial intelligence*; 2005. p. 773-778.
- Maurer, U.; Smailagic, A.; Siewiorek, D.; Deisher, M. Activity recognition and monitoring using multiple sensors on different body positions. *Proceedings of the international workshop on wearable and implantable body sensor networks*; 2006. p. 113-116.
- Mihailidis A, Barbenl J, Fernie G. The efficacy of an intelligent cognitive orthosis to facilitate handwashing by persons with moderate-to-severe dementia. *Neuropsychol Rehabil* 2004;14(1/2):135–171.
- Philipose M, Fishkin K, Perkowitz M, Patterson D, Fox D, Kautz H, Hahnel D. Inferring activities from interactions with objects. *IEEE Pervasive Comput* 2004;3(4):50–57.
- Reisberg B, Finkel S, Overall J, Schmidt-Gollas N, Kanowski S, Lehfeld H, et al. The Alzheimer's disease activities of daily living international scale (ASL-IS). *Int Psychogeriatr* 2001;13:163–181. [PubMed: 11495392]
- Rialle V, Ollivet C, Guigui C, Herve C. What do family caregivers of Alzheimer's disease patients desire in smart home technologies? *Methods Inf Med* 2008;47:63–69. [PubMed: 18213430]
- Sanchez D, Tentori M, Favela J. Activity recognition for the smart hospital. *IEEE Intell Syst* 2008;23(2): 50–57.
- Schmitter-Edgecombe, M.; Woo, E.; Greeley, D. Memory deficits, everyday functioning, and mild cognitive impairment. *Proceedings of the annual rehabilitation psychology conference*; 2008.
- Singla G, Cook D, Schmitter-Edgecombe M. Tracking activities in complex settings using smart environment technologies. *Int J BioSci Psychiatr Technol* 2009;1(1):25–35.
- Viterbi A. Error bounds for convolutional codes and an asymptotically-optimum decoding algorithm. *IEEE Trans Inf Theory* 1967;13(2):260–269.
- Wadley V, Okonkwo O, Crowe M, Ross-Meadows LA. Mild Cognitive Impairment and everyday function: Evidence of reduced speed in performing instrumental activities of daily living. *Am J Geriatr Psychiatry* 2007;16(5):416–424. [PubMed: 18448852]
- Wren, C.; Munguia-Tapia, E. Toward scalable activity recognition for sensor networks. *Proceedings of the workshop on location and context-awareness*; 2006. p. 168-185.

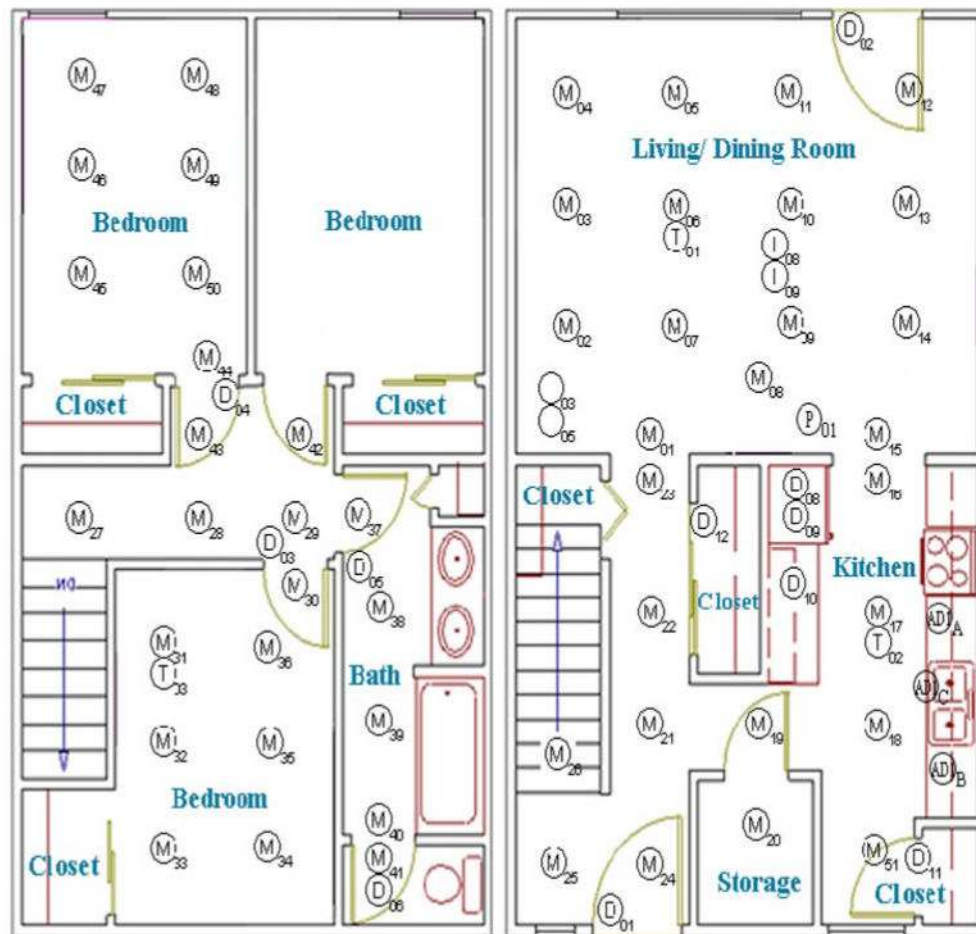


Fig. 1. The smart apartment testbed. Sensors in the apartment (*bottom*) monitor motion (*M*), temperature (*T*), water (*W*), door (*D*), burner (*AD*), and item use (*I*)



Fig. 2. Multi-resident participants performing activities in smart apartment. Person A (*upper left*) is performing “Fill medication dispenser” activity while Person B (*upper right*) is performing “Hang up clothes” activity. The sensor events corresponding to these activities is shown in the *lower left* and is visualized in the *lower right*

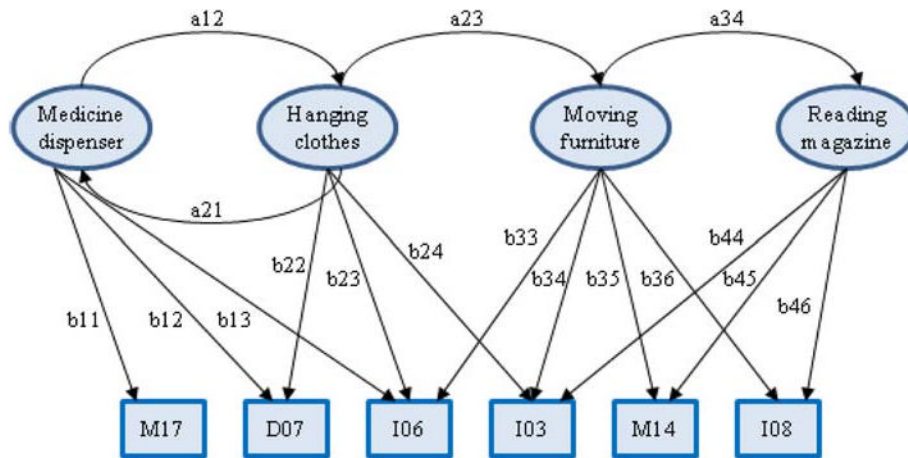


Fig. 3.

A section of an HMM for multi-resident activity data. The *circles* represent hidden states (i.e., activities) and the *rectangles* represent observable states. *Values on horizontal edges* represent transition probabilities and *values on vertical edges* represent the emission probability of the corresponding observed state

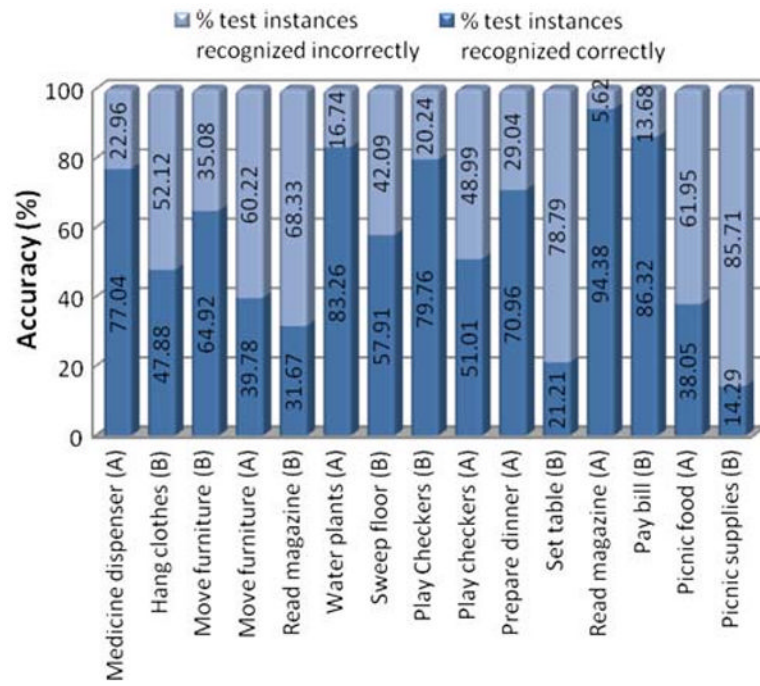


Fig. 4. Performance of a HMM in recognizing activities for multi-resident data




Fig. 5.
Performance of one HMM for each resident

Table 1

Average time (in minutes) and number of sensor events generated for each activity

Activity	Person A time	Person A #events	Person B time	Person B #events
1	3.0	47	1.5	55
2	0.7	33	0.5	23
3	2.5	61	1.0	18
4	3.5	38	2.0	72
5	1.5	41	2.0	25
6	4.5	64	1.0	32
7	1.5	37	5.0	65
8	N/A	N/A	3.0	38