Incorporating Temporal Reasoning into Activity Recognition for Smart Home Residents

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

Washington State University
Pullman, WA 99164-2752
{gsingla, djcook, schmitter-e}@wsu.edu

Abstract

Smart environments rely on artificial intelligence techniques to make sense of the sensor data that is collected in the environment and to use the information for data analysis, prediction, and event automation. In this paper we discuss an important smart environment technology – resident activity recognition. This technology is beneficial for health monitoring of a smart environment resident but accurate recognition is difficult for real-world situations. We describe our approach to activity recognition and discuss how incorporating temporal reasoning improves the accuracy of our algorithms. We validate our algorithm on real sensor data collecting in our smart apartment testbed.

Introduction

A recent convergence of technologies in machine learning and pervasive computing has caused interest in the development of *smart environments* to emerge. In addition to providing an interesting platform for developing adaptive and functional software application, smart environments can also be employed for valuable functions such as at-home health monitoring. In this project, we are using smart environments to address the health-care problem of performing automated assessment of functional health for elder adults and to provide automated assistance that will allow these individuals to remain independent in their own homes. Our target population is older adults and individuals with physical or cognitive disabilities. By 2040, 23% of the US population will be 65+ (Lanspery et al. 1997) and over 11 million people will suffer from dementia related to Alzheimer's disease (Herbert et al. 2000). To most people home is a sanctuary, yet today those who need special care often must leave home to meet clinical needs. This problem hits all ages, but is especially relevant for the quickly-growing elderly segment of the population, who are seriously affected by leaving a familiar environment: 9 out of 10 Americans want to live out their lives in familiar surroundings (Gross 2007). Given the cost of nursing home care and the importance that Americans place on remaining independent in their homes for as long as possible (AARP 2003), it is not surprising that the AARP strongly encourages increased funding for home modifications that keep older adults independent in their own homes.

In order to function independently at home, adults need to be able to complete key Activities of Daily Living, or ADLs (Diehl et al., 2005). Tracking ADL accomplishment is a time consuming task for caregivers, yet is required for formal care settings such as nursing homes. ADLs fall within categories such as medication use, telephone use, financial management, personal hygiene, hydration, and food consumption. When surveyed about the assistive technologies they desire most, family caregivers of Alzheimer's disease patients ranked activity identification, functional assessment, medication monitoring, and tracking at the top of their list of concerns (Rialle et al. 2008). The outcome of this program will be technological tools that monitor the completion of these tasks.

We hypothesize that many older adults with cognitive and physical disabilities can lead independent lives in their own homes with the aid of at-home automated assistance and health monitoring. As a first step, we want to determine how effectively ADLs can be recognized in real-world settings. Models of daily activities can be learned from sensor events collected by a smart environment. By learning models for each task, ADL initiation and completion can be automatically detected, even when the activity is incomplete, the resident is switching between tasks, and when additional people are in the environment. If we can successfully detect ADL initiation and completion in these settings then we can use the technologies to perform automatic assessment of an individual's well being and provide the foundation for reminder-based interventions.

In this paper we describe a method of learning models of activity behavior using Markov models. While this method is effective at distinguishing between simple tasks, handling real-world task recognition is more challenging. To make the models more robust we add temporal information to the models. We demonstrate that this temporal information improves activity recognition performance on real-world tasks as performed in our smart home testbed.

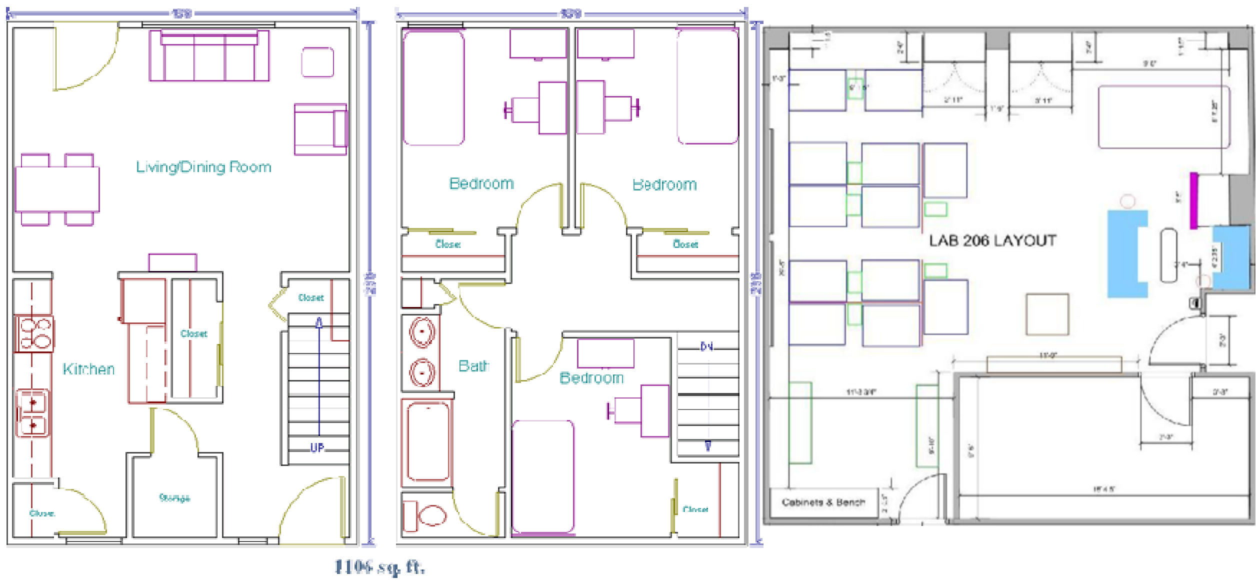

Figure 1. Two of the CASAS testbeds: a three-bedroom smart apartment (left) and a lab workplace environment (right).

Related Work

Although there is a growing interest in adding intelligence to our living and working environments, only recently has the convergence of technologies in machine learning, pervasive computing, and sensor networks made the idea of smart environments a reality. Researchers have generated ideas for designing smart environment software algorithms that track the location and activities of residents, that generate reminders, and that react to hazardous situations (Wren and Munguia-Tapia 2006).

One limiting factor of these projects is that almost none are being tested on data collected from physical environments. A few testbeds do exist in some form, although none are currently focusing on research for automated functional assessment and intervention. These include our earlier MavHome project (Youngblood and Cook 2007), the Gator Tech Smart House (Helal et al. 2005), the iDorm (Doctor et al. 2005), the DOMUS lab (Pigot et al. 2002), the Georgia Tech Aware Home (Abowd and Mynatt 2004), and the University of Colorado Adaptive Home (Mozer 2004). As a result of this and related work, researchers are now beginning to recognize the importance of applying smart environment technology to health assistance (Barger et al. 2005; Kautz et al. 2002; Larson 2007; Mihailidis e t al. 2004; Pollack 2005) and companies are recognizing the potential of this technology for a quickly-growing consumer base (Intel 2007).

CASAS

The activity recognition algorithms we introduce in this paper are part of the CASAS smart environment software architecture. In order to evaluate our algorithms, we test them using data collected from volunteer participants per-

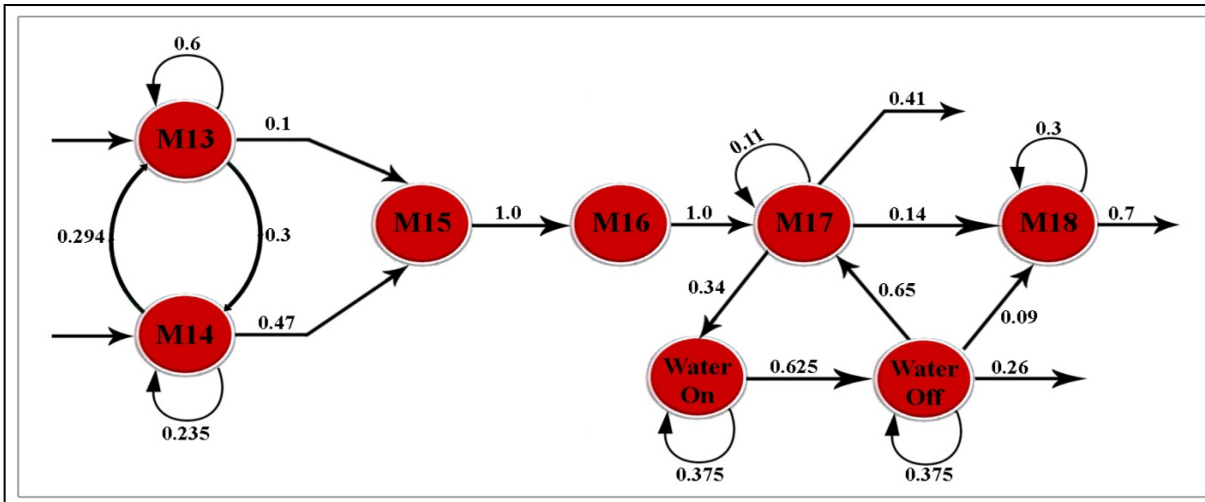

Figure 2. Markov model representing the “wash hands” activity.

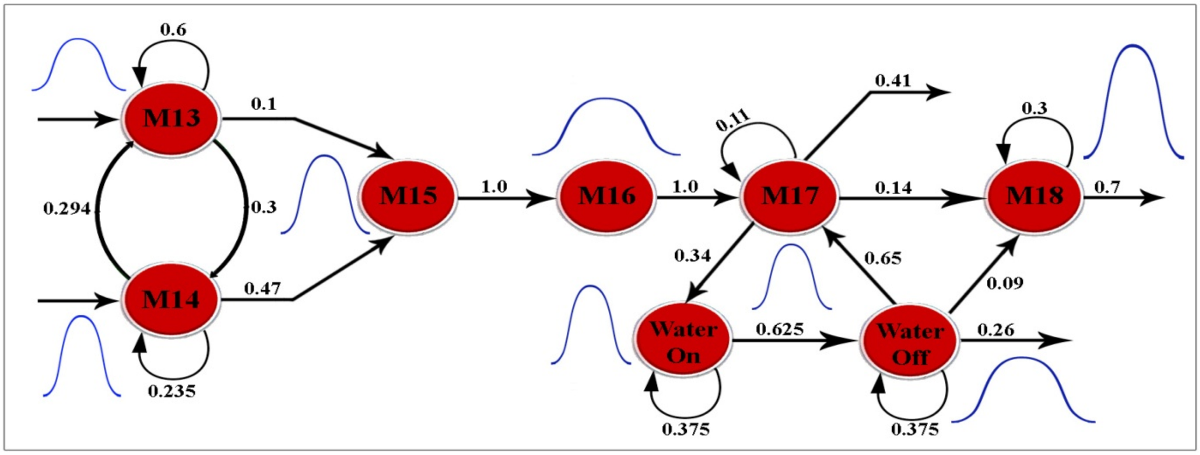

Figure 3. Markov model with each state annotated with normal distribution over time.

forming activities in our smart environment testbeds. The first physical testbed is a smart apartment on the WSU campus. The apartment includes three bedrooms, one bathroom, a kitchen, and a living/dining room. The layout of the apartment is shown in Figure 1 (left). In addition to a home environment, we have equipped the workplace environment shown in Figure 1 (right). This WSU lab contains four work areas, a common compute/print area, a kitchen, a lounge, and a meeting area. These environments are equipped with motion sensors, temperature sensors, humidity sensors, contacts switches in the doors, and item sensors on key items. We have designed special-purpose sensors to detect water usage and stove burner usage and use the Asterisk software to monitor outgoing phone usage. All of these sensors have the advantage of being non-obtrusive and relatively easy to monitor remotely.

Activity Recognition

While collecting sequences of sensor readings in a smart environment is valuable, determining what activities these sequences represent is a more challenging task. We are employing Markov models to automatically recognize the initiation and completion of ADLs. Figure 3 shows an example Markov model for this task. Researchers have investigated the recognition of resident activities using a variety of mechanisms such as naïve Bayes classifiers, Markov models, and dynamic Bayes networks (Liao et al. 2005; Philipose et al. 2004; Wren and Munguia-Tapia 2006). These approaches have been fairly effective for limited types of tasks. However, these limitations prevent the techniques from being applied in general situations. In particular, all of the approaches to date make the assumptions that the activity needs to be identified only when it is complete, that the activity is performed in a consistent manner without missing steps, that timing information is not essential for identifying activities, that activities are performed independently, and that there is only one resident to monitor at a time.

In contrast, we are designing algorithms that probabilistically identify the activity while it is being performed, as well as identify steps that are missing. The states of the Markov model can represent sensor events as shown in Figure 4 or abstractions of sensor information. With the number of tasks known ahead of time and available training data, our algorithm will construct a hidden Markov model for each task and learn the probabilistic transitions between states. Given a sequence of observed sensor events we can probabilistically determine which model best supports the sequence.

For example, the probability that the “wash hands” model in Figure 2 supports the sequence [Motion 14, Motion 15, Motion 16, Motion 17, Water On, Water Off] is 0.54, which is greater than the values from the models for cooking, making a phone call, or cleaning up. Similarly, we can probabilistically determine the belief state, or the most likely state of the model that is currently being observed. Using this information we can track the individual steps that comprise an activity and use this information to assess whether the activity was performed completely or left unfinished and whether particular critical steps were skipped.

Although a number of methods have been suggested for activity recognition, these methods do not incorporate state recognition and missing step detection. In addition, few approaches are designed to make use of the timing of the activity and steps within the activity. Duration of each subtask in an activity can be used as additional information to distinguish between overlapping activities i.e. different activities which trigger a similar set of sensors. In order to incorporate this timing information, our models annotate each state description with a normal distribution

sensor ID	date / time	reading
12048146000000B2	2008-02-12 10:50:45.673225	ON
12D27E460000000D	2008-02-12 10:50:48.903745	ON
12048146000000B2	2008-02-12 10:50:49.339849	OFF
2084A30D00000039B	2008-02-12 10:50:53.27364	0.0459382
2084A30D00000039B	2008-02-12 10:51:05.6252	0.158401

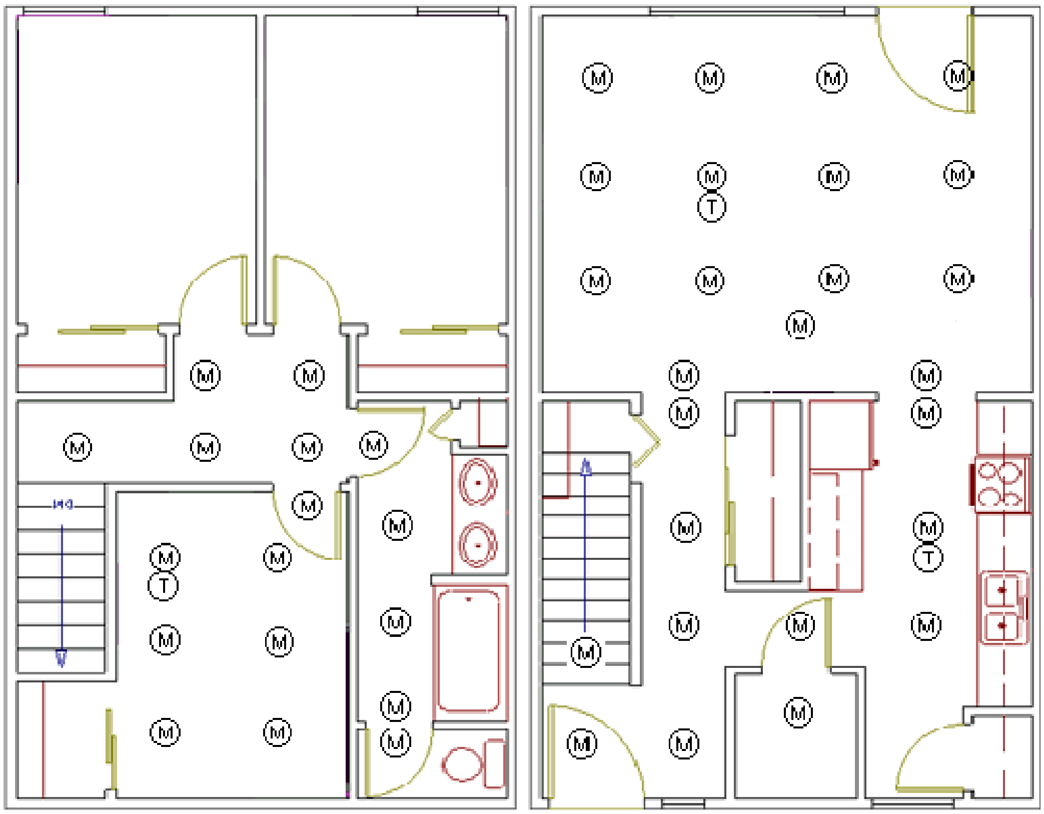

Figure 4. Recognition of the "wash hands" activity in the smart apartment. The web cam image in the upper left shows a student participant performing the task. The activity results in the sensor readings shown in the upper right (the first three readings correspond to motion sensors and the last two correspond to non-zero water flow values). A visualization of the sensor activity for the "wash hands" task is shown at the bottom.

representing the likely start time and duration of the activity initiation and of each step comprising the activity. We calculate the probability of time matches using the definition of the normal distribution and now this value also contributes to the probability of a model matching an activity in addition to the probability based on the sensor events.

Our overall objective is to design software algorithms that will monitor the overall functional wellbeing of individuals at home by detecting ADLs that are being performed by residents in a smart environment. We will test our working hypothesis that smart environment-based measurement techniques can accurately detect completed ADLs. We will ultimately us this capability to identify the current step the individual is performing within an ADL and determine which steps of the ADL were skipped or performed out of order.

sensor ID	date / time	reading
12C4395F000000F7	2008-02-12 10:53:49.31232	PRESENT
12CA7E46000000F7	2008-02-12 10:53:51.332601	CLOSE
12D27E460000000D	2008-02-12 10:53:54.815838	ON
124F7C4600000075	2008-02-12 10:54:55.23247	OFF
2084A30D00000039A	2008-02-12 10:53:56.21309	2.81481

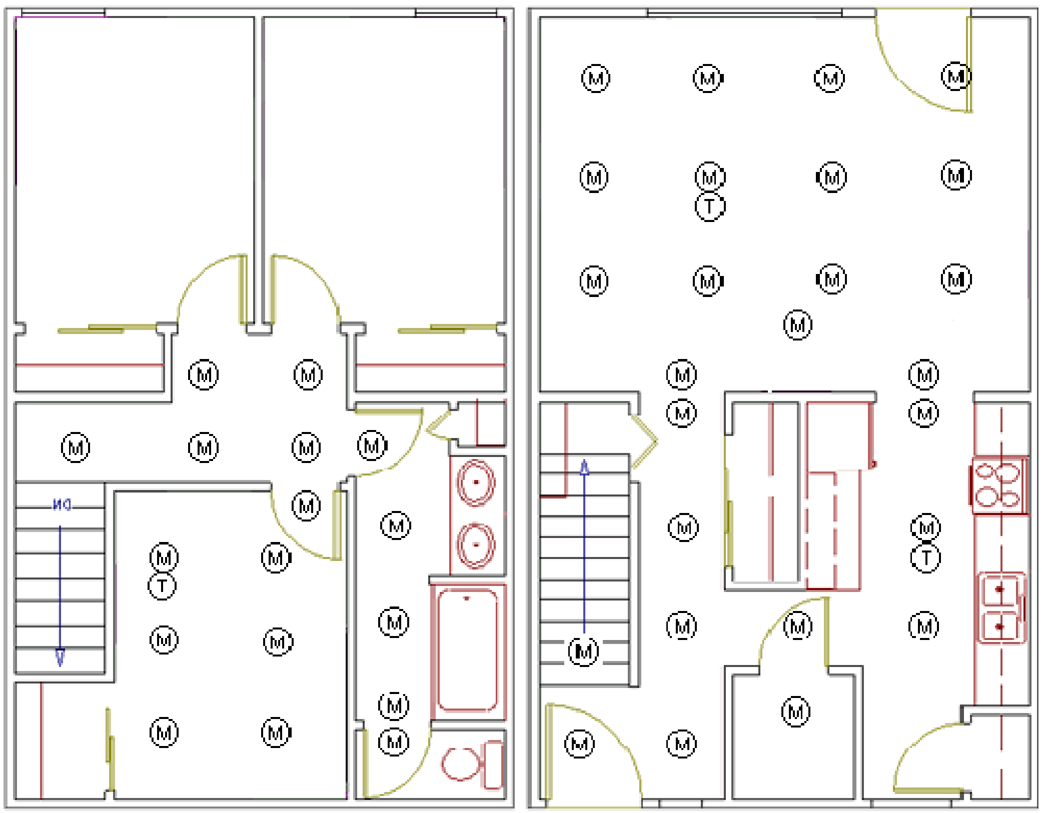

Figure 5. Recognition of the "cooking" activity in the smart apartment. The first sensor reading indicates that one of the tagged items is being used. The next entry indicates that the cabinet door was closed. The next two entries reflect that a motion sensor was activated then deactivated, and the last entry shows a non-zero reading for the stove burner. A visualization of all of the sensor values that are active at this point during the "cooking" task is shown at the bottom.

Using the sequential probability distribution that can be directly computed from the Markov model, we can observe a sequence of sensor events and identify the model (and the task that the model represents) that yields the highest probability of corresponding activity to the observation sequence. Specifically, after each sensor event we will generate a label for the activity (or set of activities) that the participant is performing, and will use a forward probability-propagating algorithm to identify the belief state (or current state) of the corresponding activities. Using this approach we can recognize if and when key ADLs were performed and by whom.

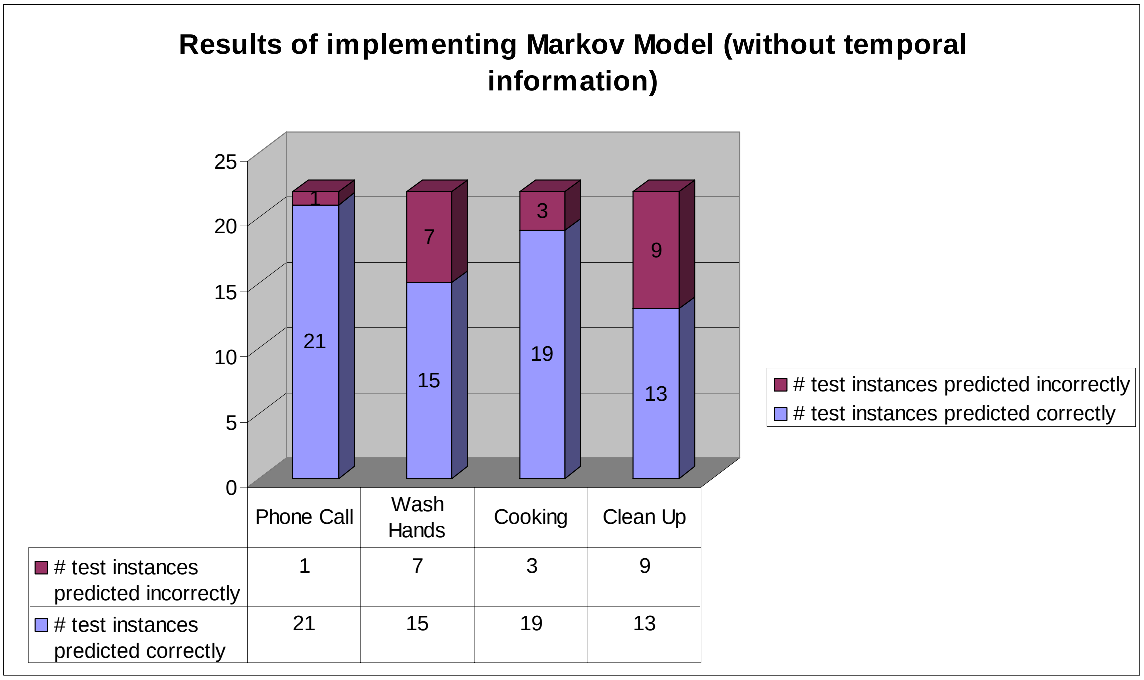

Figure 6. Bar graph showing results of using Markov model in identifying activities (without using any timing information).

Experimental Results

To validate our activity recognition algorithms we recruited 22 volunteer participants to perform a sequence of ADL activities in the smart apartment testbed. The activities we have selected include both instrumental ADLs which are disrupted in early-stage dementia and basic ADLs which are disrupted in the later stages of dementia. Specifically, the participant makes a phone call (while sitting at the dining room table) to obtain a cooking recipe. The participant then washes his/her hands, follows the directions in the recipe to cook food, takes the food with some medicine into the dining room, eats the food with the medicine, and finally cleans the dishes. Figures 4 and 5 show images from the “wash hands” and “cooking” activities together with a sample of the sensor events these sequences generate and a visualization of the sensor events.

We separated the activities into distinct event streams for training and testing. We leaned Markov models for each of the distinct activities and used the models to automatically label the sensor event streams from the test set. In the results we report below, we show accuracy results generated using three-fold cross validation on the participant data.

The four activities viz. “making phone call”, “washing hands”, “cooking” and “cleaning up” are represented on the x-axis and the y-axis depicts the accuracy in their prediction. This model shows an overall accuracy of 77.27% on the test set, with the individual accuracies of 95.45%, 68.18%, 86.36% and 59.09% respectively for the four activities. It can be observed that the model shows lowest accuracy in predicting the second and the last activity i.e. “wash hands” and “cleaning up”. As these activities are performed in the same area and they trigger similar sensors, the Markov models generated for these two activities overlap largely. Also, “wash hands” being a smaller activity forms a subset of the “cleaning up” activity and the “cleaning up” activity is thus incorrectly predicted as “wash hands” activity by the model. In the second part of our experiment, we augment our models with temporal information by associating normal distribution of the time spent in each sub-task of every activity in addition to the sequential information of sensor events. This helps in distinguishing between overlapping activities like in the above case as evidenced by the results below.

The Markov models when augmented with temporal information show an overall accuracy of 88.63%. This is an improvement of 11.36% over the previous model. It can be noticed from the results in Figure 7 that adding temporal information to the Markov models greatly enhanced their accuracy of prediction in case of similar activities. The accuracy of predicting the activity “cleaning up” showed the maximum increase from 59.09% to 88.63%. The bar graph in Figure 8 brings together the two approaches and shows a contrast between performance of Markov model with and without temporal information to facilitate easy comparison.

Conclusions

In this paper we describe an approach to recognize activities that are performed by residents of smart environments. Not only do we demonstrate that these activities can be recognized by sensors in physical environments using Markov models, but we also show that the recognition accuracy is greatly improved through the use of temporal event duration information. This increased accuracy will be important as we move on to our next steps. In particular, we will next be investigating techniques for detecting missing or incorrect steps in the activity, for detecting activities when multiple activities are interleaved, and for recognizing activities when there are multiple residents in the

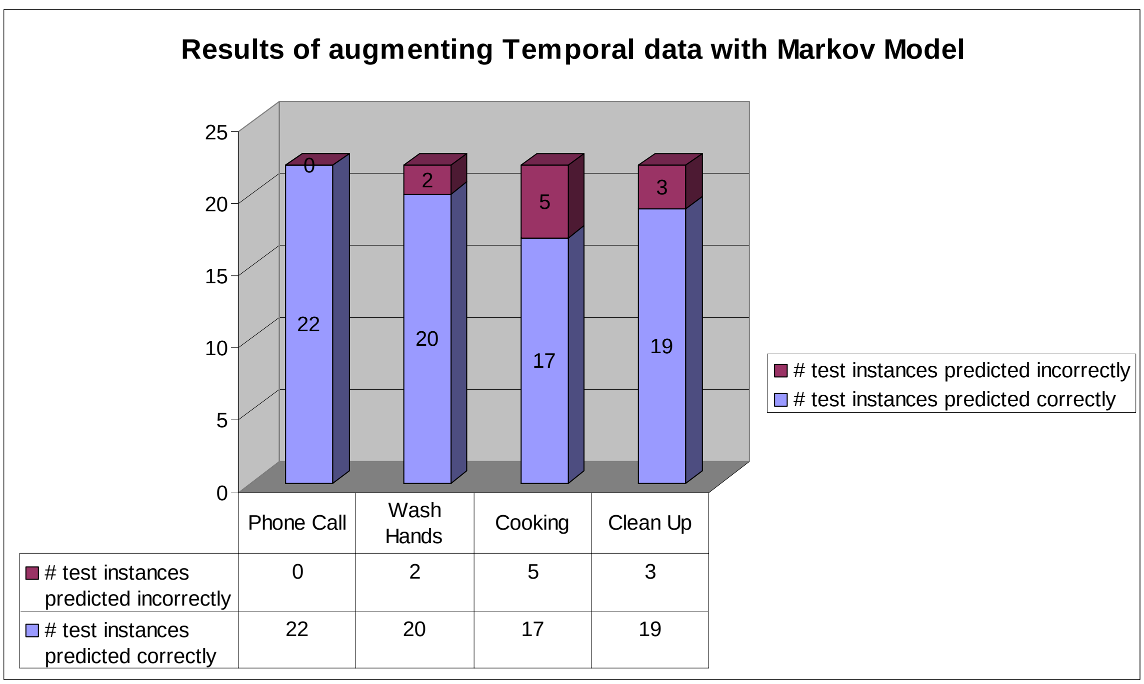

	Phone Call	Wash Hands	Cooking	Clean Up
# test instances predicted incorrectly	0	2	5	3
# test instances predicted correctly	22	20	17	19

Figure 7. Bar graph showing results of using Markov model with temporal information.

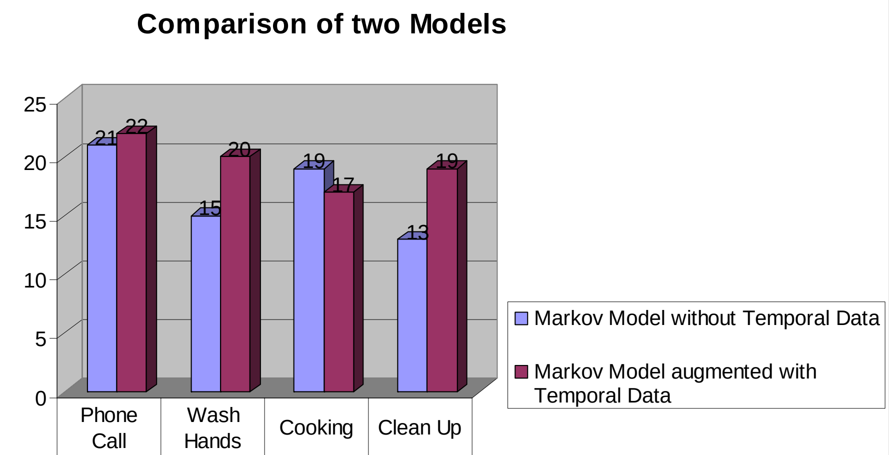

	Phone Call	Wash Hands	Cooking	Clean Up
Markov Model without Temporal Data	21	15	19	13
Markov Model augmented with Temporal Data	22	20	17	19

Figure 8. Bar graph showing comparison between the Markov models with and without temporal information.

environment.

The proposed research described here will lay the groundwork for follow-on research in which we conduct wide-spread testing of individuals in different segments of the population and with varying degrees of cognitive and physical limitations. The approaches described in this application can be used to monitor and assist with the rehabilitation progress of individuals with various types of injuries (e.g., traumatic brain injury, spinal chord injury). In addition, in-home monitoring of resident activity, diet, and exercise compliance can be extremely beneficial for diabetes patients and for individuals who suffer from drug or alcohol abuse. We also plan to extend this technology to monitor a variety of interventions outside of the home.

We believe these technologies are essential to provide accessible and low-cost health assistance *in an individual's own home*. Furthermore, investigating these issues will be imperative if we want to adequately care for our aging population and provide the best possible quality of life for them and, ultimately, for ourselves.

References

AARP. These four walls… Americans 45+ talk about home and community, 2003.

Abowd G., and Mynatt, E.. Designing for the human experience in smart environments. In D. Cook and S. Das, editors, *Smart Environments: Technology, Protocols, and Applications*, pages 153-174, Wiley, 2004.

Barger, T., Brown, D., and Alwan, M.. Health status monitoring through analysis of behavioral patterns. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 35(1):22-27, 2005.

Diehl, M., Marsiske, M., Horgas, A., Rosenberg, A., Saczynski, J., and Willis, S.. The revised observed tasks of daily living: A performance-based assessment of everyday problem solving in older adults. *Journal of Applied Gerontology*, 24:211-230, 2005.

Doctor, F., Hagras, H., and Callaghan, V.. A fuzzy embedded agent-based approach for realizing ambient intelligence in intelligent inhabited environments. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 35(1):55-56, 2005.

Ernst, R.L. and Hay, J.W.. The US economic and social costs of Alzheimer's disease revisited. *American Journal of Public Health*, 84, 1261-1264, 1994.

Gross, J.. A grass-roots effort to grow old at home. *The New York Times*, August 14, 2007.

Helal, A., Mann, W., El-Zabadani, H., King, J., Kaddoura, Y., and Jansen, E.. The Gator Tech smart house: A programmable pervasive space. *IEEE Computer*, 38(3):50-60, 2005.

Herbert, L.E., Scherr, P.A., Bienias, J.L., Bennett, D.A., and Evans, D.A. Alzheimer's disease in the US population: Prevalence estimates using the 2000 census. *Archives of Neurology*, 60:1119-1122, 2000.

Intel Proactive Health, www.intel.com/research/prohealth/cs-aging_in_place.htm, 2007.

Kautz, H., Arnstein, L., Borriello, G., Etzioni, O., and Fox, D. An overview of the assisted cognition project. *Proceedings of the AAAI Workshop on Automation as Caregiver: The Role of Intelligent Technology in Elder Care*, pages 60-65, 2002.

Kientz, J.A., Arriaga, R.I., Chetty, M., Haytes, G.R., Richardson, J., Patel, S.N., and Abowd, G.D. Grow and know: Understanding record-keeping needs for the development of young children. *Proceedings of the Conference on Computer-Human Interaction*, 2007.

Lanspery, S., Callahan, Jr., J.J., Miller, J.R., and Hyde, J.. Introduction: Staying put. In Lanspery, S. and Hyde, J., editors, *Staying Put: Adapting the Places Instead of the People*. Baywood Publishing Company, pages 1-22, 1997.

Larson, C.. In elder care, signing on becomes a way to drop by. *The New York Times*, February 4, 2007.

Liao, L., Fox, D., and Kautz, H. Location-based activity recognition using relational Markov networks. *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 773-778, 2005.

Mihailidis, A., Barbenl, J., and Fernie, G. The efficacy of an intelligent cognitive orthosis to facilitate handwashing by persons with moderate-to-severe dementia. *Neuropsychological Rehabilitation*, 14(1/2):135-171, 2004.

Mozer, M.. Lessons from an adaptive home. In Cook, D. and Das, S., editors, *Smart Environments: Technology, Protocols, and Applications*, pages 273-298, Wiley, 2004.

Noury, N.. A smart sensor for the remote follow up of activity and fall detection of the elderly. *Proceedings of the International IEEE-EMBS Special Topic Conference on Microtechnologies in Medicine and Biology*, pages 314-317, 2002.

Philipose, M., Fishkin, K., Perkowitz, M., Patterson,, D., Fox, D., Kautz, H., and Hahnel, D.. Inferring activities from interactions with objects. *IEEE Pervasive Computing*, 3(4):50-57, 2004.

Pigot, H., Mayers, A., Giroux, S., Lefebvre, B., Rialle, V., Noury, N.. Smart house for frail and cognitive impaired elders. *Proceedings of the First International Workshop on Ubiquitous Computing for Cognitive Aids (UbiCog '02)*, Goteborg, Sweden, 2002.

Pollack, M.. Intelligent technology for an aging population: The use of AI to assist elders with cognitive impairment. *AI Magazine*, 26(2):9-24, 2005.

Rialle, V., Ollivet, C., Guigui, C., and Herve, C.. What do family caregivers of Alzheimer's disease patients desire in smart home technologies? *Methods of Information in Medicine*, 47:63-69, 2008.

Sixsmith, A., and Johnson, N.. SIMBAD: Smart inactivity monitor using array-based detector. *Gerontology*, 2(1):110, 2002.

Swisher, K.. "Intelligent" appliances will soon invade homes, *The Wall Street Journal Online*, 2006.

Tam, T., Boger, J., Dolan, A., and Mihailidis, A. An intelligent emergency response system: Preliminary development and testing of a functional health monitoring system. *Gerontology*, 4(4):209-222, 2006.

Wren, C. and Munguia-Tapia, E. Toward scalable activity recognition for sensor networks, *Proceedings of the Workshop on Location and Context-Awareness*, 2006.

Youngblood, G.M. and Cook, D.. Data mining for hierarchical model creation. *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, 37(4):1-12, 2007.