

# Detection of Social Interaction in Smart Spaces

Diane J. Cook, Aaron Crandall, Geetika Singla, and Brian Thomas  
Washington State University  
Pullman, WA 99163  
djcook@wsu.edu

## Abstract

The pervasive sensing technologies found in smart environments offer unprecedented opportunities for monitoring and assisting the individuals who live and work in these spaces. An aspect of daily life that is important for one's emotional and physical health is social interaction. In this paper we investigate the use of smart environment technologies to detect and analyze interactions in smart spaces. We introduce techniques for collect and analyzing sensor information in smart environments to help in interpreting resident behavior patterns and determining when multiple residents are interacting. The effectiveness of our techniques is evaluated using two physical smart environment testbeds.

## 1. 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 applications, smart environments can also be employed for valuable functions such as at-home health monitoring and automation assistance. The long-term goal of our CASAS smart environment project [29] is to perform automated health monitoring and to provide automated assistance that will allow individuals to remain independent in their own homes. Given the aging of the population, the cost of formal health care, and the importance that individuals place on remaining independent in their own homes [1,16], these technologies will become an increasingly important component of our everyday lives.

The emphasis of smart home assistance for individuals with special needs has been to monitor completion of ADL (Activities of Daily Living) activities [5,23,27]. An area that has not received as much attention is monitoring of interactions for smart environment residents. Automating detection and

analysis of social interactions in smart spaces is the focus of this paper. Within this paper, we present and empirically validate algorithms that can visualize and analyze sensor data collected in a smart space to detect social interaction. Smart environment technologies have been employed to track activities, to monitor the well-being of residents, and to provide some context-aware services for the environment inhabitants. Our project explores a new direction for smart environment research. In particular, we hypothesize that smart environment sensor data and computational tools can be used to effectively detect certain types of social interaction in everyday environments. To validate our hypothesis, we first employ tools to visualize activity levels in smart environments and identify likely times and places of resident interactions. Second, we design an unsupervised learning algorithm to automatically identify the likely times and places of these interactions. Third, we make use of a supervised learning technique to recognize the current activities, including those that involve resident interactions. To assess the efficacy of the technologies for such real-world settings, we evaluate our algorithms using data collected in our on-campus smart apartment and our on-campus smart workplace with volunteer residents.

## **2. Social Interaction Detection**

Interaction among individuals is foundational to human health. In the movie *Cast Away*, Tom Hanks plays a FedEx employee who is stranded alone on an island. His need for interaction and personal connection is so vital that he names a volleyball “Wilson” and lets the ball play the role of his companion on the island. Ambady et al. [3] discuss how social signals are determinants of human behavior. In addition, Baker et al. [4] have shown that work environments which provide access to other people improve productivity for employees. On the other hand, socialization for individuals with special needs can be more difficult to achieve. The CDC [7] reports that older adults tend to experience greater social isolation. In turn, York and Waite [38] point out that social isolation in turn damages health. Fratiglioni et al. [14] have even shown that older people who live alone are more likely to develop dementia.

Although there is a growing interest in the design of smart environments and ambient intelligence applications, only recently have researchers used these technologies to study social behaviors. Griswold

et al. [16] explore the use of PDAs to provide awareness of colleagues' locations and working contexts. Olguin et al. [26] use Bluetooth and IR sensors to detect interactions such as proximity and face-to-face time. They use this information to identify organizational behavior and predict job satisfaction. Eagle and Pentland [13] detect social ties, and group affiliations together with daily behaviors based on mobile phone usage. Choudhury and Pentland [8] identify social network organizations among individuals by analyzing phone conversations.

While these approaches have made effective strides in automatically detecting social relationships, they tend to rely on mobile wearable or hand-held devices. We are interested in determining if resident interaction can be detected in a smart environment filled with passive, unobtrusive [18] sensors. We treat a smart environment as an intelligent agent that perceives the state of the resident and the physical surroundings using sensors and acts on the environment using controllers in such a way that the specified performance measured is optimized [9]. Researchers have generated ideas for smart environment software algorithms that track the location of single residents, that generate reminders, and that react to hazardous situations [37]. Some projects with physical testbeds have begun to emerge including the MavHome [39], the Gator Tech Smart House [17], the iDorm [12], and the Georgia Tech Aware Home [2]. Resulting from these advances, researchers are now beginning to recognize the importance of applying smart environment technology to health assistance [5,20,21,24,28] and companies are recognizing the potential of this technology for a quickly-growing consumer base [19].

In order to detect social interactions in smart spaces, we collect sensor data in our smart environment testbeds. First, we show that event density maps provide initial insights into resident behavior patterns and likely times and places for resident interaction. Second, we employ Bayesian updating to track residents through a space and automatically detect interactions based on proximity. Lastly, we conduct a controlled experiment to determine our ability to recognize interactions and caregiver assistance using hidden Markov models. All of our data and experimental validation is performed in the context of our two physical smart environment testbeds: a smart apartment and a smart workplace.

### 3. Testbeds

Our smart environment testbeds are located on the Washington State University campus and are maintained as part of our ongoing CASAS smart home project [29]. As shown in Figure 1, the smart apartment testbed contains three bedrooms, one bathroom, a kitchen, a living room, and a dining area. The apartment is equipped with motion sensors distributed approximately 1 meter apart throughout the space. In addition, we have also installed sensors to provide ambient temperature readings and custom-built analog sensors to provide readings for hot water, cold water, and stove burner use. Contact switches allow us to monitor usage of key items including a cooking pot, a medicine container, and the phone book. In addition, Insteon™ power controls monitor usage and control lighting throughout the space. Sensor data is captured using an in-house sensor network and is stored in a SQL database. Our middleware uses a XMPP-based publish-subscribe protocol as a lightweight platform and language-independent method to push data to client tools (i.e., our data analysis and application programs).

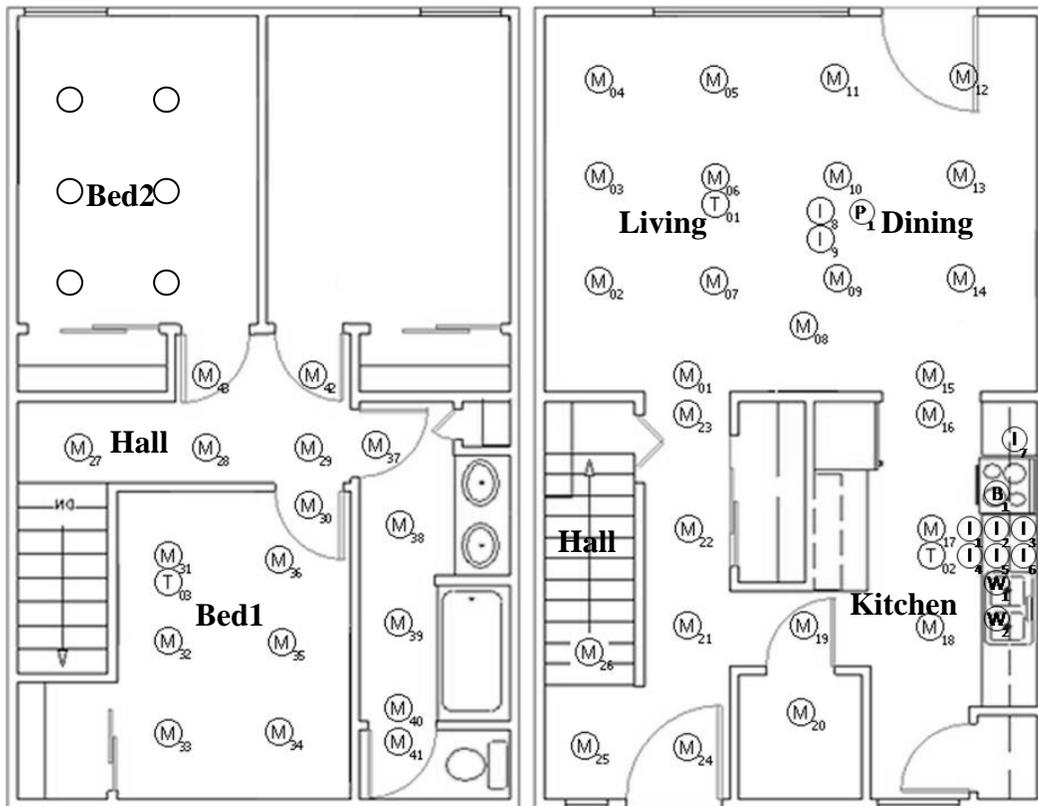


Figure 1. WSU smart apartment testbed. Sensors in the apartment monitor motion (M), temperature (T), burner (B), telephone (P), and item (I) use.

In addition, we have equipped an on-campus smart workplace environment, shown in Figure 2. This is a laboratory that is organized into four cubicles with desks and computers, an open server area, a postdoc office, a meeting area, a lounge, and a kitchen. Like the apartment, the lab is equipped with motion sensors placed approximately 1 meter apart throughout the space and magnetic sensors record door openings and closing. In addition, powerline controllers operate all of the lights in the room. Each sensor event is represented by the event's date, time, sensor ID, and sensor reading.

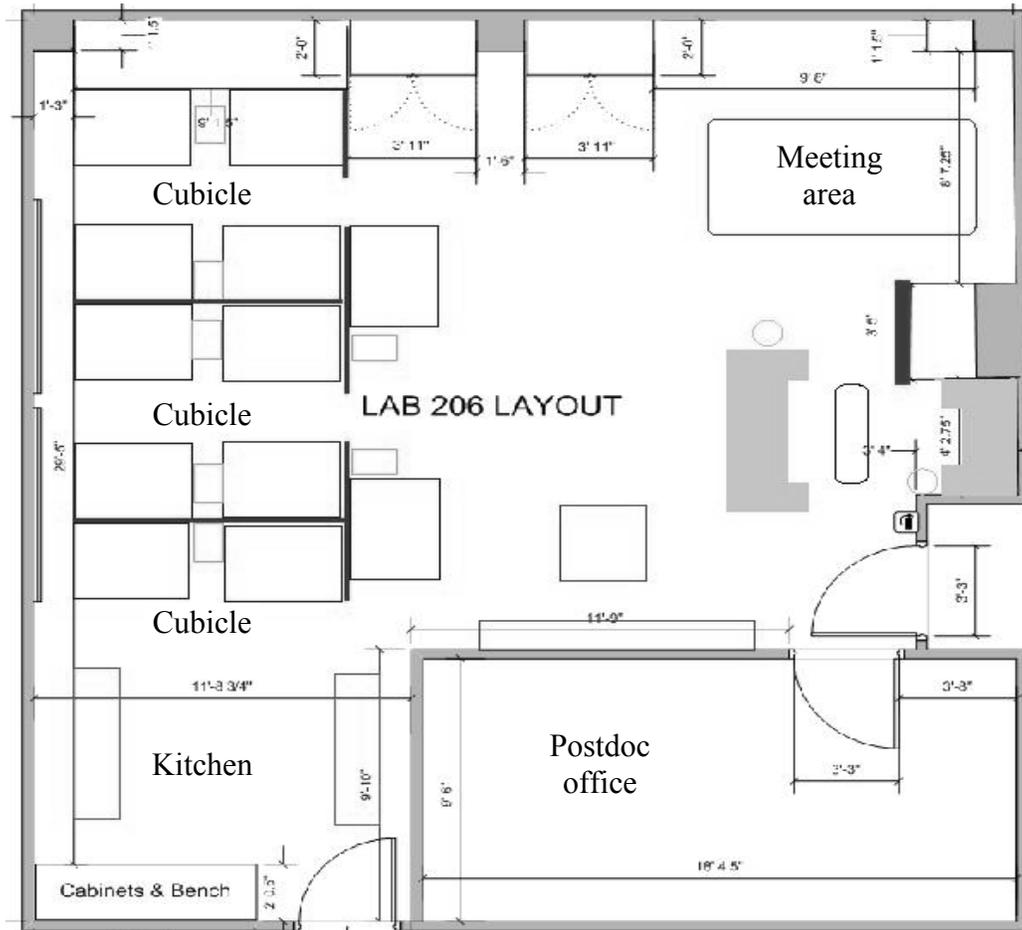


Figure 2. WSU smart workplace testbed.

The sensors that we use in these environments allow our algorithms to recognize and track daily activities. While this feature alone has benefits for monitoring the functional and physical well-being of residents, we hypothesize that the data can be used additionally to detect types of social interactions. We

do not employ cameras or microphones in these testbed. While they may offer valuable insights for social interaction detection, they are typically not well-accepted by the community that we want to serve with this technology [11] and therefore are not used as part of our smart environment testbeds.

#### 4. Visualizing Smart Environment Activities

For our first dataset, we collected data in the smart apartment while two undergraduate students were living there. While the residents did know each other, they also were taking classes; therefore, interactions between the residents were likely to occur sporadically or only at key times during the day and week. The dataset contains a total of 50,048 sensor events spread over 15 days.

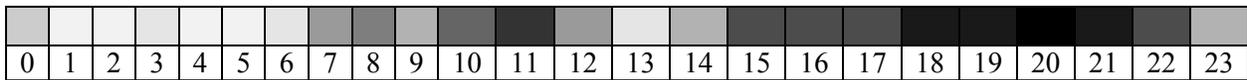


Figure 3. Hour-by-hour event density map for the smart apartment with two residents.

As the map in Figure 3 shows, the residents as a whole were most active during mid-to-late morning and around dinner time. This behavior is intuitive and consistent with the lifestyle of college students.

The next density map, shown in Figure 4, highlights the regions in which sensor events occurred, totaled over the entire data collection period and normalized to fall between 0 and 1. As the map indicates, the greatest number of events occurs in the bedrooms (the resident in the top bedroom is a very restless sleeper) and in the common areas of the living room, dining room, and kitchen.

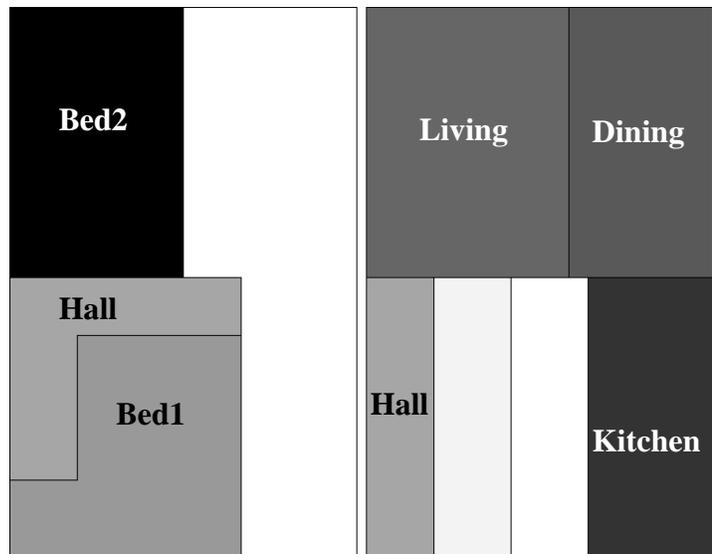


Figure 4. Region-based density map for the smart apartment with two residents.

Looking at the activity density maps, we hypothesize that the likeliest times and regions for social interaction between these two residents is in the kitchen followed by the living and dining areas, from late

afternoon until late evening. In the next section of the paper, we will determine if this hypothesis is supported based on Bayesian-based tracking of the residents.

In order to contrast social interaction in an apartment setting with a workplace, we also perform data collection in our smart workplace. We collected data in the smart workplace testbed for three and a half months, resulting in 247,545 collected sensor events. During this time, five students worked fairly regularly in the lab and another five students occasionally spent time in the space.



Figure 5. Hour-by-hour event density map for the smart workplace.

Once again, we generate event density maps averaged for each hour of the day (Figure 5) and averaged over each region in the space (Figure 6). The figures show that the workplace environment exhibits different behavioral patterns than the apartment. Unlike the apartment, there appears to be one peak of activity in the lab, around mid afternoon. Activity levels decay smoothly from the peak on either side. The region with the greatest amount of activity is the open area (the dark region in Figure 6).

Residents move through this area when they enter and exit the lab and when they move to the meeting area, other cubicles, or to get printouts. The next most active areas are the meeting area, the postdoc area, and the middle cubicle. The postdoc office and the middle cubicle house researchers that not only spend the most hours in the lab but are also the most interactive and social of the lab members. The region-based activity density map thus also mirrors our intuitive impression of where activities and interactions would likely occur in the space.

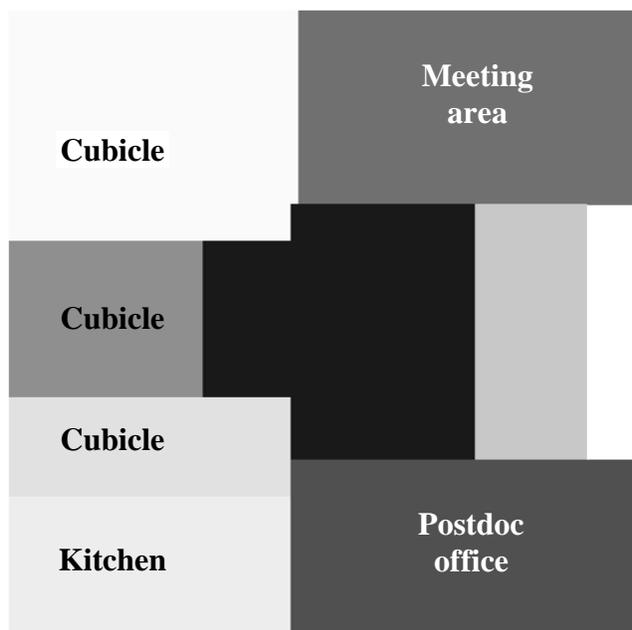


Figure 6. Region-based density map for the smart workplace.

This type of activity visualization provides us with insights on activity in the space and allows us to speculate about when interactions between residents are likely to occur. However, no actual interactions, here defined as time spent in close proximity, are detected or visualized. Automatically detecting these interactions will be the focus of the next section.

## 5. Detecting Interactions

In order to detect interactions between residents in a smart space, we need to determine the locations of each individual in the space. This can be accomplished by asking each resident to wear a locating device, but such an approach is impractical for general use and is certainly intrusive for the residents. Instead, we are interested in determining if we can detect interactions, or time spent in close proximity, using passive sensors in the space.

To detect interactions we use Bayesian updating to track the individuals through the space. Using Bayesian updating, after processing the  $t^{\text{th}}$  sensor event we calculate the probability that a resident is at location  $loc_t$  using the formula shown in Equation 1. Here  $loc_t$  refers to a possible location for the resident at time step  $t$ ,  $sensor_t$  refers to the type of sensor event that was observed at time step  $t$ , and  $loc_{t-1}$  refers to the location of the resident at the previous time step.

$$P(loc_t) = \alpha * P(sensor_t | loc_t) * \sum_{loc_{t-1}} p(loc_t | loc_{t-1}) * P(loc_{t-1}) \quad (1)$$

At the beginning of each day, the probability values of resident locations are initialized to be 0.9 for their bedroom locations (Bed1 for Resident 1 and Bed2 for Resident 2) and uniformly-distributed small probability values for the other locations. This is consistent with the actual locations of the residents observed during manual annotation of the data. For each subsequent location update, Equation 1 sums the probability of all previous locations for the resident multiplied by the probability of transitioning from the old location to the new. These probability values, based on transition frequencies found in the data set, are larger for locations that stay the same (reflecting the fact that residents stay in one location more than they move around). Smaller probability values are assigned to neighboring locations (as would occur when the resident moves around the space), and very small non-zero transition values are assigned to

location pairs that are not neighbors (a person cannot jump from one space to another disconnected space, but they can move stealthily enough that their movement would not be caught by motion sensors). This summation is multiplied by the probability that the observed sensor event would occur for a particular resident location  $loc_i$ , and the total is multiplied by a normalizing constant,  $\alpha$ , to make sure that the sum of probabilities for all possible resident locations is equal to 1. This update is performed independently for each resident after each observed sensor event.

Ironically, one difficulty we encountered using Bayesian updating for multiple residents was actually recognizing when one resident moved into close proximity with another resident. If Resident 1 is in a specific location and a sensor event occurs at that location, it initially seems more likely that Resident 1 triggered the event than that Resident 2 moved into the location. Looking at sample sensor data, we found one source of information that was not reflected in our initial probability assignments. Our previous single-resident datasets showed that when the resident was in the space a sensor event (motion, door, or item) would be generated at least every four minutes when they were awake (they were outside their bedroom) and at least every ten minutes when they were asleep (they were inside their bedroom). In the smart workplace, sensor events were generated in an individual’s location at least every two minutes. We updated our probability distributions (specifically, the probability of transitioning from one location to another) to reflect this observation. As a result, the automatically-detected interactions between residents closely approximated the interactions we detected through manual inspection of the data.

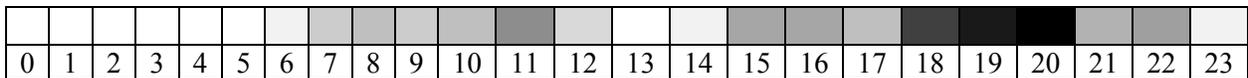


Figure 7. Resident interaction frequency in the smart apartment.

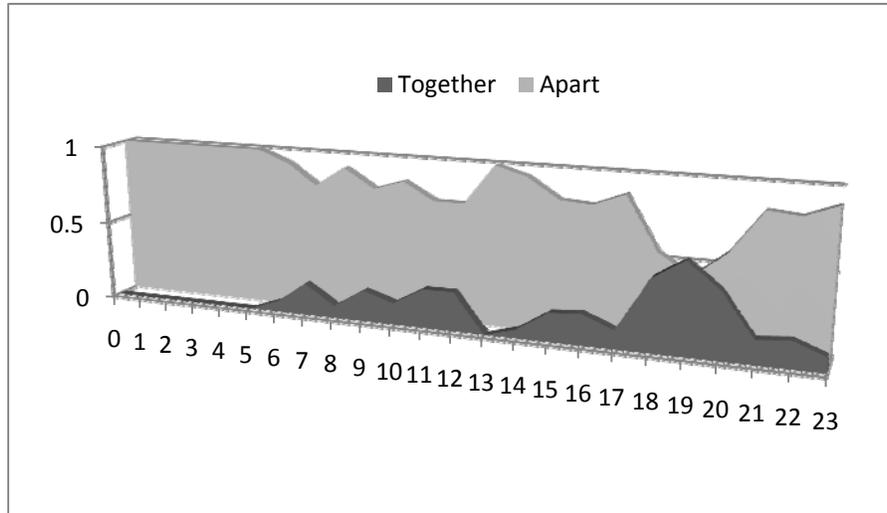


Figure 8. Relative frequency of resident interactions in the smart apartment.

Figure 7 shows the frequency of resident interactions for each hour of the day in the smart apartment, and Figure 8 shows the proportion of sensor events each hour that occur when the residents are together or are separated in the space. Both of these figures highlight the fact that almost all of the interactions occur around dinner time. Even though a great number of sensor events do occur in the middle of the night, the residents are not interacting during that time. In fact, a large number of these sensor events are hypothesized to be noisy events triggered by floor heaters that generate sudden warm air gusts near a motion sensor in the cold hours of the night. Detecting anomalous events of this type that do not fit a known activity pattern is an interesting area of additional research with application to resident security.

Similarly, we performed automatic detection of resident interaction in the smart workplace and generated interaction density maps for this testbed, as shown in Figures 9 and 10.

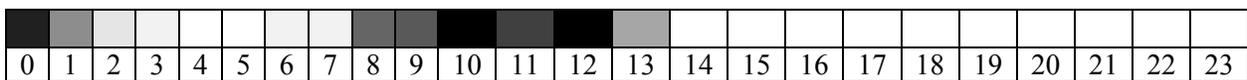


Figure 9. Resident interaction frequency in the smart workplace.

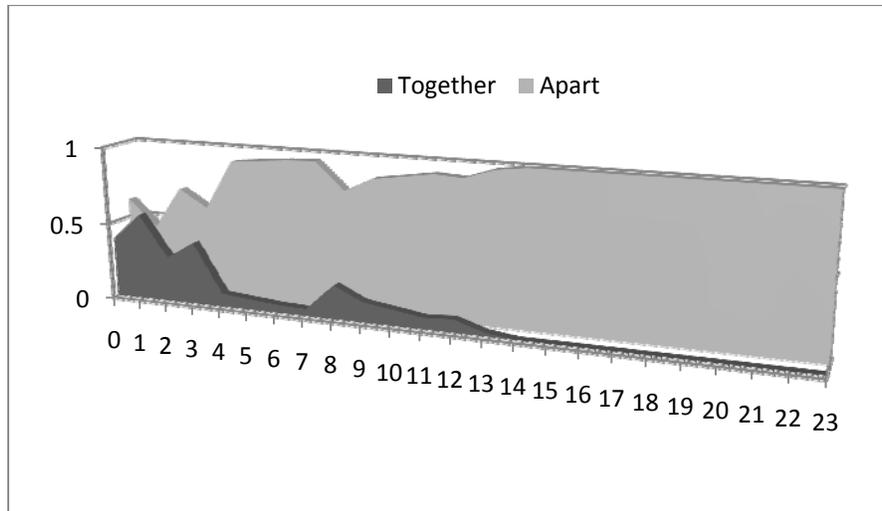


Figure 10. Relative frequency of resident interactions in the smart workplace.

The findings for the smart workplace reflect different dynamics than for the apartment. While interactions do occur in the workplace, they typically have shorter duration. There are fewer interactions overall as a proportion of the number of residents. Interestingly, most interactions occur late in the evening when students are working and close to lunch time.

This approach to unsupervised detection of resident interaction is useful for determining the amount, times, and locations of resident interactions. The information that is provided can in turn shed light on the types of interactions that most likely occur. In our final approach we design a supervised learning algorithm to recognize activities, both those that occur individually and those that involve regular or sporadic interaction with another resident.

## 6. Recognizing Interactions

For our last approach, we design a supervised learning algorithm to recognize activities that occur in a smart environment. Using this approach we can identify activities that typically involve a single resident and those that involve multiple residents. We will also use this information to verify the accuracy of our unsupervised interaction detection algorithm.

For this study, we collected data while 40 WSU undergraduates came into the smart apartment, two at a time, and performed the following set of eight activities:

1. Resident 1: Fill a medication dispenser, found in the kitchen, with pills from a bottle.
2. Resident 2 (at the same time): Hang up clothes that are laid out on the couch in the living room.
3. Resident 2: Move the couch to the other side of the living room. During this task, ask Resident 1 to come and help.
4. Resident 1: Water plants around the apartment using a watering can from the kitchen closet.
5. Resident 2 (at the same time): Sweep the kitchen floor using a broom and dust pan from the kitchen closet.
6. Resident 1 and Resident 2 (together): Play a game of checkers.
7. Resident 1: Prepare dinner using the ingredient and recipe on the kitchen counter.
8. Resident 2 (at the same time): Simulate paying bills. During the task ask Resident 1 to come help find a number in the phone book.

The selected activities include basic and complex ADL activities that are typically found in clinical questionnaires [30]. Researchers have conducted studies that assess the ability of machine learning technologies to recognize activities using wearable sensors [23], by monitoring interactions with objects in the environment [25,27], by videotaping activities [6], and by analyzing motion sensor data [7]. A variety of models including naïve Bayes classifiers [6,17,34], decision trees [23], and probabilistic model such as Markov models, dynamic Bayes networks, and conditional random fields [10,22,27,32] have been tested. While these studies have indicated the power of algorithmic methods for activity recognition, they have been tested in single-resident settings where the activities are uninterrupted. In contrast, our approach handles and is tested here on cases where some of the activities are interrupted (Resident 1's activities are interrupted when Resident 2 calls for help), some activities occur in parallel, and some activities involve multiple residents. The resident interactions occur when an activity is performed together (e.g., checkers) or when one person requests help of the other (e.g., Resident 2 asks for help finding a phone number). Recognizing these types of interactions can be helpful in assessing an older adult's need for assistance from a caregiver.

In order to recognize these activities in a smart environment, we use a portion of the data as sample data for creating a model of the activities. Specifically, we use a hidden Markov model (HMM) as a statistical model of the dynamic system. A HMM models the system using a finite set of states. Each observable and hidden state is associated with a multidimensional probability distribution over a set of parameters. The system is assumed to be a Markov process, so the current state depends only on the previous state. Transitions between states are governed by transition probabilities. We constructed two HMMs. In the first model, the *Resident Identifier* model, a hidden state is used to represent each of the two residents. In the second model, the *Activity Identifier* model, a hidden state is used to represent each of the separate activities. We use the training data to learn the transition probabilities between hidden states for each model and to learn the relationship probabilities between the hidden states and the observable states (one for each configuration of sensor values).

To label a sequence of sensor event observations with the corresponding activity, we use the Viterbi algorithm [35] to compute the most likely sequence of hidden states that correspond to a sequence of observable sensor events. Applying this algorithm to the *Resident Identifier* model, we can label each sensor event with the resident that most likely triggered the event. Applying the algorithm to the *Activity Identifier* model using the labeled event data, we can label each sensor event with the activity that was being performed when the event occurred.

Figure 11 summarizes the accuracy results for the hidden Markov model using 3-fold cross validation on our set of 20 activity sensor streams which provide a total of 160 training examples. As can be seen from this summary, the Markov model correctly identified the majority of activities, even when they were performed in parallel with other activities, were interrupted, or were performed with multiple individuals.

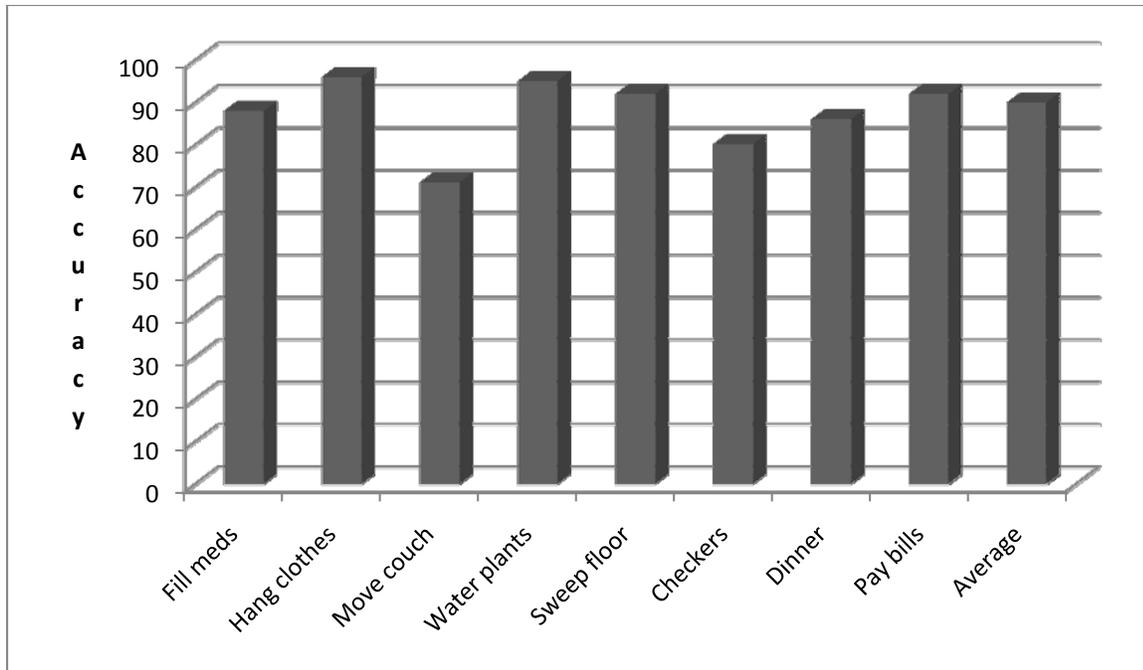


Figure 11. Activity recognition accuracy.

As the graph in Figure 11 shows, the hidden Markov model recognized the activities with good regularity, averaging an accuracy of 90%. The average precision over the 8 classes is 0.93, the average recall is 0.96, and the average f-score is 0.94. The activity that created the most difficulty for the algorithm, moving the couch, was particularly challenging because both residents moved around the space in an almost-random pattern while they positioned themselves and the couch in the apartment. The fact that even activities involving multiple residents, such as playing checkers, could be recognized indicates that resident interaction can be automatically recognized in smart spaces when it is associated with known activities. Similarly, when caregiver intervention is consistently required for an activity (e.g., for moving furniture and for paying bills), then this type of interaction will also be automatically recognized. As a point of comparison we also created a naïve Bayesian classifier (NBC) that determined activity labels probabilistically based on the number of each type of sensor event that occurred during the

activity. The NBC averaged 49% accuracy for this dataset, providing evidence that the hidden Markov model is a more effective approach for this type of classification problem.

In our final experiment, we applied our automated interaction detection algorithm to the data collected while the undergraduates performed these 8 ADL activities. Using the Bayesian updating, we identified the sensor events that occurred when the residents were in the same locations and those that occurred when they were apart (see Figure 12). We also tagged events as “coming” when one resident moved into a space that was occupied by the other resident and tagged events as “going” when one resident moved away from the other.

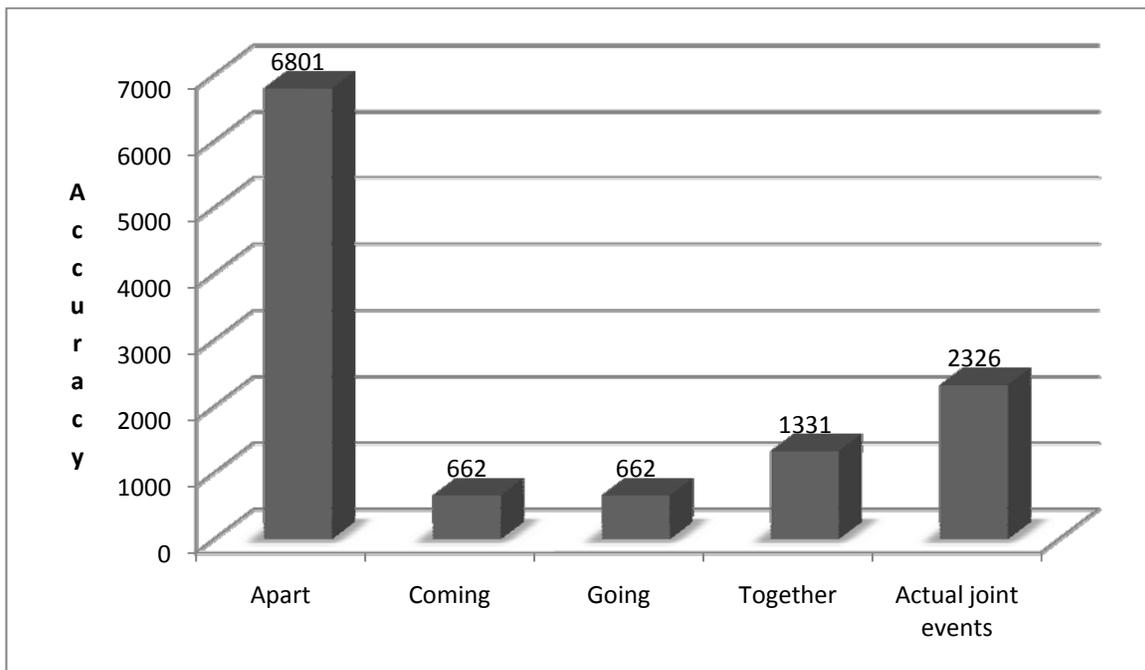


Figure 12. Detected and actual resident interactions. The number of events generated while residents interacted (2,326) is larger than the number of detected interactions (1,331).

We note that there are a much larger number of sensor events that were generated while the residents were told to interact than were automatically detected. Specifically, 2,326 sensor events were generated while the residents were playing checkers or during the times that Resident 1 assisted Resident 2 with an activity. In contrast, only 1,131 sensor events were

automatically detected as resident interactions. This is due primarily to the fact that some interactions occurred when the residents were not in close proximity. For example, 769 sensor events were generated while the residents moved a couch together. During a portion of that time the residents were not in the same location zone because they were holding the far ends of the couch. This type of discrepancy indicates that not all interactions can be determined by physical proximity – residents can interact functionally without being close to each other. We will need to continue to investigate methods by which other types of interaction can be detected.

## **7. Conclusions**

In this work we described an approach to analyzing resident interaction in smart spaces. Our algorithmic approach to interaction analysis includes visualization of sensor event density, automatic detection of close-proximity interactions, and automatic recognition of activities that involve resident interaction. We evaluated our algorithms on real data collected in our CASAS smart environment testbeds and demonstrated that each of the techniques provides a unique type of insight that is valuable for automated detection, recognition, and analysis of resident interaction.

This paper offers some algorithmic approaches to analyzing data in a very complex situation. There are clear challenges that still need to be addressed in this topic. For example, our approach has only been tested for two residents in a single environment. Detecting social interactions for three or more residents, for families, and for large groups will significantly increase the complexity of the analysis task and is a direction that we want to pursue in the future. Adding information from additional sensors, microphones, and RFID tags would also allow us to detect and analyze a greater variety of interaction types. In addition, fusing the resident identification model and activity identification model into a multi-layer hierarchical model might improve the activity recognition task for multi-resident settings. We would also like to apply these techniques to data that is collected in environments with older adults in order to analyze differences in amount and types of interactions between younger and older adults.

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 effect of social interactions on the overall health of a smart environment resident and for evaluating the effect of targeted interventions on the social health of the resident. We believe these technologies are valuable for providing automated health monitoring and assistance in an individual's everyday environments.

## References

1. AARP. These four walls... Americans 45+ talk about home and community, 2003.
2. G. Abowd and E. Mynatt. 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.
3. N. Ambady and R. Rosenthal. Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis. *Psychological Bulletin*, 2(1992):256-275.
4. W. Baker. Achieving success through social capital. Wiley, 2000.
5. T. Barger, D. Brown, and M. Alwan. Health status monitoring through analysis of behavioral patterns. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 35(2005):22-27.
6. O. Brdiczka, P. Reignier, and J. Crowley. Detecting individual activities from video in a smart home. *Proceedings of the International Conference on Knowledge-Based and Intelligent Information and Engineering Systems*, pages 363-370, 2007.
7. Centers for Disease Control and Prevention. Health aging for older adults. [www.cdc.gov/Aging](http://www.cdc.gov/Aging).
8. T. Choudhury and A. Pentland. Characterizing social networks using the sociometer. In *Proceedings of the North American Association of Computational Social and Organizational Science* (2004), 503-511.
9. D. Cook and S. Das, editors. *Smart Environments: Technology, Protocols, and Applications*. Wiley, 2004.

10. D. Cook and M. Schmitter-Edgecombe. Assessing the quality of activities in a smart environment. *Methods of Information in Medicine*, 2009.
11. G. Demiris, D.P. Oliver, G. Dickey, M. Skubic, and M. Rantz. Findings from a participatory evaluation of a smart home application for older adults. *Technology and Health Care*, 16:111-118, 2008.
12. F. Doctor, H. Hagraas, and V. Callaghan. A fuzzy embedded agent-based approach for realizing ambient intelligence in intelligent inhabited environments. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 35(2005):55-56.
13. N. Eagle and A. Pentland. Behavioral inference across cultures: Using telephones as a cultural lens. *IEEE Intelligent Systems*, 23(2008):62-64.
14. L. Fratiglioni, H. Wang, K. Ericsson, M. Maytan, and B. Winblad. Influence of social network on occurrence of dementia: A community-based longitudinal study. *The Lancet*, 355(2000):1315-1319.
15. W. Griswold, P. Shanahan, S. Brown, R. Boyer, M. Ratto, R. Shapiro, and T. Truong. ActiveCampus – Experiments in community-oriented ubiquitous computing. *IEEE Computer*, 37(2006):73-81.
16. J. Gross. A grass-roots effort to grow old at home. *The New York Times*, August 14, 2007.
17. A. Helal, W. Mann, H. El-Zabadani, J. King, Y. Kaddoura, and E. Jansen. The Gator Tech smart house: A programmable pervasive space. *IEEE Computer*, 38(2005):50-60.
18. B.K. Hensel, G. Demiris, and K.L. Courtney. Defining obtrusiveness in home telehealth technologies: A conceptual framework. *Journal of the American Medical Informatics Association*, 13(2006):428-431.
19. Intel Proactive Health, [http://www.intel.com/research/prohealth/cs-aging\\_in\\_place.htm](http://www.intel.com/research/prohealth/cs-aging_in_place.htm), 2007.
20. H. Kautz, L. Arnstein, G. Borriello, O. Etzioni, and D. Fox. An overview of the assisted cognition project. In *Proceedings of the AAAI Workshop on Automation as Caregiver: The Role of Intelligent Technology in Elder Care* (2002), 60-65.

21. C. Larson. In elder care, signing on becomes a way to drop by. *The New York Times*, February 4, 2007.
22. L. Liao, D. Fox, and H. Kautz. Location-based activity recognition using relational Markov networks. *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 773-778, 2005.
23. U. Maurer, A. Smailagic, D. Siewiorek, and M. Deisher. Activity recognition and monitoring using multiple sensors on different body positions. In *Proceedings of the International Workshop on Wearable and Implantable Body Sensor Networks (2006)*, 113-116.
24. A. Mihailidis, J. Barbenl, and G. Fernie. The efficacy of an intelligent cognitive orthosis to facilitate handwashing by persons with moderate-to-severe dementia. *Neuropsychological Rehabilitation*, 14(2004):135-171.
25. E. Munguia-Tapia, S.S. Intille, and K. Larson. Activity recognition in the home using simple and ubiquitous sensors. *Proceedings of PERVASIVE*, pages 158-186, 2004.
26. D. Olguin, B. Waber, T. Kim, A. Mohan, K. Area, and A. Pentland. Sensible organizations: Technology and methodology for automatically measuring organizational behaviors. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 39(2009):1-12.
27. M. Philipose, K. Fishkin, M. Perkowitz, D. Patterson, D. Fox, H. Kautz, and D. Hahnel. Inferring activities from interactions with objects. *IEEE Pervasive Computing*, 3(2004):50-57.
28. M. Pollack. Intelligent technology for an aging population: The use of AI to assist elders with cognitive impairment. *AI Magazine*, 26(2005):9-24.
29. P. Rashidi and D. Cook. Keeping the resident in the loop: Adapting the smart home to the user. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*, in press.
30. B. Reisberg, 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). *International Psychogeriatrics*, 13(2001):163-181.

31. V. Rialle, C. Ollivet, C. Guigui, and C. Herve. What do family caregivers of Alzheimer's disease patients desire in smart home technologies? *Methods of Information in Medicine*, 47(2008):63-69.
32. D. Sanchez, M. Tentori, and J. Favela. Activity recognition for the smart hospital. *IEEE Intelligent Systems*, 23(2008):50-57.
33. G. Singla, D. Cook, and M. Schmitter-Edgecombe. Incorporating temporal reasoning into activity recognition for smart home residents. In *Proceedings of the AAAI Workshop on Spatial and Temporal Reasoning* (2008).
34. T. van Kasteren and B. Kröse. Bayesian activity recognition in residence for elders. In *Proceedings of the International Conference on Intelligent Environments* (2008).
35. A. Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13(1967):260-269.
36. S. Wang and M. Skubic. Density map visualization from motion sensors for monitoring activity level. *Proceedings of the International Conference on Intelligence Environments* (2006), 1-8.
37. C. Wren & E. Munguia-Tapia. Toward scalable activity recognition for sensor networks, *Proceedings of the Workshop on Location and Context-Awareness* (2006).
38. E. York and L. Waite. Social disconnectedness, perceived isolation, and health among older adults. *Journal of Health and Social Behavior*, in press.
39. G. M. Youngblood and D. Cook. Data mining for hierarchical model creation. *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, 37(2007):1-12.