Collecting and Disseminating Smart Home Sensor Data in the CASAS Project

D.J. Cook, M. Schmitter-Edgecombe, Aaron Crandall, Chad Sanders, and Brian Thomas
Washington State University
Pullman, WA
cook@eecs.wsu.edu

Abstract

Although smart environment technology has rapidly been maturing, the performance of these technologies is still difficult to assess because of the limited evaluation that has been conducted. A primary limitation in evaluating technologies is the lack of rich physical datasets on which the algorithms can be tested. In this position paper we describe a publicly-available dataset that was created as part of the CASAS project and discuss challenges that are faced when creating and disseminating such data.

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. The 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.

While technologies are currently being designed to meet this need, testing of the technologies on real data is a major challenge. There is a tremendous amount of overhead in constructing a physical smart home testbed. Expertise and resources are needed to design and install the sensors, controllers, network components, and middleware just to perform basic data collections. As a result, very few physical testbeds exist. In the cases where real sensor data has been collected and analyzed, only rarely is this data made available to the community.

We argue that shared home behavior datasets are critical in order to test, compare, and enhance smart home and telemedicine technologies such as user modeling, activity recognition, assessment of resident well being, and automation assistance. In this paper we describe the home behavior datasets that we have created and discuss the challenges that are faced when generating such datasets.

2. CASAS Dataset

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 that perform robustly even when multiple residents are present and activities are interleaved, and that identify missing steps in the activity execution.

The testbed that we are using to validate our algorithms is a three-bedroom apartment located on the Washington State University campus that is part of the ongoing CASAS smart home project at WSU. The CASAS project treats environments as intelligent agents, where the status of the residents and their physical surroundings are perceived using sensors and the environment is acted upon using controllers in a way that improves the comfort, safety, and/or productivity of the residents [Cook and Das, 2004].

As shown in Figure 1, the smart apartment testbed includes three bedrooms, one bathroom, a kitchen, and a living / dining room. The apartment is equipped with motion sensors distributed approximately 1 meter apart throughout the space. In addition, we have installed digital sensors to provide ambient temperature readings, and analog sensors to provide readings for hot water, cold water, and stove burner use. Asterisk-enabled VOIP captures phone usage [Asterisk] and we use contact switch sensors to monitor usage of the phone book, a cooking pot, the medicine container, and key cooking ingredients in the apartment. Sensor data is captured using a customized sensor network and is stored in a SQL database. Our middleware uses a jabber-based publish/subscribe protocol [Jabber] as a lightweight, platform and language-independent method to push data to client tools (e.g., the visualization, data mining and activity recognition algorithms) with minimal overhead and maximal flexibility. To maintain privacy we remove participant names and identifying information and encrypt collected data before it is transmitted over the network.

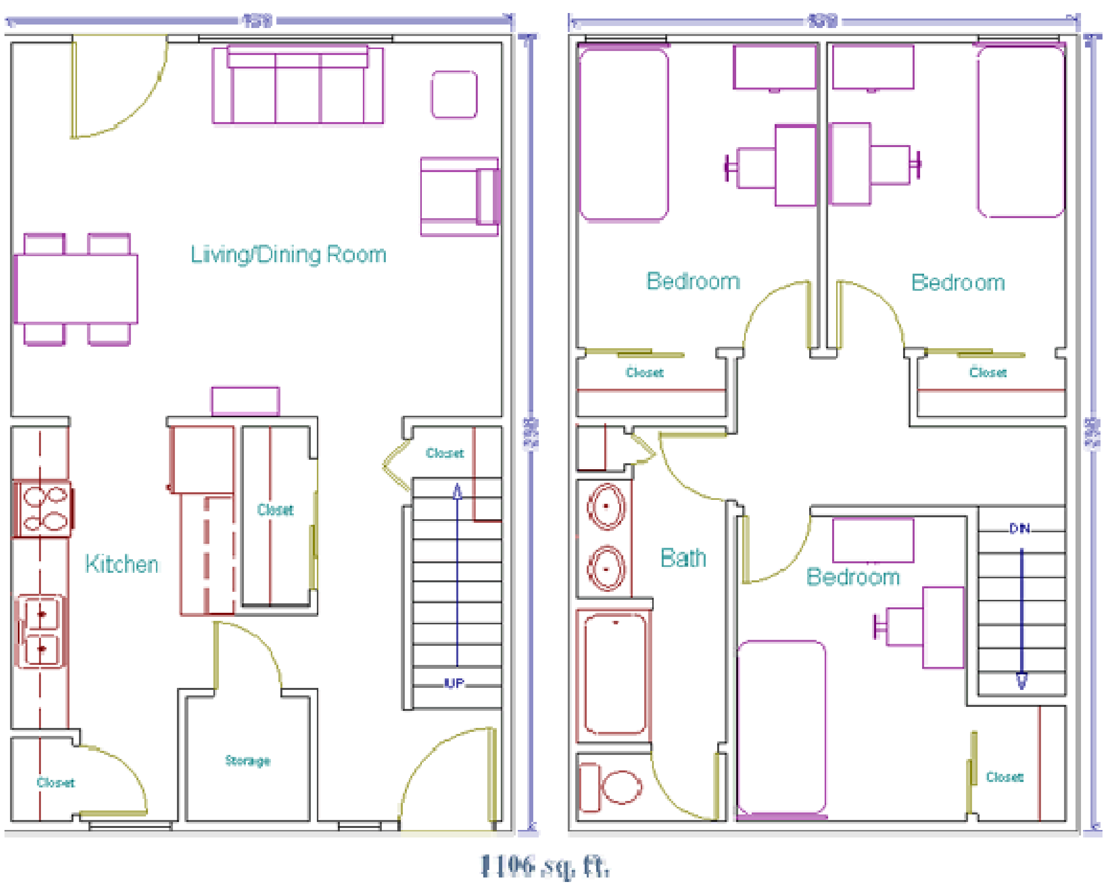

Figure 1. Three-bedroom smart apartment used for our data collection.

To provide physical training data for our activity recognition algorithms, we brought WSU undergraduate students into the smart apartment. Only the first floor of the smart apartment was used for these data collection. We generated five different data collections in all, as described in the following sections. This data is available at ailab.eecs.wsu.edu/casas/research.html [Singla et al., 2008]).

2.1. Complete Activity Sensor Dataset

The purpose of the first *normal* data collection was to generate physical sensor data that would allow us to test the accuracy of our activity recognition algorithms on complete and correct data. For this data collection, we brought one participant at a time into the environment and asked them to perform a sequence of five activities:

1. *Telephone Use*: This activity required participants to look up a specified number in a phone book, call the number, and write down the cooking directions given on the recorded message. The phone book, notepad and telephone were located on the dining room table.
2. *Hand Washing*: For this activity, participants were told to wash their hands in the kitchen sink using the soap and paper towels provided.
3. *Meal Preparation*: This activity required participants to boil water on the stove and cook oatmeal according to the recorded directions, which also specified the addition of brown sugar and raisins. The materials and utensils needed for this task were located in an identified kitchen cabinet and on the kitchen counter.
4. *Eating and Medication Use*: For this activity, participants were asked to pour themselves a glass of water from the facet, and then take the oatmeal, glass of water and medicine container to the dining room table where they were to eat and take medication. The medicine bottle was located in the same kitchen cabinet as the materials for the cooking task.
5. *Cleaning*: The cleaning activity required participants to clean the dishes and put the medicine bottle and other materials back in the cabinet.

```
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 2. Resident performing "hand washing" activity (left). This activity triggers motion sensor ON/OFF events as well as water flow sensor values (right).

The selected activities include both instrumental and basic ADLs. These ADLs are typically found in clinical questionnaires assessing everyday functional activities [Reisberg et al., 2001] and deficits in these ADLs can help identify individuals who are having difficulty living independently at home [Schmitter-Edgecombe et al., 2008].

Sensor information was collected from the downstairs rooms of the apartment and included motion, temperature, water, burner, phone usage (for completed calls), and item sensor readings. Each sensor entry is tagged with the date and time of the event, the ID of the sensor that generated the event, and the sensor reading. Figure 2 shows a person who is performing the "wash hands" activity in the smart apartment as well as a portion of the sensor events that are generated by this activity. Data was collected continuously from the beginning of the first activity to the end of the last one. This dataset contained a total of 5,312 sensor events.

2.2. Erroneous Activity Datasets

While recognizing activities is useful for evaluating the everyday functional activity performance of individuals, another important aspect of activity profiling is determining the successful completion of these activities. To support successful completion of everyday ADLs, multiple cognitive and physical skills are needed. For example, memory and executive functioning impairment have been identified as the most significant cognitive symptoms limiting autonomy in complex daily activities for persons in the early stages of dementia [Schmitter-Edgecombe et al., 2008; Tuokko et al., 2005]. One way in which such impairment manifests itself is in the inconsistent and incomplete performance of daily activities. An individual with cognitive difficulties may take longer than normal to perform an activity, may have difficulty initiating performance of an activity, may perform the steps out of order or incorrectly, or may leave the activity unfinished.

In order to provide an accurate functional assessment of individuals in their own environments, we need to not only track the execution of ADLs but also to monitor how consistently and completely they are performed. To provide data that will be useful in evaluating our algorithms, we brought an additional 20 undergraduate students into the apartment one at a time and asked them to perform the same sequence of 5 ADL activities as the normal group. For each activity, however, we selected a step from the activity to be skipped or performed incorrectly. These mistakes were selected to reflect common difficulties that can compromise everyday functional independence. Because activity errors were purposefully introduced into the performance of these activities, we refer to this dataset as the *specific error* data. Detecting such activity errors is important if we are to accurately monitor functional performance, to intervene if the error creates a hazardous situation (e.g., leaving the stove on) and to provide reminders that will enable residents to complete their daily activities more consistently and independently.

For the next 20 participants, we provided general descriptions of task completion errors and asked participants to simulate someone having difficulties with each ADL. To help guide participants with this undertaking, we first provided a scenario that included several examples of difficulties that an individual suffering from Alzheimer's disease might experience when completing everyday tasks (e.g., leaving appliances on, getting sidetracked or distracted during task completion, taking a long time to complete tasks). Participants were told to keep the scenario examples in mind as they completed the experiment. In this condition, labeled *simulation*, the types of errors that were introduced for each task were participant generated and were extremely varied.

- *Telephone Use*:
  - *Specific Error:* Dial a wrong phone number before retrying and successfully reaching the recorded message.
  - *Simulation:* Simulate someone who is having difficulty using the phone book and recording information about the recipe.
- *Hand Washing*:
  - *Specific Error:* Leave the water running after washing hands.
  - *Simulation*: Simulate someone who gets confused and becomes stuck completing the task.
- *Meal Preparation*:
  - *Specific Error*: Leave the burner on after cooking the oatmeal.
  - *Simulation*: Simulate someone having difficulty with the timing and order of steps involved in the task.
- *Eating and Medication Use*:
  - *Specific Error*: Forget to take medication with the meal.
  - *Simulation:* Simulate someone completing the task steps in a slow and inefficient manner.
- *Cleaning*:
  - *Specific Error*: Wipe off the dishes without using running water to clean them.
  - *Simulation*: Simulate someone who becomes confused and moves off task and is later directed back to the task.

2.3. Interweaved Activity Dataset

Most activity recognition algorithms have been tested in situations where a single individual is performing activities in an uninterrupted, sequential manner. Our next dataset increases the complexity of the activity modeling task by collecting smart environment sensor data while the participant interweaves ADL activities. For this study we brought the participants into the environment one at a time and asked them to perform a new set of 7 ADL activities including filling a medication dispenser, filling out and addressing a birthday card, selecting an outfit for a job interview, watching a DVD, watering the plants, answering a phone call, a sweeping the floor. We first asked the participants to perform each activity separately then asked them to perform the set of activities a second time, ordering and interleaving the activities in any manner that was comfortable to them.

The participants took maximum advantage of the opportunity to interweave activities. In fact, many participants performed several activities concurrently, such as talking on the phone while watching the DVD and watering the plants. A total of 280 data sets were generated, representing 7 activities, each performed twice by 20 participants.

2.4. Multiple Resident Activity Data

Our final dataset focuses on a different complexity issue for home behavior, that of recognizing activities when multiple residents co-exist in a space. For this study, two undergraduate students lived in our smart apartment for 8 weeks during the summer term. We continuously collected data while the students lived there. During this time we recorded data collected from both floors of the apartment. This was useful in tracking the residents for the purpose of not only recognizing activities but also determining which resident triggered the sensor event.

3. Challenges in Generating and Disseminating Home Behavior Data

Our experiences in generating smart home data have provided us with insights on the challenges of such physical data collections. Some of the particular challenges we faced include the following:

- Ensuring clean data. Any type of failure in the environment would invalidate the data that was collected. This included a failed sensor, network issues, camera failure, or database crashes. This type of setback dramatically increases the time required to collect data, which is likely one of the many reasons that few physical data collections exist.
- Annotating the data. In order to train a learning algorithm, we need to not only collect sensor data but need to correctly label it with the class we are trying to learn. In our case, this means annotating sequences of sensor data with the corresponding activity that is being performed. Other approaches to this problem require that the participant self-annotate the data by recording the activity they are performing or that the participant perform the exact set of activities that are requested. Neither approach is practical for deployment in homes of elder adults. Annotating the data by analyzing the sensor data is another approach. However, this is very time consuming and subject to annotation errors. This continues to be a challenge for our ongoing work.
- Generating sufficiently varied data. The reason that we created the datasets that we have described here is to ensure that our algorithms are robust before we deploy them into the homes of elder adults. However, the datasets have certainly not captured all of the variations that will be encountered in everyday settings. Challenges will arise when the technology is moved to different homes, with different sizes and layouts, different family dynamics, and different daily routines. While we would like to create datasets that characterize these variations, this is not practical. Collecting smart home data is time consuming and creates large volumes of raw sensor data. Large data repositories and tools to clean and process the data are needed before such data collections can be effectively maintained and used.

We feel that generating and disseminating smart home datasets is very important if we want to create robust, usable smart environment technologies. We also feel that creating such public datasets will foster collaboration and improve technology evaluation. Several research groups have already downloaded the CASAS datasets for use in their own research projects. However, a much larger, collaborative effort is needed in order to generate sufficient data for the rigorous testing of smart environment technologies. This level of testing will help us transition smart environment technologies from the lab into the homes of the individuals that need them.

References

[Cook and Das, 2004] D. Cook and. S. Das, editors. *Smart Environments: Technology, Protocols, and Applications*. Wiley, 2004.

[Doctor et al., 2005] F. Doctor, H. Hagras, 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(1):55-56, 2005.

[Mihailidis et al., 2004] 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(1/2):135-171, 2004.

[Philipose et al., 2004] 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(4):50-57, 2004.

[Reisberg et al, 2001], B. Reisberg, et al. The Alzheimer's disease activities of daily living international scale (ASL-IS). *International Psychogeriatrics, 13,* 163-181, 2001.

[Schmitter-Edgecombe et al., 2008] M. Schmitter-Edgecombe, E. Woo, and D. Greeley. Memory deficits, everyday functioning, and mild cognitive impairment. *Proceedings of the Annual Rehabilitation Psychology Conference*, 2008.

[Singla et al., 2008] G. Singla, D. Cook, and M. Schmitter-Edgecombe. Incorporating temporal reasoning into activity recognition for smart home residents. *Proceedings of the AAAI Workshop on Spatial and Temporal Reasoning*, 2008.

[Tuokko et al., 2005] H. Tuokko, C. Morris, and P. Ebert. Mild cognitive impairment and everyday functioning in older adults. *Neurocase*, 11:40-47, 2005.

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