Object-Use Activity Monitoring: Feasibility for People with Cognitive Impairments

Mark R. Hodges
Computer Science and Engineering
University of Michigan
Ann Arbor, MI, USA
hodgesm@umich.edu

Mark W. Newman and **Martha E. Pollack**
School of Information
University of Michigan
Ann Arbor, MI, USA
{mwnewman, pollackm}@umich.edu

Abstract

A significant amount of prior work has addressed the question of whether wireless sensor networks can collect sufficient data about how people interact with everyday objects to support accurate interpretation of their behavior. One commonly discussed application for object-use-based activity monitoring is the development of systems to support people with cognitive impairment. However, most prior studies have involved only cognitively intact subjects. In this paper, we report on an experiment in which we attempt to replicate an earlier study that used object-use activity monitoring to identify individuals, but we used subjects with traumatic brain injuries (TBI). We describe the work, which is currently ongoing, and provide preliminary answers to the key questions of whether this approach to activity monitoring is feasible for people with TBI, and whether the data collected provides a level of identification accuracy equivalent to that of the unimpaired subjects.

Introduction

There has been a great deal of prior work studying the use of wireless sensor networks to interpret the behavior of people as they perform daily activities. A particularly successful paradigm has been to model activity based on object usage (Philipose et al. 2004; Tapia, Intille, and Larson 2004; Wang et al. 2007). The intuition here is that an effective way to understand what someone is doing is to collect and analyze information about the objects with which they are interacting. A recent study by Hodges and Pollack (Hodges and Pollack 2007), two authors of this paper, exploited the regularities between object usage and activity performance in a different way: given data about object use collected during the performance of a *known* activity, they inferred the identity of the person performing the action, calling the regularities between a person's interaction patterns and their performance of an activity their "object-use fingerprint." They showed that even with very simple machine-learning techniques, they could identify subjects about three-quarters of the time, a rate of success that was well above chance.

While much of the prior work has been fundamental in nature, one commonly discussed motivation is to support the development of applications for people with cognitive impairments, for example, to give them guidance about when and how to perform their necessary activities of daily living (Mihailidis et al. 2007; Pollack 2005) or to gather information about how they are doing in these activities in order to provide their therapists and caregivers with a more accurate picture of their current status (Haigh and Kiff 2004). However, only a few studies to date have actually used object-based activity monitoring methods with people who are cognitively impaired (Mihailidis et al. 2007; Levinson 1997), and thus significant questions remain about the feasibility of methods for use in this target population. In this paper, we we report on an experiment in which we attempt to replicate the Hodges and Pollack (henceforth, HP) study, but we used subjects with traumatic brain injuries (TBI). After providing a brief review of the related literature, we describe our experiment and the data analysis, both of which are currently ongoing.

The initial HP study used a tethered RFID reader, attached to a glove, to obtain object-use information by detecting proximity to RFID tags attached to kitchen objects. That technology was preliminary, and infeasible for the current study, where it was important to avoid the possible tripping danger of a wire. We therefore replaced the tethered glove with a wireless RFID reader, and completed an initial study to validate that the change in technology does not affect the accuracy results.

After discussing the comparison of the two technologies, we turn to our central question, which is whether it is feasible to use object-use-based activity monitoring methods with TBI patients. We ask both whether the overall approach is feasible, i.e., whether patients will tolerate the apparatus, and whether the data obtained with TBI patients provides the same level of accuracy of identification as that obtained with unimpaired people. Our answer to the first question is a resounding yes: we have already collected data from 24 subjects. However, conditions that arose in the collection of the data led us to have to modify the everyday task being performed, and this has a significant impact on the second question. Specifically, the modifications to the task result in much lower accuracy for *both* the unimpaired and the cognitively impaired subjects, and it is not yet clear whether there are significant differences in accuracy across the two groups at these levels.

Background

Human activity recognition is an active field of research that uses various sensors to collect data about what people are doing, and applies algorithms to process that raw data into information describing their activities. The level of abstraction of the recognized activities may vary, from motor movement, e.g., recognizing whether a person is jumping or walking (Ben-Arie et al. 2002); to low-level goal-directed activity, e.g., distinguishing amongst taking medication, making cereal, or eating cereal (Pentney et al. 2006); to higher-level goal-directed activity, such as identifying a person's likely destination as she moved about a city (Liao, Fox, and Kautz 2006).

The types of sensors that are used in such studies also vary widely, and may include everything from video cameras (Ben-Arie et al. 2002), to light, barometric pressure, and temperature sensors (Lester et al. 2005) to GPS (Liao, Fox, and Kautz 2006). One increasingly common approach, however, especially for the identification of low-level goal-directed activities, is to use sensors that detect direct interaction with objects, such as RFID readers (Fishkin, Philipose, and Rea 2005a), contact switches (Tapia, Intille, and Larson 2004), and electric current and water flow detectors (Logan et al. 2007). In each case, the sensors measurements provide a proxy for object usage: RFID readers, for example, measure proximity of a person's hand to an object, while accelerometers measure the movement (presumably by the person) of an object. Fairly standard machine learning methods are generally used to interpret the collected data (Philipose et al. 2004; Tapia, Intille, and Larson 2004; Fishkin, Philipose, and Rea 2005b; Logan et al. 2007). Recently, common sense databases have been developed to support the interpretation process via the use of common-sense knowledge of context (Pentney et al. 2006; Wang et al. 2007; Wyatt, Philipose, and Choudhury 2005).

In a study reported last year, Hodges and Pollack asked a related question: whether an individual–rather than an activity–could be identified based on object-usage observations. As noted above, they obtained a high level of identification accuracy using a simple decision-tree algorithm on data collected while subjects performed the task of making coffee. Further details of the HP study are provided later in this paper.

There are a number of possible applications for activity monitoring, but one that is frequently mentioned is the development of systems to support people with cognitive impairment. Systems are planned to encompass a range of goals, such as assisting people in navigating through their communities to reach common destinations such as work or the grocery store (Liao, Fox, and Kautz 2006); providing them with reminders of important daily tasks such as taking medicine (Pollack 2005); guiding them through the multi-step activities such as hand-washing or meal preparation (Mihailidis et al. 2007); or ensuring safety by providing reports to caregivers and/or triggering an alarm in dangerous situations, such as when a dementia patient starts to leave her home (Haigh and Kiff 2004). Another important goal is to provide improved reports on the functional behavior of people with cognitive impairment to their professional caregivers (Glascock and Kutzik 2000), noting in particular when their behavior begins to deviate from baseline. The original HP study is relevant to this latter type of system, because the ability to recognize individuals is a partial indication of the ability to determine a baseline from which there may be identifiable changes.

Experimental Setup

The original HP study involved ten unimpaired subjects, who were asked to make a cup of coffee in a University laboratory setting, in which the objects used–e.g., the spoon, coffee, coffee-maker, etc.–had RFID tags attached. The subjects all wore a glove with an attached RFID reader, which was connected to a computer by a wire; they each performed the coffee-making task ten times, mostly on different days.

In the current experiment, our goal was to collect similar data from patients who had been diagnosed with traumatic brain injury (TBI) and were undergoing treatment at a local rehabilitation clinic. While the original study used a wired RFID glove, for this study we used a more technologically advanced device: an "iBracelet" constructed by Intel Research Seattle (Smith et al. 2005). It was very important to use the iBracelet in place of a wired device with the TBI patients, to avoid tripping hazards. However, this necessitated our determining the extent to which the original HP results might be influenced by the shift in technology. To this end, in addition to collecting data from TBI patients, we recruited additional unimpaired subjects, to perform the coffee-making task, in the kitchen laboratory, but wearing the iBracelet.

To date, we have collected data from five additional unimpaired subjects, each of whom performed the coffee-making task five times, on five different days, while wearing the iBracelet. Data collection is ongoing, and we anticipate obtaining data from at least five additional unimpaired subjects. In addition, we are in the process of collecting data from 25 TBI patients and this paper presents the results based on 11 of them, each of whom performed the task five times on five different days [1].

All subjects, both unimpaired and TBI patients, were given a brief tour of the instrumented environment before their first trial, and those who did not know how to make coffee were given basic instructions. These instructions were as general as possible. For example, subjects were told to "put water in the reservoir at the back of the coffee-maker," rather than indicating exactly how the water should be put there, so that they would choose for themselves whether to use the coffee cup or coffee carafe to transport the water from the sink. No physical demonstration of the coffee-making process was given.

Note that the TBI patients performed the coffee making tasks in a kitchen at the rehabilitation clinic that is used for occupational therapy. We endeavored to match as closely

[1] We have finished collection with 24 of the 25 patients and 15 of the 24 patients completed the study by performing the task five times. These results are based on 11 of those 15 patients because the remaining 4 patients finished their participation too recently for their results to be included.

as possible the physical environment in both settings (the kitchen laboratory used by the unimpaired subjects and the occupation therapy kitchen used by the TBI patients). However, there were some differences that nonetheless emerged, and we had to control for them in our data analysis.

First, in the kitchen laboratory, the coffee and supplies were stored in a cabinet under the coffee maker. However, the therapists overseeing the care of our TBI subjects advised us that some of them might have trouble completing the tasks if the needed supplies were out of sight, and thus, in the occupational therapy kitchen, we put all supplies on the counter next to the coffee maker.

Second, because of the therapists's concerns about patients possibly burning themselves, we stopped the trials with TBI patients at the point at which they started the coffee maker; in contrast, the unimpaired subjects continued through the point of pouring the coffee. The result of these changes is that we have significantly less data per trial with the TBI patients than we do with the unimpaired patients (and similarly, than was present for the subjects in the original HP study).

Finally, the original HP study collected data on the use of sugar and creamer but ignored this data in the analysis, because it was considered "cheating" by making the problem too easy; we do the same in this study (i.e., we collect the data about sugar and creamer use, but exclude it from our analysis). However, the original study did take into account information about whether the subjects chose to grind their own beans or use pre-ground coffee; because 90% of them ground their own beans, this was not considered to be a problem. In the current study, with both the unimpaired and the TBI patients, we avoided this issue by providing only pre-ground coffee.

In the next section, on Data Analysis, we explain how we control for each of these three differences.

In sum, there are three sets of data that we analyze: the data originally collected in the HP study with unimpaired users wearing a tethered RFID glove; the additional data we are collecting with unimpaired users wearing the iBracelet; and the data we are collecting with TBI patients wearing the iBracelet. These data sets are listed in Table 1. Recall that while we have at present collected data from 24 TBI patients, as mentioned above, in the current report on work in progress we are making use of only the data from the first 11 to complete the study because data from the other 13 is either incomplete or extremely recent. Table 2 summarizes the objects that had RFID tags attached, noting where there are differences across the data sets.

Data Set	Glove-Control	Bracelet-Control	Bracelet-TBI
From	Hodges-Pollack	This study	This study
# Subjects	10	5	11
# Trials/Subject	10	5	5
TBI Subjects	No	No	Yes
RFID Reader	Wired glove	Bracelet	Bracelet
Full/Short Trial	Full	Full	Short

Table 1: Description of the Data Sets

Coffee Maker
Coffee Carafe
Mug
Spoon
Cabinet (University lab only)
Coffee Grounds
Whole Beans (HP study only)
Filters
Grinder (HP study only)
Creamer (University lab only)
Sugar (University lab only)

Table 2: List of Tagged Objects

Data Analysis

We performed the same data analysis as that done in the original HP study, (Hodges and Pollack 2007), computing several types of features and learning decision trees from them using the C4.5 decision tree algorithm. Here we briefly review the features developed in the HP study; more detail can be found in the original paper. We also describe the methods we used to control for the differences amongst the data sets.

Observation Granularity

Following the original study, we make use of a notion of observation granularity, which refers to the level of abstraction of an observation. For example, when a subject interacts with a particular tag, we obtain information not only about tag itself (e.g., Tag # 14, which may be the left-most one on the lid of the coffee-maker), but also about the larger object to which it is attached (in this case, the coffee-maker). We considered three levels of abstraction:

1. Tag: Interaction with an individual tag affixed to an object.
2. Group: Interaction with any of a group of tags that are equivalent except for the orientation of the object (e.g., the tag on the left side of the coffee grounds and the tag on the right)
3. Object: Interaction with an object; that is, any of the tags on the object were detected.

At times, these levels may be functionally equivalent. The coffee carafe, for example, has only a single tag, so there is no difference between tag, group, and object interactions on the carafe. For other objects, like the coffee maker or the container of coffee grounds, the use of multiple levels of granularity allows for different patterns of interaction to be detected.

Feature Type

At each level of granularity (tag, group, or object), we measure five different types of features, again following the original study:

1. Detected: A binary feature that is positive iff there was any interaction with an “entity,” i.e., a tag, a group or an object.
2. Count: A scalar feature that records the number of times interaction with an entity was observed.
3. Total Duration: A scalar feature that records the total amount of time interaction occurred with an entity.
4. Average Duration: A scalar feature representing the average time of interaction with an entity. (Note that this is a computed feature, equal to the Total Duration divided by the Count.)
5. Order: A binary feature that is positive iff an arbitrary two- or three-entity ordering is observed. Note that orderings are defined only within a single level of granularity: for example, we track whether there is a sequence of interactions with, say Tag #14 and then Tag #20, or with the coffee maker and then the spoon, but we do not have a feature representing an interaction with Tag #14 followed by one with the coffee maker.

Minimizing Differences Amongst Data Sets

In order to control for the differences in experimental set up described in the previous section, we did some systematic pre-processing of our data.

First, to correct for the fact that some objects were only available in some of the data sets (see Table 2), whenever we compared two data sets, we ignored any sensor data that was only available in one of them. This means that when we compared the *Glove-Control* and *Bracelet-Control* data sets (see Table 1), we omitted triggering of sensors on the grinder and the whole beans; when we compared the *Bracelet-Control* and the *Bracelet-TBI* data sets, we omitted the data that indicates triggering of the cabinet, creamer, and sugar.

To control for the shorter trials of the TBI patient, resulting from their stopping at the point at which they turned on the coffee maker, we truncated the data collected from the control subjects to match. Specifically, when we compare the *Bracelet-Control* and the *Bracelet-TBI* data sets, we truncate the trials in the former at the start of the longest gap in sensor triggers, which, we assume, correlates with the point at which the subject turns on the coffee. As part of this ongoing study, we plan to verify that assumption.

Finally, we need to adjust for the fact that we have a different amount of data in each data set: in the *Glove-Control* data (the original HP data), we have ten subjects who each performed ten trials; in the *Bracelet-Control* data we have five subjects, with five trials each; and in the *Bracelet-TBI* data we have (at the time of the analysis) eleven subjects with five trials each. In order to make fair comparisons, when we use the *Glove-Control* and the *Bracelet-TBI* data, we thus randomly select five subjects, and five trials (for *Glove-Control*) for each. To ensure that we have not skewed the data, we repeat the random selection of subjects and trials a total of 100 times and report the average result.

Table 3 summarizes our resulting comparisons. The first data column show the parameters of the original HP data.

Data Set	G-C	G-C	B-C	B-C	B-T
# Subjects	10	5*	5	5	5*
# Trials/Subject	10	5*	5	5	5
Full/Short	Full	Full	Full	Short*	Short
RFID Reader (Glove/Bracelet)	G	**G**	**B**	B	B
TBI	No	No	No	**No**	**Yes**

Table 3: Description of Analyses. Attributes being compared are in bold, attributes being adjusted from their original data set are marked with an asterisk; the adjustments are discussed in the data analysis section.

The second and third data columns show the data that we used when comparing *Glove-Control* and *Bracelet-Control* to determine the impact of the technology change on overall accuracy. As can be seen, the main difference is that we select five trials from five subjects from *Glove-Control* to match data-set size (and as just described above, repeat this process one hundred times). The last two data columns show what we used when we compare *Bracelet-Control* and *Bracelet-TBI*: here we are selecting 5 subjects from the *Bracelet-TBI* set (100 times, but using all 5 trials from the 5 selected subjects), and we are truncating the data from each trial in *Bracelet-Control* to mitigate the fact that the TBI trials ended with starting the coffee-maker.

Results

We first briefly provide the results from the original HP study, as a baseline. We then compare the accuracy of identification of control (unimpaired) subjects wearing the tethered glove and the wireless bracelet, and finally turn to the result of the experiment with TBI patients.

Original Hodges-Pollack Data

Table 4 shows the results obtained in the original HP study, according to the set of features used for classification. The results are an average of a ten-fold cross validation. Note that accuracy ranges from a low of 58% when only the average time per detection for each entity is considered, to a high of 77% when all of the features except order are used. The original paper presents a hypothesis about why omitting order helps; in short, the belief is that inclusion of the order feature leads to overfitting. Note that with ten subjects, performance at chance would be 10%.

Effect of the Technology Change

We next consider whether the use of the iBracelet, in place of the RFID glove, has an impact on accuracy. The iBracelet has a major advantage in not being tethered. In addition, it has a longer range, meaning fewer tags need to be used to reliably detect interaction with an object–conversely, however, with fewer tags there is a less information about exactly where on an object an individual is touching. As a result, there may potentially be a decline in the identification accuracy with the iBracelet relative to the glove.

Table 5 presents the result of our comparison of the *Glove-Control* and the *Bracelet-Control* data sets. As described

Features Used	Accuracy
Full Feature Set	73%
All But Order	77%
Object-level	69%
Group-level	72%
Tag-level	72%
Detected	59%
# Detections	65%
Tot. Time	75%
Avg. Time	58%
Order	70%

Table 4: Accuracy Using Original HP Data with 10 Subjects and 10 Trials per Subject (Hodges and Pollack 2007)

Features Used	Glove-Control Accuracy	Bracelet-Control Accuracy
Full Feature Set	68.2%	71.3%
All But Order	64.8%	64.8%
Object-level	71.7%	49.7%
Group-level	69.5%	39.2%
Tag-level	69.9%	69.9%
Detected	59.2%	42.4%
Count	60.5%	68.7%
Tot. Time	63.8%	66.3%
Avg. Time	58.7%	28.4%
Order	69.4%	55.6%
Object Order	71.9%	44.2%

Table 5: Accuracy Using Glove vs Bracelet as RFID Reader

earlier, we obtained these results by randomly selecting five subjects and, for each, five trials from the *Glove-Control* data (i.e., the original HP data), and then performing five-fold cross validation, repeating this overall process 100 times. In the case of the *Bracelet-Control* data, we only have five subjects with five trials each, and so here we simply use all the data in a five-fold cross-validation experiment. Because order is a computationally expensive feature to compute, as addressed in the original HP study, we include the additional feature set of order at the object level only, because it is faster to compute with a reduced number of entities.

We stress that this is very preliminary data: with only five subjects it is very difficult, if not impossible, to draw general conclusions. Indeed, the variance across folds in the validation process is quite high. Nonetheless, so far it appears that the differences across technology are not large: with the glove we are able to identify subjects with nearly 72% accuracy, and with the iBracelet, we achieve just over 71% accuracy, albeit using different feature sets. Additional analysis will be done as we expand the size of the *Bracelet-Control* data set.

Identification Accuracy with TBI Patients

Finally, we turn to the question of using object-use activity monitoring with cognitively impaired patients. Our first question is a qualitative one: is it reasonable to assume that such patients will wear a sensor bracelet while conducting an activity of daily living? Our answer to this is an unqualified "yes," at least for TBI patients: we have successfully collected data from 24 subjects (15 of whom completed the study and 22 of whom made coffee at least three times).

Our second question was whether identification accuracy with the TBI patients would be comparable to that of the unimpaired subjects. Unfortunately, we are unable to answer that question directly, for reasons we will describe in a moment; reasons which, however, taught us something else interesting.

Table 6 summarizes the comparison of the *Bracelet-Control* and the *Bracelet-TBI* data sets, again, with the controls described in Table 3; we again report on the average of five-fold cross validation, where for the *Bracelet-TBI* case these are averaged over 100 random selections of five subjects.

The main thing to note is the significant drop in accuracy relative to the prior comparison: instead of topping out at levels over 70%, we are now able only to achieve accuracy below 50%. This is still better than chance, which is only 20%, but not sufficient for any real system. Moreover, the variance across trials is again quite high, making it difficult to assess the reliability of the results.

The cause of the drop in accuracy is clear: by removing the data that indicates contact with the cabinet, as well as the sugar and creamer, and also removing all data that occurs after the coffee maker is started, we end up with significantly less information about each subject. Note that the modifications in the data were a direct result of advice given to us by the TBI patients' therapists about what types of tasks were appropriate for them to perform. The lesson is an important one for designers of systems for people with cognitive impairment: task selection is critical, and we cannot assume that the same tasks that are reasonable for study in unimpaired people are also always appropriate for another population.

Discussion

We have presented preliminary results of an ongoing study aimed at determining whether object-use activity monitoring can be used in systems designed for people with cognitive impairment. Toward that end, we replicated the study of Hodges and Pollack (Hodges and Pollack 2007), using traumatic brain injury patients as subjects.

As an initial step, we examined the differences in identification accuracy between analyses of data collected with the tethered RFID glove used in the original study and analyses of data collected with a wireless iBracelet. We do not yet have sufficient data to make definitive claims, but the preliminary results show that accuracy levels stay approximately the same, despite differences in the ranges of the two devices.

We found that TBI patients are able to perform activities

Features Used	Bracelet-Control Accuracy	Bracelet-TBI Accuracy
Full Feature Set	44.4%	38.4%
All But Order	43.7%	35.4%
Object-level	41.1%	41.8%
Group-level	41.1%	38.2%
Tag-level	47.7%	33.1%
Detected	37.8%	30.2%
Count	45.8%	36.8%
Tot. Time	47.9%	34.0%
Avg. Time	27.9%	37.2%
Order	46.3%	31.6%
Object Order	44.8%	40.4%

Table 6: Accuracies Identifying Unimpaired Individuals and Individuals with TBI from Different Feature Sets

while wearing an iBracelet without difficulty. However, it turned out that we had to modify the activity performed–making coffee–in some important ways to make it feasible for most TBI patients, and the result was a loss in data that significantly decreased our ability to perform identification. We note that when we manipulated the data from the unimpaired subjects to control for the changes in task activity as performed by the TBI subjects, we saw a similar drop in identification accuracy. Thus, our conclusion–and again, it is a preliminary conclusion pending the collection of more data–is not that it is more difficult to identify TBI patients than to identify control subjects based on their object-use patterns *per se*. Rather, we note that it may be more challenging to select appropriate tasks for activity monitoring amongst cognitively impaired people. As a result, our next steps include both investigating the use of more powerful learning algorithms, which can potentially do a better job with the impoverished data set than the basic C4.5 approach, and to evaluate other types of everyday tasks that may provide richer data.

Acknowledgements

The authors would like to acknowledge Intel Research Seattle who donated the iBracelet for use in this study.

References

Ben-Arie, J.; Wang, Z.; Pandit, P.; and Rajaram, S. 2002. Human activity recognition using multidimensional indexing. *IEEE Trans. Pattern Anal. Mach. Intell.* 24(8):1091–1104.

Fishkin, K. P.; Philipose, M.; and Rea, A. 2005a. Hands-on rfid: Wireless wearables for detecting use of objects. In *ISWC '05: Proceedings of the Ninth IEEE International Symposium on Wearable Computers*, 38–43. Washington, DC, USA: IEEE Computer Society.

Fishkin, K. P.; Philipose, M.; and Rea, A. 2005b. Hands-on rfid: Wireless wearables for detecting use of objects. In *ISWC '05: Proceedings of the Ninth IEEE International Symposium on Wearable Computers*, 38–43. Washington, DC, USA: IEEE Computer Society.

Glascock, A., and Kutzik, D. 2000. Behavioral telemedicine: A new approach to the continuous nonintrusive monitoring of activities of daily living. *Telemedicine Journal* 6.

Haigh, K. Z., and Kiff, L. M. 2004. The independent lifestyle assistanttm (i.l.s.a.): Ai lessons learned. *IAAI 04*.

Hodges, M. R., and Pollack, M. E. 2007. An 'object-use fingerprint': The use of electronic sensors for human identification. In Krumm, J.; Abowd, G. D.; Seneviratne, A.; and Strang, T., eds., *Ubicomp*, volume 4717 of *Lecture Notes in Computer Science*, 289–303. Springer.

Lester, J.; Choudhury, T.; Kern, N.; Borriello, G.; and Hannaford, B. 2005. A hybrid discriminative/generative approach for modeling human activities. *Proceedings of International Joint Conference on Artificial Intelligence*.

Levinson, R. 1997. Peat - the planning and execution assistant and trainer. *Journal of Head Trauma Rehabilitation* 12.

Liao, L.; Fox, D.; and Kautz, H. 2006. Location-based activity recognition. In Weiss, Y.; Schölkopf, B.; and Platt, J., eds., *Advances in Neural Information Processing Systems 18*. Cambridge, MA: MIT Press. 787–794.

Logan, B.; Healey, J.; Philipose, M.; Tapia, E. M.; and Intille, S. 2007. A long-term evaluation of sensing modalities for activity recognition. In *Proceedings of Ubicomp 2007*.

Mihailidis, A.; Boger, J.; Canido, M.; and Hoey, J. 2007. The use of an intelligent prompting system for people with dementia. 14:34–37.

Pentney, W.; Popescu, A.-M.; Wang, S.; Kautz, H. A.; and Philipose, M. 2006. Sensor-based understanding of daily life via large-scale use of common sense. In *AAAI*.

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

Pollack, M. E. 2005. Intelligent technology for an aging population: The use of ai to assist elders with cognitive impairment. *AI Magazine* 26:9–24.

Smith, J. R.; Fishkin, K. P.; Jiang, B.; Mamishev, A.; Philipose, M.; Rea, A. D.; Roy, S.; and Sundara-Rajan, K. 2005. Rfid-based techniques for human-activity detection. *Commun. ACM* 48(9):39–44.

Tapia, E. M.; Intille, S. S.; and Larson, K. 2004. Activity recognition in the home using simple and ubiquitous sensors. In *Pervasive*, 158–175.

Wang, S.; Pentney, W.; Popescu, A.-M.; Choudhury, T.; and Philipose, M. 2007. Common sense based joint training of human activity recognizers. *Proceedings of IJCAI 2007*.

Wyatt, D.; Philipose, M.; and Choudhury, T. 2005. Unsupervised activity recognition using automatically mined common sense. *Proceedings of AAAI 2005*.