

# Resident and Caregiver: Handling Multiple People in a Smart Care Facility

Aaron S. Crandall and Diane J. Cook

School of Electrical Engineering and Computer Science  
Washington State University, 99164, USA.  
acrandal@wsu.edu & cook@eecs.wsu.edu

## Abstract

Intelligent environment research has benefited medical care in a number of ways, including emergency detection, comfort and accessibility. However, most of these techniques have been applied in the context of a single resident, leaving out situations where there is more than one person in the living space. A current looming issue for intelligent environment systems is performing these same techniques when multiple residents or care providers are present in the environment. In this paper we investigate the problem of attributing sensor events to individuals in a multi-resident intelligent environment. Specifically, explore and contrast using two different classification techniques. The naïve Bayesian and Markov Model classifiers present different capabilities and features for identifying the resident responsible for a unique sensor event. We present results of experimental validation in an intelligent workplace testbed and discuss the unique issues that arise in addressing this challenging problem.

## Introduction

With the introduction of smart home technologies into homes and care facilities, the possibilities for customizing system behavior have increased dramatically. Significant headway has been made in tracking individuals through spaces using wireless devices (Bahl and Padmanabhan 2000)(Priyantha, Chakraborty, and Balakrishnan 2000)(Yin, Yang, and Shen 2007) and in recognizing activities within the space based on video data (Feng et al. 2001)(Intille 2002)(Snidaro, Micheloni, and Chivedale 2005), motion sensor data (Jakkula and Cook 2007)(Wren and Tapia 2006), or other sources of information (Moncrieff 2007)(Orr and Abowd 2000). However, much of the theory and most of the algorithms are designed to handle one individual in the space at a time. Passive tracking, activity recognition, event prediction, and behavior automation becomes significantly more difficult when there are multiple residents in the environment. Since care of people commonly includes multiple residents, a resident and care provider, or a number of residents and care providers, having the ability to discern individual activities and events is a base requirement for real-world deployment of smart home technologies.

The goal of this research project is to model and automate resident activity in multiple-resident intelligent environments. There are simplifications that would ease the complexity of this task. For example, residents could be asked to wear devices that enable tracking them through the space (Hightower and Borriello 2001)(Yin, Yang, and Shen 2007). This particular solution is impractical for situations in which individuals do not want to wear the device, forget to wear the device, let the device's power source die, or enter and leave the environment frequently. Similarly, capturing resident behavior with video cameras aids in understanding resident behavior even in group settings (Krumm et al. 2000). However, surveys with target populations have revealed that many individuals are adverse to embedding cameras in their personal environments (Intille 2002), although new techniques using silhouette-only imaging has shown itself to be more palatable. As a result, our aim is to identify the individuals and their activities in an intelligent environment using passive and low profile sensors.

To achieve this overall goal, our first step has been to design an algorithm that maps sensor events to the resident that is responsible for triggering said sensor event. This information will allow our algorithms to learn profiles of resident behaviors, identify the individuals currently in the environment, monitor their well-being, and automate their interactions with the environment. Some previous works have focused on passive multi-resident systems (Cook and Das 2007), and give some indication of techniques that have succeeded on real-world data sets for activity recognition (Lu, Ho, and Fu 2007).

To date, the focus has often been on looking at global behaviors and preferences with the goal of keeping a group of inhabitants satisfied (Roy et al. 2005). In contrast, our research is focused on identifying an individual and logging their preferences and behaviors in the context of the multi-resident spaces. This will bring simpler, more private smart home technologies to care facilities, and individual homes that have multiple residents.

The solutions used in this work revolve around using very simple passive sensors, such as motion, contact, door sensors, appliance interaction and light switches to give a picture of what is transpiring in the space. These information sources offer the benefits of being fixed, unobtrusive and robust devices, as well as being commonly found in living

spaces. Examples of the motion detectors and light switches we use in our testbed are shown in Figure 2.

Smart homes and medical care systems are often targeted towards recognizing and assisting with the Activities of Daily Living (ADL's) are used by the medical community to categorize levels of healthy behavior in the home. The ability of smart homes to help disabled and elderly individuals to continue to operate in the familiar and safe environment is currently one of the greatest reasons for their continued development, alongside energy efficiency technologies. So far, most smart home research has focused on monitoring and assisting a single individual in a single space. Since homes often have more than a single occupant, building solutions for handling multiple individuals is vital. Dealing with multiple inhabitants has rarely been the central focus of research so far, as there have been numerous other challenges to overcome before the technology can effectively handle multiple residents in a single space.

Since smart home research has the ultimate goal of being deployable in real-world environments, seeking solutions that are as robust as possible is always a factor in the systems we engineer. With that in mind, building an entirely passive solution gives the advantage of keeping the technology separate from the inhabitants while they go about performing their daily routines. This lets the smart home feel as "normal" as possible to the residents and their guests. By reducing the profile of the new devices as much as possible, people's behavior should be less effected by the technology that surrounds them.

In this paper we present a pair of solutions to part of the problem described above. Specifically, we apply two alternative supervised machine learning algorithms to the task of mapping sensor events to the resident responsible for the event. Each of the algorithms show different kinds of capabilities in determining differences in behavior, which will make them more or less suitable for use in different smart home environments.

The solutions offer the advantage of using previous behavioral data collected from the set of known residents without requiring significant additional actions to be performed by the residents. This historical behavior is used to train the learning algorithms for use in future real-time classification of the individuals and can be updated over time as new data arrives.

Here we present the results of using the naïve Bayesian and Markov Model classifiers to learn resident identities based on observed sensor data. Because the algorithms are efficient and robust, we hypothesize that they will be able to accurately handle the problem of learning resident identities and be usable in a real-time intelligent environment. We validate our hypothesis using data collected in a real smart workplace environment with volunteer participants.

## Data Gathering Environment

The smart home testbed environments at Washington State University consist of a lab space on campus and a town home off campus. These testbeds are part of WSU's CASAS smart environments project. For our study, we used the lab

space on campus, as there are multiple faculty, staff, and students who regularly enter the space and a number of different kinds of activity take place throughout the rooms. The space is designed to capture temporal and spatial information via motion, door, temperature and light control sensors. For this project we focus on events collected from motion sensors and resident interaction with lighting devices. Part of the testbed layout for both sensors and furniture is shown in Figure 1. The rest of the space is very similar with desks, tables and cubicles being the predominate features.




Figure 1: 2D view of inner office furniture and sensor locations.

Throughout this space, motion detectors are placed on the ceilings and pointed straight down, as shown in Figure 2. Their lenses are occluded to a smaller rectangular window giving them roughly a 3'x3' coverage area of the corresponding floor space. By placing them roughly every four feet, they overlap (between a few inches, up to a foot) and allow tracking of an individual when a motion occurs. The

motion sensor units are able to sense when a motion as small as reaching from the keyboard to a mouse. With this level of sensitivity, sensors around work spaces trip even when people sit quietly in a private space to work at a computer.

To provide control and sensing over the lighting, Insteon™ brand switches (similar to X10 devices) are used to control all of the ceiling and desk lights in the room. These switches communicate with a computer and all interactions with them are logged. Figure 2 shows examples of both the motion and light switch sensors.



Figure 2: CASAS sensors: motion detector and Insteon light switch.

The entire lab space, including the portion shown in Figure 1, has two doors with simple magnetic open/closed sensors affixed to them. These record door openings and closings via the same bus as the motion detectors.

By being able to log any major movement through out the space, as well as device interactions, this system captures basic temporal and spatial activity that can be used to identify individuals based on behavior. The tools used in this project are designed to exploit both the spatial and temporal differences, such as personal work spaces and activity times, to accurately classify a given individual.

### Data Representation

The data gathered by CASAS for this study is represented by a quintuple:

1. Date
2. Time
3. Serial Number
4. Event Message
5. Annotated Class (Resident ID)

The first four fields are generated automatically by the CASAS data collection infrastructure. The annotated class field is the target field for this problem and represents the resident ID, to which the sensor event can be mapped.

<i>Date</i>	<i>Time</i>	<i>Serial</i>	<i>Message</i>	<i>ID</i>
2007-12-21	16:41:41	07.70.eb:1	ON	abe
2007-12-21	16:44:36	07.70.eb:1	OFF	abe
2007-12-24	08:13:50	e9.63.a7:5	ON	john
2007-12-24	14:31:30	e9.63.a7:5	OFF	john

Table 1: Example of data used for naïve Bayes classifier training.

Training data was gathered during several weeks in the lab space by asking individuals working in the lab to log their presence by pushing a unique button on a pin pad when they entered and left the space. During post processing, the database was filtered to only use sensor events during the time windows when there was a single resident in the space. The corresponding data for the given time frame was then annotated and supplied as training data to our machine learning algorithm. The total time frame for data collection was three weeks, and over 6000 unique events were captured and annotated as training data. For an example of the resulting quintuples, see Table 1.

### Naïve Bayes Data Representation

Building more complex parsings of the data was done with a number of strategies that were designed to capture the differences in behavior between individuals. Primarily, these strategies revolved around using the data and time information to give the classifier additional information in the form of "feature types", as shown in Table 2. The times that different people work, especially in any kind of care facility, are very helpful in discriminating between residents and care providers.

### Markov Model Data Representation

One of the biggest advantages of using a Markov Model over a naïve Bayes classifier is context. To represent the same data for a Markov Model, it is presented as a series of evidence. In a smart home context, this is given as the series of sensor events that a resident has caused.

For this work, the data set was broken into the event series caused by our subjects. The series was noted with the ID of the person and used to train their personal classifier.

Additional features that could be explored in these kinds of classifiers include more temporal information about time of day and the time it took for an individual to perform activities, both of which will give even greater differentiation between behavior patterns.

### The Classifiers

The classifiers used for this comparison are a naïve Bayes classifier and a Markov Model-based classifier. These kinds of classifiers have been used with great effect in other smart home research projects (Tapia, Intille, and Larson 2004). The two classifiers made use of the same data to accomplish the same goal, but they are able to take into account very different strategies of determining the current resident.

#	Feature Type	Example
1	Plain	07.70.eb:1#ON
2	Hour of Day	07.70.eb:1#ON#16
3	Day of Week	07.70.eb:1#ON#FRI
4	Part of Week	07.70.eb:1#ON#WEEKDAY
5	Part of Day	07.70.eb:1#ON#AFTERNOON

Table 2: Feature types used for classifier training

In the first approach, a simple naïve Bayes classifier was trained, where the features were built from the event information, with the given class as the individual to whom the event is associated with. This required it be distilled to only a single feature paired to a given class. The class is set by the annotation, but the feature chosen can be built from a combination of the fields. The resulting feature to class pairing is used for the classic naïve Bayes statistical classifier based primarily on frequency of occurrence.

For the simplest interpretation, only the serial number coupled with event message was used, see Table 2, row 1. This simple feature set provides a good baseline to compare more complex parsings with. The more complex parsings, such as "Part-of-Week" (ie WEEKDAY or WEEK-END) capture more information about the given behavior, and can give the classifier more information for correct future classifications. Depending on the facets of the data set, different kinds of feature types can give the classifier better or worse results.

The different feature choices available (ie Simple vs Hour of Day, etc.) divide the data up in different ways. Each way captures the behaviors or the residents with varying degrees of accuracy, depending on the feature types chosen and the behavior of the individuals in the data set. The purely statistical nature of a naïve Bayes classifier has the benefit of being fast for use in prediction engines, but lacks the ability to handle context in the event stream that could be advantageous in discerning different behaviors.

Conversely, the simple Markov Model classifier is built around whole series of events, giving it the ability to generate classifications with more context. In this approach, one model was trained for each resident, where the states of the model map directly to the sensors in the system. This direct exposure of the possible states makes this implementation a simple Markov Model.

When given a test series of events, it was applied to each model in turn using the forward backward algorithm, and the model with the highest resulting probability was chosen as the guess. This kind of system requires more calculation time to run than the naïve Bayes solution, but it is highly parallelizable by the number of residents.

We selected Markov models for our second approach because this representation encapsulates additional contextual information. As a result, the context of the sensor event is used when labeling the sensor event. By adding transitions between states in the Markov Model, the spatial and temporal relationships between sensor events are captured. Thus, by taking more of both the physical and the temporal information into account, our smart environment will more

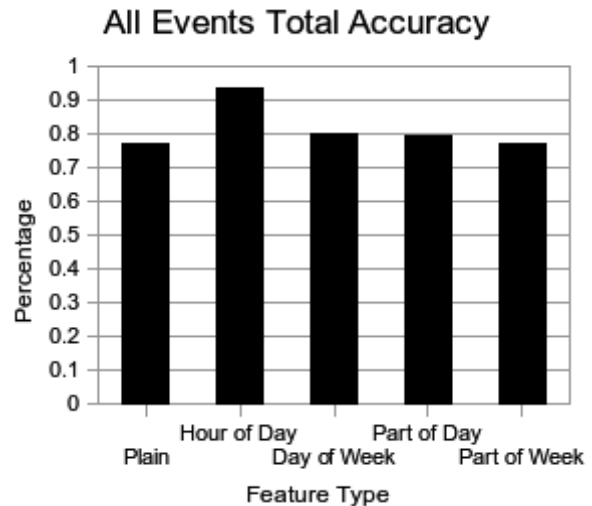


Figure 3: Average accuracy rates by feature type.

effectively scale to a larger number of residents. The classification algorithm will also be able to express and learn a mapping based on more subtle differences in behavior between the residents in the space.

## Results

Figure 3 shows the classification accuracy of our naïve Bayesian classifier for the three residents we tested in our lab space, while Figure 4 gives the results for all three residents over a number of event series sizes.

The most direct comparison between these two classifiers is to take the "plain" feature type for the naïve Bayes and compare it to the middle event series sizes of the Markov Model. Neither one of these take into account information about the time of day, nor the time taken for a given event series. Looking at these two, the naïve Bayes classifier runs about 76%, while the Markov Model is about 84%. This difference is significant when using the classifiers without additional temporal information. It is likely attributed to the Markov Model taking into account more contextual information about the physical layout and ordering of the sensors.

The naïve Bayes classifier can be enhanced to include more temporal information via adding it to the feature format. In particular, we added the date and time of each sensor event, as shown in Table 2. The classifier can now use time of day or day of week information to differentiate between the behaviors of the various individuals. For example, John always arrived early in the day, while Abe was often in the space late into the evening. Finding the correct features to use for this kind of capturing of the behavior can be done by balancing the overall correct rate and false positive rate against one another.

For this data, it was found that using the hour of day provided the best improvements. These are shown in both the overall accuracy (Figure 3), and in each individual's results (Figure 5 and Figure 6).

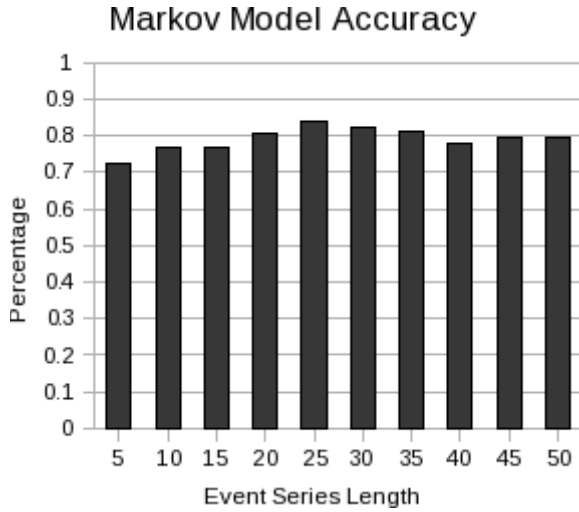


Figure 4: Markov Model classification varied by event stream size.




Figure 6: Charlie's rate of correct classification across feature types.

Note that the classification accuracy is quite high for the John values, but so is the false positive rate (Figure 7). This is because our John participant was responsible for most (roughly 62%) of the sensor events in the training data. As a result, the apriori probability that any sensor event should be mapped to John is quite high and the naïve Bayesian classifier incorrectly attributes Abe and Charlie events to John as well. On the other hand, while Charlie has a much lower correct classification rate, he also has a lower false positive rate.



Figure 5: John's rate of correct classification across feature types.



Figure 7: John's rate of false positives across feature types.

The choice of feature descriptors to use is quite important and has a dramatic effect on the classification accuracy results. Looking at the accuracy rate as effected by the feature type chosen as shown in Figure 3, it shows that using hour-

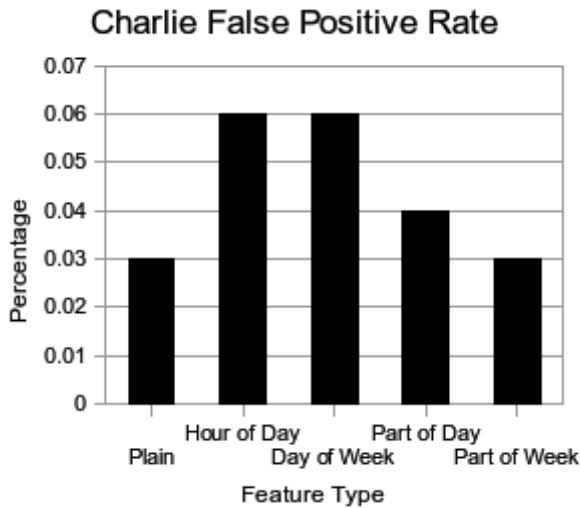


Figure 8: Charlie’s rate of false positives across feature types.

of-day increases the average identification significantly. Additionally, by using hour-of-day, the false positive rate drops dramatically, as shown in Figure 9. When the right features are selected from the data set, the classifier is able to make better overall classifications.



Figure 9: Average false positive rates by feature type.

Choosing the best feature type to pick means balancing the accuracy against the false positive rate. A visual way of showing this kind of balancing is shown in Figure 10. By choosing time-of-day the benefits to the accuracy rate will probably outweigh the increase in false positive rate. In this case, a 2.5x increase in accuracy balances against a 2x increase in false positives. Unless the final application is highly dependent on the certainty of the predictions, it should be a simple algorithm to determine which feature

type is most advantageous.

If the intelligent environment can take false positive rates into account, this information about false positives can be leveraged accordingly via a belief value. As we move towards the use of ensembles of classifiers to build the smart home identification systems, the ability to determine how accurate a result is likely to be will become a necessity. This value will need to be taken into account by any fusion engine we use to make the final identification determination, and could be passed onto any historical data gathered for further training, leading to even the classifiers being built with data giving them a value of believed accuracy.

### Charlie Event Classification




Figure 10: Overall classification rates for all features for Charlie.

With the initial introduction of temporal information, the classifiers begin to see correctness rates of over 93% and false positive rates below 7%. A prediction engine that relies on these classifiers can have a high degree of confidence that it is correctly identifying an individual and attributing events to them. This then leads to proper ADL detection, activity trending, and better information provided to care providers using a low profile, passive smart home infrastructure.

### Time Delta Enhanced Classification

Adding more features to our data set did improve the resident classification accuracy. Due to the nature of the naïve Bayes classifier, the features can only be complexified a little bit before they are no longer useful. To add additional information to the data, short events were left out. The length of an event was calculated to be the length of time from when someone first tripped a given sensor until they tripped another sensor. This had the effect of giving the system less information about “mixed” areas, where people often walk through quickly.

This was based around the fact that the majority of our events were short. The breakdown is shown in Figure 11, and further inspection showed the the majority of the short

events were the ones that confused the naïve Bayes style of classification.

To begin, we removed from our data set any motion sensor events whose durations, or time elapsed between events, fell below two standard deviations from the mean, leaving the longest deltas. With an even more reduced set in hand, the data splitting, training and testing were all done the same way as before with the full data set.




Figure 11: Count of lengths an individual spends on any sensor.

The resulting classifier only used a handful of the available sensors throughout the living space, but the accuracy and false positive rates improved dramatically. This is attributed to the fact that motion sensors in shared spaces or walkways will mostly have very small time deltas associated with them. Since these sensors are also the ones with the highest false positive rates in the full set classifier, removing these sensor events will improve the overall performance of the classifier. Note that with this filtered-data approach, sensor events with short durations will not be assigned a mapping to a specific resident. However, by combining this tool with one that tracks inhabitants through the space (Jakkula, Crandall, and Cook 2007), only a handful of sensor events need to be classified as long as they have a high accuracy. For the naïve Bayes classifier, applying this kind of filtration was very valuable. The accuracy rates exceeded 98%, with correspondingly low false positive rates in some instances (See Figure 12).

This kind of filtration does not assist the Markov Model classifier. The inclusion of short events is valuable to the Markov Model because it uses the context of the whole event series to differentiate between individuals. Removing these events renders the classifier unusable, as there is too little evidence to process.

The Markov Model takes into account either traversing sensors quickly, or remaining on a single sensor by design. This allows the overall architecture of the system to no longer require this first stage of data analysis and modifi-

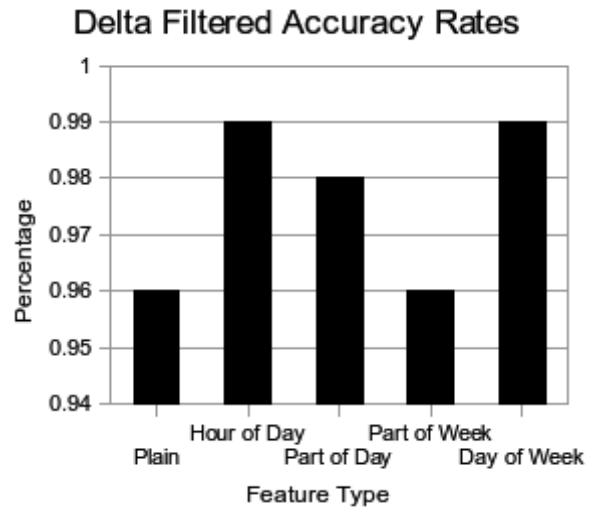


Figure 12: Delta filtered classification accuracy results.

cation for improvements. It also reduces the reliance on external tools to help manage the physical location and tracking of individuals. These kinds of external tools will be required to differentiate streams of events, but the Markov Model classifier will not require them to track and identify nearly as much as the naïve Bayes system that is relying on a form of delta filtering to garner accurate results with.

## Conclusions

In this paper, we design and evaluate two alternative machine learning approaches to identifying individuals in a smart environment. The approaches build upon our earlier work using a naïve Bayes classifier, and introduce the use of a Markov model. Both classifiers learned accurate concept descriptions, but their internal design infuses them with different classification capabilities for this task.

Using a real-world testbed with real-world activity, the classifiers performed well. With simple, raw smart home sensor data the naïve Bayes classifier was showing an average accuracy over 90% for some feature selections. After applying some filtration to the data set to exaggerate the behavior of the inhabitants, accuracy rates over 95% and false positive rates under 2% were possible.

The Markov Model classifier, without the benefit of as much temporal information managed to get near the 90% number, but took into account more of contextual behavior inherently. This provides a much better starting point for future designs, as less external management of the data needs to take place for the system to function.

Additionally, the Markov Model style of classifier should scale better with both the size of the space and the number of inhabitants. With a small home and a limited number of people, the naïve Bayes should be functional enough, but for larger facilities and more individuals, making a system able to seek out more subtle features to classify upon will be of paramount importance.




Figure 13: Delta filtered classification false positive results.

The naïve Bayes requires that different features are shoehorned into place and compared against one another to determine which ones best fit the current historical data set. While this is not as time consuming as many other techniques would be, it still requires more passes and complexity to accomplish. By moving towards a Markov Model solution, we were able to gather a number of these externally built features without additional effort, allowing for more subtle features to be used in the classification.

When using these kinds of classifiers in a medical or security environment, the false positive rate is very important. The system should hold off on making decisions instead of guessing because being incorrect about whether someone is acting correctly or having the system call security without need can lead to not just poor performance, as it would for an energy efficiency or comfort application, but it can begin to impact diagnosis or responsiveness of the emergency response system. These kinds of real world issues impact deep within smart home technology choices.

Choosing the best time-based features can strongly influence the performance of any temporally-dependent environment, and this is no exception. Whether the final application needs a very high level of certainty for one or more of the residents or can trade that certainty off for higher accuracy across all individuals is up to the needs of the final smart home application. Fortunately, as the systems move towards ensembles of tools, the systems should be able to seek a number of features and choose the best between them. Additional benefits are derived by choosing classifiers that take more context and timing information into them by design, instead of relying on complexifying the features given to a simpler classifier, such as naïve Bayes.

Both the naïve Bayes and the Markov Models have advantages and disadvantages. With identifying individuals in a space, there are very subtle physical and temporal facets to seek out and different kinds of classifiers will work best together, keeping their deficiencies mitigated by allowing

them to focus on what kinds of features they are best and leveraging.

## Future Work

To continue to grow the capabilities of these kinds of classifiers, a number of things can help. Additional data with more individuals will show how robust of a solution this is. Differentiating between two people with very similar schedules might be very difficult for this kind of tool. Comparing this tool as a baseline solution with Hidden Markov or Bayesian Network based solutions will allow the continued research to show how much contextual information assists with the classification of individuals.

Applying this classifier to a larger preference and decision engine is a must. Adding this tool to a passive tracking solution will give significantly more information to any individual's history for future learning and prediction systems that are deployed in the CASAS testbed. Comparing it to a system without this kind of identification process, or one based on device tracking, will be a significant step for smart home research.

## Acknowledgments

This work is supported by NSF grant IIS-0121297.

## References

- Bahl, P., and Padmanabhan, V. 2000. Radar: An in-building rf-based user location and tracking system. In *Proc. IEEE Infocom*, 775–784. IEEE CS Press.
- Cook, D., and Das, S. 2007. How smart are our environments? an updated look at the state of the art. *Journal of Pervasive and Mobile Computing*.
- Feng, W.; Walpole, J.; Feng, W.; and Pu, C. 2001. Moving towards massively scalable video-based sensor networks. In *Proc. Workshop on New Visions for Large-Scale Networks: Research and Applications*.
- Hightower, J., and Borriello, G. 2001. Location systems for ubiquitous computing. *Computer* 32(8):57–66.
- Intille, S. S. 2002. Designing a home of the future. *IEEE Pervasive Computing* vol. 1:80–86.
- Jakkula, V., and Cook, D. 2007. Using temporal relations in smart environment data for activity prediction. In *Proceedings of the 24th International Conference on Machine Learning*.
- Jakkula, V.; Crandall, A.; and Cook, D. 2007. Knowledge discovery in entity based smart environment resident data using temporal relation based data mining. *IEEE Workshop on Temporal-Spatial Data Mining*.
- Krumm, J.; Harris, S.; Meyers, B.; Brumitt, B.; Hale, M.; and Shafer, S. 2000. Multi-camera multi-person tracking for easy living. In *Proc. 3rd IEEE Int'l Workshop Visual Surveillance*, 3–10. Piscataway, N.J.: IEEE Press.
- Lu, C.; Ho, Y.; and Fu, L. 2007. Creating robust activity maps using wireless sensor network in a smart home. In *Proceedings of the 3rd Annual Conference on Automation Science and Engineering*. Scottsdale, AZ: IEEE.



- Moncrieff, S. 2007. Multi-model emotive computing in a smart house environment. *Journal of Pervasive and Mobile Computing* vol. 3.
- Orr, R., and Abowd, G. 2000. The smart floor: A mechanism for natural user identification and tracking. In *Proc. 2000 Conf. Human Factors in Computing Systems (CHI 2000)*. New York: ACM Press.
- Priyantha, N.; Chakraborty, A.; and Balakrishnan, H. 2000. The cricket location support system. *Proc. 6th Ann. Int'l Conference on Mobile Computing and Networking* 32–43.
- Roy, N.; Roy, A.; Basu, K.; and Das, S. K. 2005. A co-operative learning framework for mobility-aware resource management in multi-inhabitant smart homes. In *Proc. of IEEE Int'l Conference on Mobile and Ubiquitous Systems: Networking and Services*, 393–403.
- Snidaro, L.; Micheloni, C.; and Chivedale, C. 2005. Video security for ambient intelligence. *IEEE Transaction on Systems, Man and Cybernetics, part A* Vol. 35:133–144.
- Tapia, E. M.; Intille, S. S.; and Larson, K. 2004. *Activity recognition in the home setting using simple and ubiquitous sensors*, volume LNCS 3001 of *Proc. PERVASIVE*. Berlin / Heidelberg: Springer-Verlag. 158–175.
- Wren, C., and Tapia, E. 2006. Toward scalable activity recognition for sensor networks. *Hazas, M., Krumm, J., Strang, T. (eds.) LoCA 2006* 3987:168–185.
- Yin, J.; Yang, Q.; and Shen, D. 2007. Activity recognition via user-trace segmentation.