Health Monitoring in an Agent-Based Smart Home by Activity Prediction

Sajal K. Das and Diane J. Cook
Center for Research in Wireless Mobility and Networking (CReWMaN)
Department of Computer Science Engineering
The University of Texas at Arlington
{das, cook}@cse.uta.edu

Abstract

To many people, home is a sanctuary. For those people who need special medical care, they may need to be pulled out of their home to meet their medical needs. As the population ages, the percentage of people in this group is increasing and the effects are expensive as well as unsatisfying. We hypothesize that many people with disabilities can lead independent lives in their own homes with the aid of at-home automated assistance and health monitoring. In order to accomplish this, robust methods must be developed to collect relevant data and process it dynamically and adaptively to detect and/or predict threatening long-term trends or immediate crises. The main objective of this paper is to investigate techniques for using agent-based smart home technologies to provide this at-home health monitoring and assistance. To this end, we have developed novel prediction algorithms that will determine common activities of an inhabitant and report significant activity anomalies that may indicate health crises. Specifically, we address the following technological challenges: 1) developing secure situation- or context-aware methods to collect at-home health and activity data, 2) learning patterns from the collected data, 3) identifying long-term trends of increasing or deteriorating health, 4) detecting and responding to anomalies in the data, and responding to predicted health problems, and 5) providing reminder and automation assistance to promote independent living at home. The proposed solution approaches are being tested in simulation and with volunteers at the UTA's MavHome site, an agent-based smart home project funded by NSF.

1 Introduction and Motivation

We live in an increasingly connected and automated society. We are investigating monitoring and automation assistance in our most personal environment: *the home*. In particular, we will investigate methods of adapting such a smart home environment to perform health monitoring and assistance for persons with disabilities and for aging adults. This integration of engineering and life science builds upon UTA's MavHome project [9], a home environment that perceives the state of the home through sensors and intelligently acts upon the environment through controllers.

As Lanspery and Hyde [17] state, "For most of us, the word 'home' evokes powerful emotions [and is] a refuge". They note that older adults and people with disabilities want to remain in their homes even when their conditions worsen and the home cannot sustain their safety. In a national survey, researchers found that 71% of the respondents felt strongly that they wanted to remain in their current residence as long as possible, and another 12% were somewhat likely to remain there [1]. Nearly 1/4 of the respondents expected that they or a member of their household would have problems getting around their house in the next five years. Of these respondents, 86% stated that they had made at least one modification to their home to make it easier to live there, and nearly 70% believe that the modifications will allow them to live in the current homes longer than would have otherwise been possible. A separate study supported these results and found that the most common modifications were an easy-to-use climate control system and a personal alert system.

Zola [26] maintains that the problems of aging and disability are converging. Improvements in medical care are resulting in increased survival into old age, thus problems of mobility, vision, hearing, and cognitive impairments will increase [19, 20]. As the baby boomers enter old age, this trend will be magnified. By 2040, 23% will fall into the 65+ category [17]. An AARP report [1, 2] strongly encourages increased funding for home modifications that can keep older adults with disabilities independent in their own homes.

While use of technology can be expensive, it may be more cost effective than the alternative [13]. Nursing home care is generally paid either out-of-pocket or by Medicaid. Typical nursing home costs are about $40,000 a year, and the $197 billion of free care offered by family members comes at the sacrifice of independence and job opportunities by the family caregivers.

In this paper, our goal is to assist the elderly and individuals with disabilities by providing home capabilities that will monitor health trends and assist in the inhabitant's day to day activities in their own homes. The result will save money for the individuals, their families, and the state. We are seeking to meet this goal using the MavHome smart home environment. MavHome is equipped with sensors that record inhabitant interactions with many different devices, medicine-taking schedules, movement patterns, and vital signs. We have developed novel algorithms that learn patterns of activities from this data. We will apply these capabilities to health monitoring in the following steps: (i) Perform secure, situation (or context)-aware collection of inhabitant health and activity data; (ii) Use our data mining and prediction techniques to learn patterns in collected data; (iii) Identify trends that could indicate health concerns or a need for transition to assisted care; (iii) Detect anomalies in regular (monitored) patterns that may require intervention; and (v) Provide reminder and automation assistance for inhabitants.

We will report experimental results on simulated data and on data collected from student volunteers in smart apartments (MavPad) on UTA campus. By investigating these issues we can offer the community an intelligent system with learning algorithms that not only perform their individual tasks well, but also form a synergistic whole that is stronger than the parts. We will show the benefits that can be derived from these algorithms in intelligent environments and other learning-driven situations for aging adults and persons with disabilities.

The rest of the paper is organized as follows. Section 2 gives an overview of the MavHome project while Section 3 explains how to identify significant episodes in smart homes. Section 4 deals with prediction of inhabitants' mobility and activities. Section 5 summarizes the capabilities of MavHome to assist the elderly and disabled.

2 Overview of the MavHome Smart Home

We define an intelligent environment as one that is able to acquire and apply knowledge about its inhabitants and their surroundings in order to adapt to the inhabitants and meet the goals of comfort and efficiency [7]. These capabilities rely upon effective prediction, decision making, robotics, wireless and sensor networking, mobile computing, databases, and multimedia technologies. With these capabilities, the home can adaptively control many aspects of the environment such as climate, water, lighting, maintenance, and multimedia entertainment. Intelligent automation of these activities can reduce the amount of interaction required by inhabitants, reduce energy consumption and other potential wastages, and provide a mechanism for ensuring the health and safety of the environment occupants [6].

Compared to existing smart home research projects and/or prototypes, MavHome is unique in combining technologies from artificial intelligence, machine learning, databases, robotics, and multimedia computing to create a smart home that acts as an intelligent agent, as shown in Figure 1. MavHome's smart home capabilities are organized into an agent-based software architecture that seamlessly connects needed components while allowing improvements to be made to any of the supporting technologies. that are separated into four cooperating layers. The *Decision layer* selects actions for the agent to execute. The *Information layer* collects information and generates inferences useful for decision making. The *Communication layer* routes information and requests between agents. The *Physical layer* contains the environment hardware including devices, transducers, and network equipment. The MavHome software components are connected using a CORBA interface.

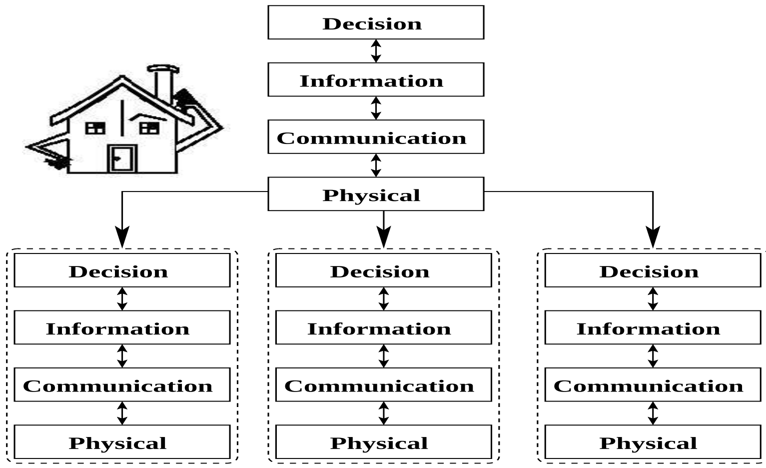

Figure 1: MavHome agent architecture.

Because controlling an entire house is a very large and complex learning and reasoning problem, the problem is decomposed into reconfigurable subareas or tasks. Thus the Physical layer for one agent may in actuality represent another agent somewhere in the hierarchy, which is capable of executing the task selected by the requesting agent.

Perception is a bottom-up process. Sensors monitor the environment (e.g., lawn moisture level) and, if necessary, transmit the information to another agent through the Communication layer. The database records the information in the Information layer, updates its learned concepts and predictions accordingly, and alerts the Decision layer of the presence of new data. During action execution, information flows top down. The Decision layer selects an action (e.g., run the sprinklers) and relates the decision to the Information layer. After updating the database, the Communication layer routes the action to the appropriate effector to execute. If the effector is actually another agent, the agent receives the command through its effector as perceived information and must decide upon the best method of executing the desired action. Specialized interface agents allow interaction with users, robots, and external resources such as the Internet. Agents can communicate with each other using the hierarchical flow shown in Figure 1.

3 Learning to Identify Significant Episodes

In order to maximize comfort, minimize cost, and adapt to inhabitants, a smart home must rely upon tools from artificial intelligence such as data mining and prediction. Prediction can be used to determine the inhabitant's next action. Specifically, MavHome needs to identify repetitive tasks performed by inhabitants that establish a baseline for learning trends in behaviors, detecting anomalies, and determining repetitive tasks worthy of automation by the home. The home can make this prediction based solely on previously-seen inhabitant activities and the current state of the inhabitant and the house.

A smart home inhabitant performs various routine activities, which may be considered as a sequence of events, with some inherent pattern of recurrence. This repeatability leads us to the conclusion that the sequence can be modeled as a stationary stochastic process. We can then perform inhabitant action prediction by first mining the data (using episode discovery techniques) to identify sequences of actions that are regular and repeatable enough to generate predictions, and by second using a sequence matching approach (LeZi-update or Active LeZi) to predict the next action in one of these sequences.

Our Episode Discovery (ED) data mining algorithm is based on the work of Srikant and Agrawal [23] for mining sequential patterns from time-ordered transactions. We move an examination window through the history of events or inhabitant actions, looking for episodes (sequences) within the window that merit attention, or *significant episodes*. Each candidate

Table 1: Scenario prediction results.

Scenario	1	2	3	4	5	Average
Events	12958	12884	12848	13058	12668	12883
Significant Episodes	13	13	13	13	13	13
IPAM Percentage Correct	39%	42%	43%	40%	41%	41%
IPAM+ED Percentage Correct	77%	84%	69%	73%	65%	74%
BPNN Percentage Correct	62%	64%	66%	62%	64%	64%
BPNN+ED Percentage Correct	84%	88%	84%	84%	88%	86%
ED Processing Time (secs)	11	9	10	9	9	10

episode is evaluated using the Minimum Description Length (MDL) principle [21]. The MDL principle favors patterns that can be used to minimize the description length of a database by replacing each instance of the pattern with a pointer to the pattern definition. A detected regularity factor (daily, weekly, or other time frame) further compresses the data because the sequence can be removed without storing a pointer to the sequence definition, and thus increases the value of a pattern. Deviations from the pattern definition in terms of missing events, extra events, or changes in the regularity of the occurrence add to the description length because extra bits must be used to encode the change, thus lowering the value of the pattern. The larger the potential amount of description length compression a pattern provides, the greater the impact that results from automating the pattern.

Our ED algorithm successfully identified daily and weekly patterns in synthetic data based on the MavHome scenario described earlier. We also used ED to mine data that was collected in the MavHome environment from six students during 2003. These students indicated in advance their likely daily and weekly activities, although the timing, regularity, and order of events would vary throughout the semester. The dataset contains 618 interactions that are members of patterns occurring once a week, multiple times a week, and randomly. ED successfully identified the patterns of three of the inhabitants as weekly significant episodes, and marked which of the 618 interactions contributed to the significant episodes [14].

The knowledge that ED obtains by mining the user action history can be used in a variety of ways. First, the mined patterns provide information regarding the nature of activities in the home, which can be used to better understand lifestyle patterns and aid in designing homes and devices for the home. Second, the significance of a current event as a member of a discovered pattern can be used in controlling the home, to determine whether this task is worth attempting to automate. Third, knowledge of the mined sequences can improve the accuracy of predicting the next action, by only performing prediction for events known to be part of a common pattern. We demonstrate the ability of ED to perform the third task, improving the accuracy of prediction algorithms, by adding the mined results as a preprocessor to two simple prediction algorithms. Action sequences are first filtered by the mined sequences. If a sequence is considered significant by the mining algorithm, then predictions can be made for events within the sequence window.

To test the filtering capabilities of ED, we coupled it with the IPAM sequential predictor [24] and a back-propagation neural network (BPNN). We created a sequence of 13,000 actions based on five randomly-generated scenarios, a situation in which these algorithms by themselves may not perform well. ED discovered 14 episodes in the data sets, and appreciably improved the accuracy of both algorithms across all five scenarios, as can be seen in Table 1. By using ED, we improve the accuracy of the prediction algorithms by reducing the total number of incorrect predictions, leading to inaccuracies in learned health trends, detected anomalies, and automated patterns.

4 Learning to Predict Inhabitant's Mobility and Actions

Prediction is an important component in a variety of domains in artificial intelligence and machine learning, that allows intelligent systems to make more informed and reliable decisions. Certain domains require that prediction be performed on sequences of events that can typically be modeled as stochastic processes. Especially common is the problem of sequential prediction: given a sequence of events, how do we predict the next event based on a limited known history? This is true, for example, when predicting inhabitant actions in a smart environment such as MavHome. Prediction can be performed of upcoming inhabitant movement activities based on observed past movement/activity patterns.

The fundamental framework due to Bhattacharya and Das [3, 4] demonstrate how optimal, on-line learning and prediction algorithms can be designed for user mobility in a wide area cellular network. This adaptive framework is based on the fundamental hypothesis that *a good leaner is a good predictor*, thus applying information theoretic concepts and Lempel-Ziv (LZ78) family of text compression techniques [25], on the mobility profiles collected over time. We have also shown how to adapt this framework for managing indoor mobility and context-aware resource management in smart home environments [22].

We have recently extended this framework and designed an algorithm, named *Active LeZi (ALZ)*, for sequential action prediction from an information theoretic standpoint. For any sequence of events that can be modeled as a stochastic process, this algorithm employs the power of Markov models to optimally predict the next symbol in any stochastic sequence. Inhabitant action prediction, similar to other prediction problems, requires that the prediction algorithm be capable of incrementally gathering information and deliver real time, "online" predictions. Active LeZi is based on the LZ78 data compression algorithm, which employs incremental parsing, thereby addressing this requirement. The smart home provides a ready environment for employing sequential prediction algorithms such as ALZ. The home will have to make inhabitant predictions based only on previously-seen inhabitant actions.

Consider a sequence of events being generated by an arbitrary deterministic source, which can be represented by the stochastic process $X = \{x_i\}$. The sequential prediction problem can then be stated as follows. Given the sequence of symbols $\{x_1,\ x2,\ \ldots\ x_i\}$, what is the next symbol x_{i+1}? Well-investigated text compression methods, such as LZ78, have also been established as good predictors [4, 11]. According to information theory, a predictor that builds a model whose entropy approaches that of the source achieves greater predictive accuracy. Also, it has been shown that a predictor with an order that grows at a rate approximating the entropy rate of the source is an optimal predictor [11]. Another motivation to look to the field of text compression is that such algorithms are essentially incremental parsing algorithms, providing a basis for online processing. ALZ addresses this prediction problem based on the above-mentioned motivations.

The LZ78 data compression algorithm due to Lempel and Ziv [25], is a dictionary-based text compression algorithm that performs incremental parsing of an input sequence and, as a result, gradually changes the Markov order at the appropriate rate. This algorithm parses an input string $x_1,\ x_2,\ \ldots,\ x_i$ into c(i) substrings $w_1,\ w_2,\ \ldots,\ w_{c(i)}$ such that for all j¿0, the prefix of the substring w_j (i.e., all but the last character of w_j) is equal to some w_i for $1 < i < j$. Because of this prefix property, parsed substrings can efficiently be maintained in a trie.

Consider the sequence of input symbols $x^n = aaababbbbbaabccddcbaaaa$. An LZ78 parsing of this string would yield the following set of phrases: $a, aa, b, ab, bb, bba, abc, c, d, dc, ba, aaa$. This algorithm maintains statistics for all contexts seen within the phrases. For example, the context a occurs 5 times (at the beginning of the phrases a, aa, ab, abc, aaa), the context bb is seen 2 times (bb, bba), etc. These context statistics are stored in a trie.

LZ78 suffers from two drawbacks: in any LZ parsing of an input string, all the information crossing phrase boundaries is lost. In most situations, there might be significant patterns crossing phrase boundaries that affect the next symbol in the sequence. Second, although

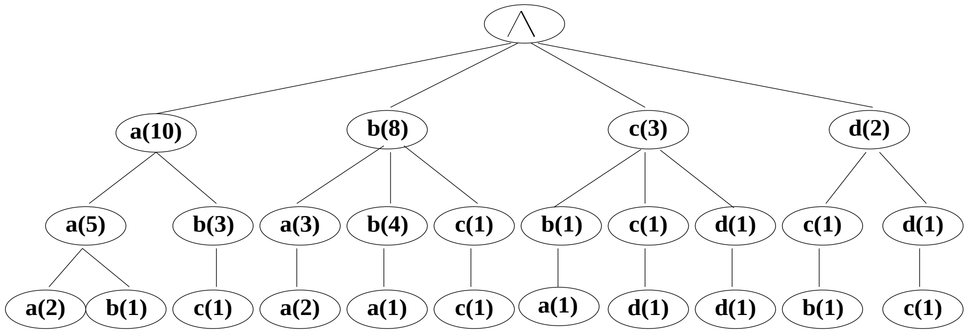

Figure 2: Trie formed by the ALZ parsing of the string *aaababbbbbaabccddcbaaaa*.

LZ78 asymptotically approaches optimal predictability, its convergence rate is slow [11].

Our Active LeZi (ALZ) algorithm is an enhancement of LZ78 that addresses these drawbacks. As the number of states in an input sequence grows, we can see that the amount of information being lost across the phrase boundaries increases rapidly. Our solution to this problem involves maintaining a variable length window of previously-seen symbols. We choose the length of the window at each stage to be equal to the length of the longest phrase seen in a classical LZ78 parsing of the symbol history. The reason for selecting this window size is that the LZ78 algorithm is essentially constructing an (approximation to an) order-(k-1) Markov model, where k is equal to the length of the longest LZ78 phrase seen so far, which builds a better approximation to the order-k Markov model. Therefore, we gain a better convergence rate to optimal predictability as well as greater predictive accuracy. Figure 2 shows the trie formed by the ALZ parsing of the input sequence aaababbbbbaabccddcbaaaa, which results in an order-2 Markov model.

In order to predict the next event of the sequence for which ALZ has built a model, we calculate the probability of each next action (symbol) occurring in the sequence, and predict the one with the highest probability. For sequential prediction, to achieve good convergence rates to optimal predictability, the predictor must "lock on" to the minimum possible set of states that is representative of the sequence being considered by using a "mixture" of all possible order models in assigning the next symbol its probability estimate. For this goal, we employ the Prediction by Partial Match (PPM) family of predictors.

PPM algorithms combine weighted different-order Markov models to build a probability distribution. In our predictive scenario, ALZ builds an order-k Markov model. The window maintained by ALZ represents the set of *contexts* used to compute the probability of the next symbol. In our example sequence aaababbbbbaabccddcbaaaa, the last phrase aaa (which is also the current ALZ window) is used. Within this phrase, the contexts that can be used are all suffixes within the phrase, except the window itself (i.e., aa, a, and the null context).

Suppose the probability that the next symbol is an a is being computed. From Figure 2 we see that an a occurs two out of the five times that the context aa appears, the other cases producing two null outcomes and one b. Therefore, the probability of encountering an a at the context aa is 2/5, and we now fall back (escape) to the order-1 context (i.e., the next lower-order model) with probability 2/5. At the order-1 context, we see an a five out of the ten times that we see the a context, and of the remaining cases, we see two null outcomes. Therefore we predict the a at the order-1 context with probability 5/10, and escape to the order-0 model with probability 2/10. At the order 0 model, we process an a for 10 out of 23 symbols seen so far, and we therefore predict a with probability 10/23 at the null context. The blended probability of seeing an a as the next symbol is therefore $\frac{2}{5} + \frac{2}{5}\{\frac{5}{10} + \frac{2}{10}(\frac{10}{23})\}$. This blending strategy assigns greater weight to higher-order models, while lower-order models are suppressed owing to the null context escape probability. This is in keeping with the advisability of making the *most informed* decision.

Figure 3 shows the performance of ALZ tested on data obtained from a Synthetic Data

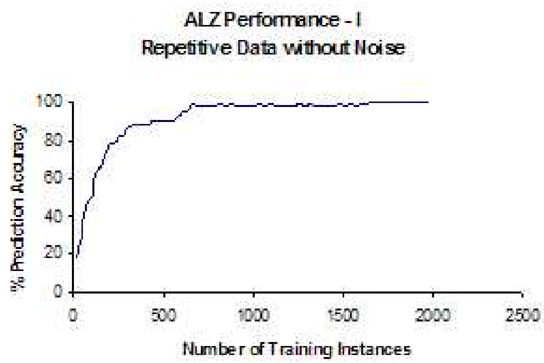

(a)

(b)

Figure 3: ALZ learning curve on data a) without noise and b) with noise.

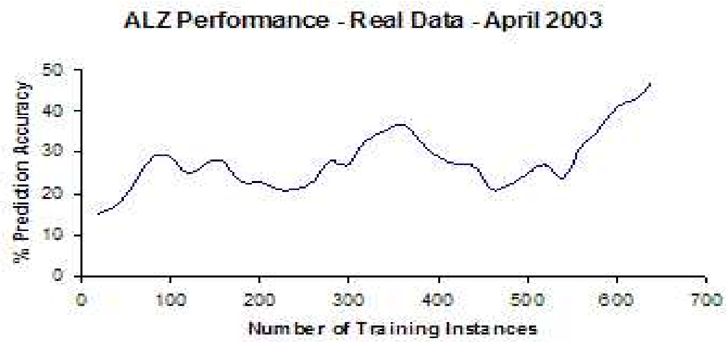

Figure 4: ALZ performance on real data from April, 2003.

Generator containing in the first case a high degree of repetitiveness and little noise, and containing in the second case a lower percentage of repeated events and more noise. The ALZ learning curves were plotted by testing the number of correct predictions from the next 100 events, while increasing the training data set. ALZ's performance converges quickly to 100% in the first experiment (Figure 3a), and eventually converges to 86% accuracy in the second experiment (Figure 3b).

In addition to the tests performed on the simulated data, ALZ was also tested on data collected in an actual smart home environment - the UTA MavHome Lab. This data was collected from monitoring the interactions of six different participants with devices, based on regular daily and weekly patterns of activity. These represent data that will typically be seen in a home environment – with all the associated randomness and noise within the system. As we can see from Figure 4, the algorithm converges to about 47% accuracy on one month of data consisting of 750 data points.

We have also enhanced ALZ to learn a Gaussian distribution representing relative time intervals between actions in the observed sequences, each characterized by a mean μ and standard deviation σ. Here, the relative frequency counts of the various phrases are stored as before in the trie and used to incrementally build Gaussians that represent the observed normal distribution of the relative time of occurrence of the last event in the corresponding phrase. As with the prediction of the event itself, information stored at various orders of the trie are blended to strengthen the resulting prediction. This enhancement was tested using synthetic data consisting of a smart home scenario with repeating events, where the order of events was constant, but the time difference between successive events were drawn from different distributions depending on the context in which that event was taking place. Results from this test show that in 92% of the cases, the next event occurred within the mean +/- 2 standard deviations of the predicted time.

As an additional experiment, we use ED to filter data for prediction by ALZ using actual collected MavHome activity data. Based on one month of collected data with an extremely large amount of noise (there are 20 students and 50 devices in the environment), ALZ achieves a predictive accuracy of 30%, while ALZ with ED achieves an improved predictive accuracy of 44%. These results show that even a powerful predictor such as ALZ benefits from being combined with the results of a data mining algorithm.

5 Using a Smart Home to Assist Elderly and Disabled

The data mining, prediction, and multiagent technologies available in MavHome can be employed to provide health care assistance in living environments. Specifically, models can be constructed of inhabitant activities and used to learn activity trends, detect anomalies, intelligently predict possible problems and make health care decisions, communicate with caregivers, and provide automation assistance for inhabitants with special needs.

A variety of approaches have been investigated in recent years to automate caregiver services. Many of the efforts offer supporting technologies in specialized areas, such as using computer vision techniques to track inhabitants through the environment and specialized sensors to detect falls or other crises. Some special-purpose prediction algorithms have been implemented using factors such as measurement of stand-sit and sit-stand transitions and medical history, but are limited in terms of what they predict and how they use the results. Remote monitoring systems have been designed with the common motivation that learning and predicting inhabitant activities is key for health monitoring, but very little work has combined the remote monitoring capabilities with prediction for the purpose of health monitoring. Some work has also progressed toward using typical behavior patterns to provide reminders, particularly useful for the elderly and patients suffering from various types of dementia [12, 15, 18].

Our approach differs from these earlier explorations in that we are combining capabilities in the areas of data collection, remote monitoring, prediction, data mining, and knowledge engineering to provide predictive health monitoring assistance for inhabitants with disabilities and caregivers. Instead of designing one specialized component, we demonstrate here that a smart environment can accomplish all of the tasks needed to identify patterns indicating or predicting a sudden or slow change in health status, to supply caregivers with periodic or emergency information, and to provide inhabitants with needed automation assistance.

5.1 Capability 1: Perform Secure, Context-Aware Data Collection

The first capability we are developing is that of remote collection of long-term activity and health status data. Data will consist of monitored inhabitant activities, vital signs, and interactions with the environment, and will be collected using context-aware technology.

Our smart home environment is currently equipped to gather the following types of information: 1) inhabitant usage of any electrical device in the home, 2) usage of water in the home and amount, 3) temperature settings, 4) inhabitant weight, 5) inhabitant movement throughout the home, 6) prescribed and actual medicine dispensing schedule, 7) time, duration, and intensity of exercise, and 8) use of food items in kitchen. Active sensors including wearable vital sign monitors can be integrated to further refine the model.

We have designed efficient algorithms for collecting real-time data that are context- and/or situation-aware. For example, the inhabitant's current activity (e.g., cooking vs. watching television) or location in the environment (e.g., bedroom vs. navigating the stairs) can affect the choice of sensors to use, and thus represent a defined context. Our data mining and prediction algorithms are highly scalable, a desirable feature when numerous tiny, portable sensor devices are involved. Furthermore, the algorithms create personal profiles and hence provide customized solutions to individual patients. There may also be a need for ad hoc collaboration of various entities as the emergency need arises. These problems can be elegantly tackled with the help of pervasive computing and communications technology [10, 16].

Use of sensors in smart homes and on elderly or people with disabilities is critical for collecting, storing, and processing appropriate data intelligently and in a timely manner. Data collection and processing needs to be backed by an infrastructure that allows anytime, anywhere reliable access to data sources (e.g., national data bases or vital records). This is crucial for the purpose of data mining, intelligent decision making and profiling. Privacy of the information is highly required since personal information will be collected during the monitoring phase.

5.2 Capability 2: Identify Trends in Long-Term Data

The second health monitoring capability uses our data mining and prediction techniques to identify trends in long-term data. Trends in time-varying data can be discovered and predicted using the same data mining and prediction techniques described earlier. Instead of capturing individual events and readings to store in the trie and use for prediction, changes in these values over varying periods of time are captured, stored, and predicted. In particular, for each feature of interest, the qualitative change in value (increase, decrease, same) and quantitative relative change (numeric difference) is recorded. Because the chosen time steps are limitless in value and yet critical to learning trends, we allow the user to determine the time step of choice (e.g., daily, weekly, monthly) and use auto-correlation techniques to automatically determine the time step most indicative of a temporal trend in the data.

The user can select features to monitor for trend analysis, including 1) change in mobility (schedule, time, room location, total movement), 2) change in amount of exercise, 3) change in deviation from prescribed medicine schedule, 4) change in amount of smart home requested assistance (reminders or automation assistance), 5) change in nutrition (types of food, number of calories), and 6) change in amount and types of activities.

Using the ED algorithm, long-term trends can be discovered from the raw data. The algorithm can also be used to search for changes of a type of duration indicated by the user. In addition, the prediction algorithm can be used to predict the upcoming changes in these features for the next time step. Finally, we can use data from individuals assessed by practitioners to learn classes of patient types, such as patients who require a move to an assisted care facility. Time-varying data captured for a particular individual can then be classified based on the learned models.

The result of this technology can be used in a number of ways. For patients recovering from an illness or accident, the algorithms can be used to determine whether they are regaining strength and vitality at an acceptable rate. For patients at risk of various types of accidents and health risk situations, the trends can be used to determine the current health risk and predicted short term health risk. These analyses can aid a caregiver in deciding whether the individual requires a change in medication, activity schedule, care scheduling, or environment.

5.3 Capability 3: Detect Anomalies in Current Data

A direct outcome of our work in data mining prediction is the ability to detect anomalies in collected data. For this purpose, we use the intuitive notion of an anomaly as *a surprising or unusual occurrence*. With this notion in mind, the anomaly value of each monitored activity or captured feature value can be collected and used to provide health status information.

An important consideration in any anomaly-detection system is the regularity, or predictability, of the data. Generally, the more predictable the data, the easier it is to detect anomalies. As a result, the anomaly value of each data point is influenced first by the degree of membership of the current state in a significant episode, second by the significance value of the episode itself, and third by the deviation of the observed activity from the predicted activity given the known episode.

The degree of membership of the current state in any of the discovered significant episodes is calculated as the probability of occurrence of the episode given the immediate history of observed activities. A probability is returned for each episode. This probability is then multiplied with the significance of the episode itself to calculate the regularity of the data context. The episode yielding the highest regularity value, episode i, represents the baseline for the anomaly calculation. Once membership in episode i is established, the ALZ trie is used to provide the probability of the next predicted event e. If the event that actually occurs has a very low probability given the context, the resulting event constitutes an anomaly. If the regularity of the context is high and the probability of the observed event is low, this event should be flagged as an anomaly. Thus the anomaly value of observed event e can be calculated as $A = \frac{(Member(e, Episode_i) * Significance(Episode_i))}{P(e|Episode_i, ImmHistory(e))}$.

Anomalies can be be handled in a manner appropriate to the nature of the event and anomaly value. The criticality of events can be set by the user. For example, an anomaly in predicted interactions with the house will typically have a low criticality and thus simply be recorded as an anomaly or prompt a reminder of the inhabitant of their normal routine. For a more critical anomaly, such as a sudden disruption in vital signs, the home will first attempt to contact the inhabitant. If contact cannot be made, the home will contact the caregiver. Web cameras placed throughout the house can be used by the caregiver to check on the inhabitant in such a situation and intervene with a medical difficulty.

5.4 Capability 4: Design Reminder Assistance System

An additional benefit that our smart home data mining and prediction capabilities can offer is to provide a reminder system for inhabitants. Reminders can be triggered by two situations. First, if the inhabitant queries the home for his/her next routine activity, the activity with the highest probability will be given based on the ALZ prediction. Second, for episodes with a high degree of significance and criticality, if the inhabitant deviates from the normal routine (creates an anomaly), the house will initiate contact with the inhabitant to remind him/her of the next expected activity. Such a reminder service will be particularly beneficial for individuals suffering from dementia. The reminder can enable the individual to feel secure about performing daily activities, and prevent accidents such as forgetting to turn off the water in the bathtub or keeping the doors unlocked when leaving the house.

5.5 Capability 5: Provide Automation Assistance

Intelligent automation of home activities can reduce the amount of manual home interaction required by inhabitants. Earlier we developed a decision-making algorithm to automate activities in the house [8]. Selection of actions for the house to take is based on suggestions from a reinforcement learning algorithm, which minimizes necessary manual intervention with the house and reduces actions with costly effects. The decision maker relies on input from ALZ to suggest the typical next inhabitant action, and uses results from the ED data mining algorithm to segment the learning space and improve the quality of the learned results.

Users can specify specific activities they would like the house to automate (temperature control, control of hard-to-reach devices, etc.). MavHome will then automate the activities based on learned preferences and action patterns. In addition, if the inhabitant does not respond to a reminder of a critical event, MavHome will automate an activity to ensure safety of the individual. Such actions include shutting off the bathwater, turning off the stove, or locking up the house when it is empty.

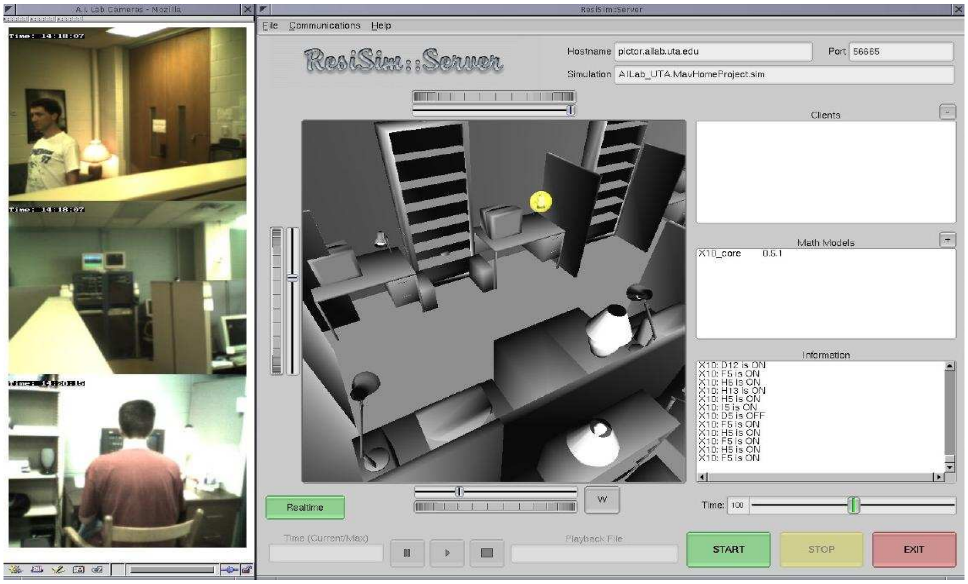

Figure 5: MavHome environment and interface.

6 Conclusion

Initial algorithms are in place for data mining, mobility and action prediction, decision making, and automation. These capabilities are being used in the MavHome lab (MavLab) and apartment (MavPad). Figure 5 shows lights in the entryway (top left) and on Ryan's desk (bottom left) turning on in reaction to inhabitant activities (these automated actions are also reflected in the simulator). We are collecting health-specific data and testing using volunteers living in MavPad and MavHome, as well as recruited residents of the C.C. Young Retirement Community in Dallas, Texas.

Acknowledgments: This work is supported by US National Science Foundation under ITR grant IIS–0121297.

References

[1] AARP, "Fixing to Stay: A National Survey of Housing and Home Modification Issues," 2000.

[2] AARP, "These Four Walls ... Americans 45+ Talk About Home and Community," 2003.

[3] A. Bhattacharya and S. K. Das, "LeZi-update: An information-theoretic approach to track mobile users in PCS networks," *Proc. Ann. Int'l Conference on Mobile Computing and Networking*, August 1999, pp. 1-12.

[4] A. Bhattacharya and S. K. Das, "LeZi-update: An information-theoretic approach for personal mobility tracking in PCS networks," *ACM-Kluwer Wireless Networks* (WINET), vol. 8, no.2, pp.121-137, 2002.

[5] M. Chan, C. Hariton, P. Ringeard, and E. Campo, "Smart House Automation System for the Elderly and the Disabled," *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, pp. 1586-1589, 1995.

[6] D. J. Cook and S. K. Das, "Health Monitoring in an Agent-Based Smart Home," To be presented at the *International Conference on Aging, Disability and Independence* (ICADI), Washington, DC, Dec 2003.

[7] D. J. Cook and S. K. Das (Eds.), *Smart Environments: Architectures, Protocols and Applications*, John Wiley, to appear, Sept 2004.

[8] D. J. Cook, M. Huber, K. Gopalratnam, and E. Heierman, "Learning to Control a Smart Home Environment," *IEEE Transactions on Systems, Man and Cybernatics*, to appear.

[9] S. K. Das, D. J. Cook, et al., "The Role of Prediction Algorithms on the MavHome Smart Home Architectures," *IEEE Wireless Communications* (Special Issue on Smart Homes), Vol. 9, No. 6, pp. 77-84, Dec 2002.

[10] S. K. Das, M. Kumar, and B. Shirazi, "Pervasive Community Computing (PCC) Paradigm with Application to Health care," *Proc. International Conference on Ubiquitous Computing* (Keynote Talk), Seoul Korea, Oct 2003.

[11] M. Feder, N. Merhev, and M. Gutman, "Reflections on 'Universal Prediction of Individual Sequences'," *IEEE Information Theory Society Newsletter*, 1994.

[12] T. T. Fulmer, et al. "An Intervention Study to Enhance Medical Compliance in Community-Dwelling Elderly Individuals," *Journal of Gerontological Nursing*, 25(8): 6-14, 1999.

[13] P. J. Grayons, "Technology and Home Adaptation," In S. Lanspery and J. Hyde (Eds.), *Staying Put: Adapting the Places Instead of the People*, pp. 55-74, Baywood Pub. Co., 1997.

[14] E. Heierman and D. J. Cook, "Improving Home Automation by Discovering Regularly Occurring Device Usage Patterns," *Proceedings of the International Conference on Data Mining*, 2003.

[15] H. Kautz, L. Arnstein, G. Borriello, O. Etzioni, and D. Fox, "An Overview of the Assisted Cognition Project," *Proceedings of the AAAI Workshop on Automation as Caregiver*, 2002.

[16] M. Kumar, S. K. Das, et al, "PICO: A Middleware Framework for Pervasive Computing," *IEEE Pervasive Computing*, Vol. 2, No. 3, pp. 72-79, July-Sept 2003.

[17] S. Lanspery and J. Hyde (Eds.), *Staying Put: Adapting the Places Instead of the People*, Introduction, pp. 1-22, Baywood Pub. Co., 1997.

[18] M. E. Pollack, et al., "Autoreminder: A Planning, Monitoring, and Reminding Assistive Agent," *Proceedings of the International Conference on Intelligent Autonomous Systems*, 2002.

[19] J. Pynoos, "Neglected Areas in Gerontology: Housing Adaptation," In *Panel Presentation at the Annual Scientific Meeting of the Gerontological Society of America*, 2002.

[20] J. Pynoos and V. Regnier, "Design Directives in Home Adaptation," In S. Lanspery and J. Hyde (Eds.), *Staying Put: Adapting the Places Instead of the People*, pp. 41-54, Baywood Pub. Co., 1997.

[21] J. Rissanen, *Stochastic Complexity in Statistical Inquiry* World Scientific Pub. Co., 1989.

[22] A. Roy, S. K. Bhaumik, A. Bhattacharya, K. Basu, D. J. Cook, and S. K. Das, "Location Aware Resource Management in Smart Homes," *Proceedings of First IEEE International Conference on Pervasive Computing and Communications* (PerCom'03), Ft Worth, Texas, pp. 481-488, Mar 2003.

[23] R. Srikant and R. Agrawal, "Mining Sequential Patterns: Generalizations and Performance Improvements," *Proceedings of the Fifth International Conference on Extending Database Technology*, 1996.

[24] M. C. Torrance, "Advances in Human-Computing Interaction: The Intelligent Room," *Proceedings of the Working Notes of the CHI'95 Research Symposium*, 1995.

[25] J. Ziv and A. Lempel, "Compression of Individual Sequences via Variable Rate Coding," *IEEE Transactions on Information Theory*, Vol. 24, pp. 530-536, 1978.

[26] I. K. Zola, "The Convergence of Aging and Disability," In S. Lanspery and J. Hyde (Eds.), *Staying Put: Adapting the Places Instead of the People*, pp. 25-40, Baywood Pub. Co., 1997.