Health-Status Monitoring Through Analysis of Behavioral Patterns

Tracy S. Barger, Donald E. Brown, *Fellow, IEEE*, and Majd Alwan, *Member, IEEE*

***Abstract*—With the rapid growth of the elderly population, there is a need to support the ability of elders to maintain an independent and healthy lifestyle in their homes rather than through more expensive and isolated care facilities. One approach to accomplish these objectives employs the concepts of ambient intelligence to remotely monitor an elder's activities and condition. The SmartHouse project uses a system of basic sensors to monitor a person's in-home activity; a prototype of the system is being tested within a subject's home. We examined whether the system could be used to detect behavioral patterns and report the results in this paper. Mixture models were used to develop a probabilistic model of behavioral patterns. The results of the mixture-model analysis were then evaluated by using a log of events kept by the occupant.**

***Index Terms*—Ambient intelligence, human behavior models, mixture models.**

I. Introduction

The population of Americans over 65 numbered 35 million in 2000, comprising 12.7% of the total population and is projected to more than double to over 70 million by 2030, while the 85+ population is expected to increase from 4.2 million in 2000 to 8.9 million in 2030 [1]. With this proliferation of the elderly population also comes an increased need for services for the elderly, including assisted-living facilities. However, many elderly desire to stay in their own private residences for as long as possible, and thus methods are needed to allow them to do so safely and at reasonable costs.

One possible method to help enable elders to live independently is to employ the concepts of ambient intelligence by installing remote monitoring technologies in elders' homes. The technologies could alert relatives, caregivers, or health-care personnel of any change in an elder's normal activity pattern. The monitoring technologies should maximize the privacy of elders while still providing information of any problems or deviations from normal habits. As a result, simple and inexpensive motion-detection sensors are likely preferred to more invasive technologies such as video recording. However, the ability of basic sensors to detect a person's behavioral patterns needs to be examined.

The primary objective of this paper is to examine whether a system of basic motion sensors could detect behavioral patterns and, thus, to provide the foundation for an ambient intelligence approach to elder care. A mixture-model framework was used to develop a probabilistic model of behavior and was tested on data from the University of Virginia's SmartHouse system. The results were then compared to a user log to provide validation of the patterns. The effect of analyzing behavior during work and off-days separately was also examined.

The remainder of this paper is organized as follows. Section II describes the background of the SmartHouse project and related literature. Section III describes the mixture-model analysis methodology. Section IV presents the results of the mixture-model analysis. Section V compares the mixture-model results to the user log. Section VI presents some conclusions.

II. Background

The Medical Automation Research Center (MARC) at the University of Virginia established a SmartHouse project to evaluate the ability to use a system of various sensors to provide monitoring health checks. A prototype of the system was installed in the residence of a volunteer subject to examine the use of such technology. This SmartHouse system consists of a series of motion detection and on–off switches. Eight motion-detection sensors have been installed, one in each room of the house (bedroom, bathroom, office, living room, kitchen, and laundry room/back door area), one at the front door, and one in the shower. These sensors fire whenever there is movement within the sensor's area of coverage. The switches are installed primarily in the kitchen and indicate actions such as the opening of a kitchen cabinet or the microwave. Sensor readings are collected continuously and consist of the sensor and the time at which the sensor readings were recorded.

There have been several other research projects that have investigated the use of various sensor technologies on monitoring of daily activity [4], [5], [13]–[15], [18]. Data-analysis techniques have included using plots and histograms to observe periodic patterns and infer activities [15], [18], comparing sensor readings to a daily activity log [13], [14], and using neural networks to predict how long a person will spend in or out of a single-room environment based on a succession of previous times spent in and out of a room [5]. Other research efforts have examined cyclical behavioral patterns, using probability estimates to detect deviations from the normal amount of time a person spends in a room in an hour [16].

Manuscript received February 5, 2004; revised June 11, 2004 and July 6, 2004. This work was supported in part by the Medical Automation Research Center, University of Virginia. This paper was recommended by Guest Editor G. L. Foresti.

T. S. Barger was with the Systems and Information Engineering Department, University of Virginia, Charlottesville, VA 22904 USA. She is now with Northrop Grumman Corporation, Chantilly, VA 20151 USA (e-mail: tsb6e@alumni.virginia.edu).

D. E. Brown is with the Systems and Information Engineering Department, University of Virginia, Charlottesville, VA 22904 USA (e-mail: brown@virginia.edu).

M. Alwan is with the Medical Automation Research Center, University of Virginia, Charlottesville, VA 22904 USA (e-mail: ma5x@virginia.edu).

Digital Object Identifier 10.1109/TSMCA.2004.838474

This paper expands upon these previous studies by applying a probabilistic approach to real data collected from a multiroom environment. It improves upon the studies performed on real data in a home-monitoring environment by utilizing a formal statistical methodology. Compared to the study using neural networks to predict the time spent in or out of a room, this work allows for a multiroom environment. The paper differs from the research on cyclical behavioral patterns that applied an algorithm to simulated data by using real data and not dividing the day into discrete hour time periods. This allows for the examination of behavioral patterns that may fluctuate between multiple discrete hours.

Examining the research in other monitoring systems outside of the home-health domain shows that there are a couple of examples where clustering and mixture-model methods have been applied to analyze sensor data. Mathews and Warwick examined the use of clustering techniques for identifying maintenance requirements on industrial machinery [12]. Goodman employed mixture models to classify target types (e.g., personnel, vehicle, background) based on data collected from unattended ground sensors [9]. These two papers demonstrate the use of clustering techniques in a sensor system domain and show the potential for their application in the field of home-health monitoring.

III. Mixture-Model Methods

A. Data

A total of 65 days of data collected from the SmartHouse motion sensors during a 12-week period were used in the mixture-model analysis presented here. The data were randomly split into a training and test set using an approximate 3/5 to 2/5 ratio to divide the days. Data from 40 of the days constituted the training set, while data from the remaining 25 days were used as the test set.

An observation in the data represents the time spent in one room before moving to the next room. We considered four attributes associated with each observation: 1) the sensor location; 2) the start time; 3) the length of time spent in the room; and 4) the activity level while in the room. The activity level is defined as the number of sensor firings while in the room, divided by the time (in seconds) spent in the room.

The observation data set was derived from the raw motion sensor data collected in a log file on a personal computer in the SmartHouse. The log file was chronologically ordered with each line displaying the sensor ID and the time when motion was detected. The sensor ID maps to one of the eight motion-detection sensors and was used to determine the room where motion occurred. Consecutive detections of motion in the same room were then combined into one observation. For example, consider the sample sequence of motion detections shown in Table I. The resulting observation used in the mixture-model analysis is shown in Table II. The observation represents an event occurring in the kitchen, starting at a time of 15 s, lasting 80 s, and having an activity level of 0.05.

TABLE I
Sample Sequence of Motion Sensor Data

Room	Time (s)
Living Room	3
Kitchen	15
Kitchen	27
Kitchen	59
Kitchen	92
Living Room	95

TABLE II
Sample Observation Derived From Sequence of Motion Data in Table I

Sensor Location	Start Time	Length (s)	Activity Level
Kitchen	15	80	0.05

B. Mixture Models

Mixture models were applied to the sensor data to develop a probabilistic model of event types which could represent underlying behavioral activities of the subject. The mixture-model approach is an unsupervised learning method of event estimation with other such methodologies including self-organizing maps [11] and k-means clustering [10].

The mixture-model approach served to cluster the observations with each cluster considered to be a different event type. An event type is thus defined to be a cluster of data observations (e.g., long lengths of time spent in the kitchen in the evening), while an activity is defined as an actual behavior performed by the subject (e.g., preparing dinner). Assuming that the sensor readings came from a mixture of different activities, then the distribution of sensor properties should vary based on the activity. For example, consider the activities of checking e-mail and putting a book away on the office bookshelf. Both of these activities would occur in the office and be detected by the office motion sensor; however, the length of time spent in the office would be expected to be longer for checking e-mail. Assuming that these were the only two activities, a graph showing the frequency of each length of time spent in the office would be expected to have two peaks, one representing the normal time spent for checking e-mail and the other for putting a book away. Mixture models can then be applied to the combined distribution to separate sensor firings occurring during e-mail checking from those resulting from putting a book away. With a mixture model the combined distribution is assumed to be a mixture of K different groups. The appropriate number of K clusters must be determined along with the density-function parameters for each group and the assignment of each observation to the appropriate group.

More formally, the mixture-model approach assumes that the density function for the data can be modeled as a mixture of K individual density functions, with each density function representing a separate cluster. The density function for the data can be written as

$$f(x) = \sum_{k=1}^{K} \tau_k \varphi_k(x_i \mid \mu_k, \Sigma_k). \tag{1}$$

The subscript k specifies a particular cluster, while x is a random vector for the observations. The mixing proportion is τ_k with $0 < \tau_k < 1$ for all k and $\Sigma_k \tau_k = 1$. The term φ_k represents the Gaussian density function of a particular cluster k [3].

Each cluster can be represented as a Gaussian model of type

$$\varphi_k(\mathrm{x} \mid \mu_k, \Sigma_k) = (2\pi)^{-p/2} |\Sigma_k|^{-1/2} \times \exp\left\{ -\frac{1}{2}(x_i - \mu_k)^T \Sigma_k^{-1} (x_i - \mu_k) \right\}. \quad (2)$$

The cluster has p dimensions and is centered at mean μ_k with the variance-covariance matrix, Σ_k, determining the scaling properties [2]. For the data used in this analysis, μ_k is a vector of the means of the time of day (μ_T), length of time in the room (μ_L), and activity level attributes (μ_A)

$$\mu_k = \begin{bmatrix} \mu_T \\ \mu_L \\ \mu_A \end{bmatrix}. \quad (3)$$

The covariance matrix can be decomposed into the form: $\Sigma_k = \lambda_k \Delta_k A_k D_k^T$ with the D_k matrix determining the orientation of the cluster, the A_k matrix describes the shape, the λ_k scalar determines the volume, and the superscript T denotes the transpose of the D_k matrix [2]. The orientation, shape, and volume were allowed to vary between the clusters for the analysis.

The maximum likelihood criterion was used to determine the mixing proportions and density parameters. The objective was then to maximize the log-likelihood of the parameters given the data. This is represented as

$$l(\theta \mid x_1, \ldots, x_n) = \sum_{i=1}^{N} \ln \left[\sum_{k=1}^{K} \tau_k \varphi_k(x_i \mid \mu_k, \Sigma_k) \right] \quad (4)$$

where θ denotes the set of parameters [3].

The expectation maximization (EM) algorithm was applied to calculate the parameter estimates and the assignment of observations to groups that maximized the log-likelihood [6]. This algorithm consists of an iterative two-step process. When applied to mixture modeling, the E-step involves computing the probability that observation i belongs to group k given the current parameter estimates and assigning each observation i to the group k for which the probability is the highest.

$$z_{ik} = \begin{cases} 1, & \text{if } x_i \text{ belongs to group } k \\ 0, & \text{otherwise}, \Pr(z_{ik} \mid \mu_k \Sigma_k) \end{cases}. \quad (5)$$

In the M-step, the log-likelihood criterion is used to compute the parameter estimates given the current set of assignments of observations to groups

$$l(\theta) = \sum_{i=1}^{N} \sum_{k=1}^{K} z_{ik} \log \left(\tau_k \varphi_k(x_i | \mu_k \Sigma_k) \right). \quad (6)$$

The number of groups can be determined by comparing the Bayesian information criterion (BIC) value for different numbers of groups with

$$\mathrm{BIC} = 2 \log p(x|M) + \mathrm{const} \simeq 2 l_M(x, \theta) - m_M \log(n) \quad (7)$$

where $l_M(x, \theta)$ is the maximized mixture likelihood of the model, m_M is the number of parameters estimated in the model, and n is the number of observations [6].

C. Application of Mixture Models

The EM algorithm was applied to the training set data to develop the mixture-model parameters and assign observations in the training set to clusters. The algorithm was run separately on each sensor type. The number of clusters was selected by choosing the model with the highest BIC value. The M-step of the algorithm was then used to assign the observations in the test set to the clusters. This process was repeated for the subset of data corresponding to each individual cluster. For a cluster to be further subdivided, the BIC value for two or more clusters must exceed that for one cluster. This iterative step helps prevent difficulties caused by clusters of different sizes.

After using mixture models to divide the data into clusters, the clusters representing potentially significant event types must be determined. Some groups may consist of simply random events or several different types of underlying events that may not be part of a consistent pattern of behavior. For example, if 50-s events in the kitchen occur at about the same density throughout the day, they may all be placed in one group. Such events could represent many types of underlying activities from walking through the kitchen on the way to the back door, to getting a quick snack, to checking on dinner. Random events are assumed to most likely be those of short time lengths that do not occur consistently at the same time of day.

Groups were considered to be significant if they met either a length or time-range criterion and occurred on at least 25% of the days. To determine whether the length of time was significant, the mixture-model-clustering algorithm was applied to the mean length of time values for all groups. The groups were split into two clusters with groups in the cluster with the longer lengths of time considered to be significant. A similar approach was used to determine whether the range of start times was significant. The range of start times was calculated for all groups, and mixture modeling was used to divide the groups into two clusters. The groups in the cluster with the smaller time ranges were considered to be significant.

D. Analysis of Work and Off-Days

There may be significant differences in behavioral patterns between days when the subject works and days when the subject is off and, thus, before trying to discover patterns, it was important to first categorize the days into work and off-days. However, the subject works a variable schedule, and whether a day is actually a work or off day is not known; thus, unsupervised learning methods must be applied. The main indication of a work day is a lack of activity in the house during the daytime hours. As a result, the number of motion sensor firings during the daytime hours could be used to classify the days. However, even on work days, there could be some sensor firings during the day due to occurrences, such as the opening of the front door or a sensor firing because of an outside event such as a door slamming in the basement apartment. Also, on off-days, errands might be run, leaving several hours without activity. Hence, classifying the days could not be simply based on examining the data for an 8-h period of no activity. Instead, a clustering approach was applied to the hourly counts of the number of motion-sensor firings. This allowed for grouping the days with low activity without setting an arbitrary cut-off for the number of sensor firings in a particular hour that can occur.

The number of motion-sensor firings was calculated for each discrete hour, and agglomerative hierarchical clustering using Ward's method was applied to the hours from 7 am. to 7 pm. This method begins by placing each observation in its own cluster. Clusters are then iteratively merged until all clusters are combined. Ward's method attempts to minimize the within-cluster variance by merging the clusters with the minimum sum of squares [17].

This clustering method divided the days into two primary clusters with 37 days being classified as off-days and 28 identified as work days. The mixture-model analysis was then applied separately to the work and off-days and the combined data set. Results of the mixture-model analysis on these sets were compared and are described in the next section.

IV. Mixture-Model Results

A. Metrics

Several performance metrics were used to evaluate the results of the mixture-model analysis. One performance metric for each individual cluster is the uncertainty values for the assignment of the observations to the cluster. For each cluster, the probability that the observation belongs to the cluster was calculated. The uncertainty is then one minus this probability value. Let U_k represent the uncertainty for cluster k, N represent the number of observations assigned to cluster k, and p_{ik} the probability that observation i belongs to cluster k

$$U_k = \sum_{i=1}^{N} (1 - p_{ik}). \tag{8}$$

A second metric compares the estimated probabilities of a cluster in the training and test sets. For example, based on an event space $\{A, B, \dots\}$ consisting of the different clusters, the estimated probability of event A is the number of occurrences of event $A(N_A)$ divided by the total number of event occurrences (N_all)

$$P(A) = N_A/N_\mathrm{all}. \tag{9}$$

The probability of event A for the training and test sets can then be compared using a percent error metric

$$E = |P(A)_\mathrm{train} - P(A)_\mathrm{test}|/P(A)_\mathrm{train}. \tag{10}$$

Similarly, the probability that event A occurs at least once a day can be compared between the training and test sets. In this case, the probability of event A occurring on a day was estimated by dividing the number of days on which event A occurred (D_A) by the total number of days in the data set (D_all)

$$P(A) = D_A/D_\mathrm{all}. \tag{11}$$

TABLE III
Significant Clusters That Meet Evaluation Criteria for Combined Data

% Obs. Error < 0.2	% Days Error < 0.2	Time Mean Diff < Stdev	Length Mean Diff < Stdev	Uncert. Train < 0.2	Uncert. Test < 0.2	All Criteria
25	25	43	43	29	29	12

Another metric involves comparing the parameter values of the mean length of time and mean time of day properties as calculated by the mixture-model algorithm ($\mu(L)_\mathrm{train}$ and $\mu(T)_\mathrm{train}$) to the estimated values of the parameters from the test set clusters ($\mu(L)_\mathrm{test}$ and $\mu(T)_\mathrm{test}$). The differences can be normalized by dividing by the parameter standard deviations to provide an estimate of consistency between the properties of the test set clusters and the original model

$$L_\mathrm{Diff} = |\mu(L)_\mathrm{train} - \mu(L)_\mathrm{test}|/\sigma(L)_\mathrm{train} \tag{12}$$
$$T_\mathrm{Diff} = |\mu(T)_\mathrm{train} - \mu(T)_\mathrm{test}|/\sigma(T)_\mathrm{train}. \tag{13}$$

B. Results From the Combined Data Set

The data were grouped into 139 clusters. The mean uncertainty level for each observation in the training set was 0.1016 and the standard deviation was 0.1468. For the test set, the mean uncertainty level for each observation was 0.1087 and the standard deviation was 0.1509. For the training set, 72.76% of the uncertainty values were less than 0.20 while 73.08% of the test set observations had uncertainty values less than 0.20. The percent of values exceeding 0.50 uncertainty was 1.90% for the training set and 2.14% for the test set.

Of the 139 clusters, 44 were determined to be significant. Fifteen of these 44 clusters exceeded a mean uncertainty level of 0.20, indicating a relatively high level of uncertainty for some of the clusters for the training set data. These clusters were evaluated using the test set. Comparing the percentage of days on which a cluster occurred in the training and test sets showed that many of the clusters showed a consistent rate of occurrence; however, 19 had percent errors exceeding 0.20. For the measure representing the percentage of days on which at least one observation in the cluster occurred, there were also 19 clusters with percent errors exceeding 0.20. When comparing the time of day attribute between the test set and model parameters, one cluster differed from the mean parameter by more than one standard deviation while six additional clusters differed from the mean parameter by more than one half the standard deviation. For the length of time parameter, one cluster differed from the mean parameter by more than three times the standard deviation. All of the other clusters had an estimated mean within one standard deviation of the mean parameter value with four clusters exceeding one half of the standard deviation. Similar to the training set, the mean uncertainty levels for clusters in the test set included both low and high levels of uncertainty with the mean uncertainty exceeding 0.2 for 15 clusters. Table III summarizes the number of clusters that met each of these criteria.

TABLE IV
COMPARISON OF OFF, WORK, AND COMBINED MODEL RESULTS

Data Type	Num Clust.	Num. Sign. Clust.	Sign. Clust. Pass Test	Mean Uncert. Train	Std. Dev. Uncert. Train	Mean Uncert. Test	Std. Dev. Uncert. Test
Off	111	40	11	0.07176	0.11811	0.08059	0.12786
Work	78	49	12	0.03456	0.08069	0.04820	0.10518
Combined	139	44	12	0.10161	0.14682	0.10868	0.15091

Of the significant groups, 12 had uncertainty levels under 0.20 for both the training and test sets, had error rates of less than 0.20 for the percentage of observations and percentage of days measures, and were within one standard deviation of the parameter mean for the time of day and length of time parameters. These 12 groups include four bedroom clusters, three office sensor clusters, one living room cluster, one front door sensor cluster, one laundry room cluster, and two kitchen clusters.

C. *Effect of Evaluating Work and Off-Days Separately*

When the off and work day data were analyzed separately, the mean uncertainty for both the training and test sets was less than for the combined data set, showing an improvement in performance. The mean uncertainty for the work data at 0.034 56 was especially low. Furthermore, the standard deviations of the parameter values for the separate off and work day models were in general smaller than for the combined data. As a result, despite the smaller number of observations, there were a similar number of significant clusters and significant clusters that met all the evaluation criteria. Table IV summarizes the comparison between the off-day, work-day, and combined data models.

The results of the off, work, and combined data sets were further compared on each of the evaluation criteria. The results showed a significantly higher rate of errors in the percentage of observations measure for the separate off and work day models. Only 32.5% of the significant clusters for the off-day data and 28.6% of the clusters for the work-day data had percent errors less than 0.2 for the measure. This compares to 56.8% for the combined model. Another measure where the error rates were higher for the separate models was in the comparison of parameter values to the test set mean time and mean length values. For the combined model, there was only one cluster for both attributes where the difference exceeded one standard deviation. In comparison, for the work-day data set fourteen observations exceeded a difference of one standard deviation for the time of day parameter and ten observations exceeded a difference of one standard deviation for the length of time parameter. The uncertainty metric was the main one that showed improvement from the combined model to the separate models. Tables V and VI show the number of significant clusters that met each evaluation criteria for the off and work day data sets.

V. COMPARISON TO USER LOG

An activity log from the subject was collected for a period of 37 days that did not correspond to the days used in the mixture-model analysis. The log was created by the subject using a customized PDA, where several event types, as well as the location of the activity, could be selected. The subject was asked to log activity entries in as close a time proximity to the activity event as possible. The PDA automatically stamps the entry with the correct date and time. The log consisted of 448 entries where one of 26 activity types was selected. These entries were compared to the significant clusters that met all of the evaluation criteria for the combined, off-day, and work-day data sets.

TABLE V
SIGNIFICANT CLUSTERS THAT MEET EVALUATION CRITERIA FOR OFF DAYS

% Obs. Error < 0.2	% Days Error < 0.2	Time Mean Diff < Stdev	Length Mean Diff < Stdev	Uncert. Train < 0.2	Uncert. Test < 0.2	All Criteria
13	23	38	36	36	35	11

TABLE VI
SIGNIFICANT CLUSTERS THAT MEET EVALUATION CRITERIA FOR WORK DAYS

% Obs. Error < 0.2	% Days Error < 0.2	Time Mean Diff < Stdev	Length Mean Diff < Stdev	Uncert. Train < 0.2	Uncert. Test < 0.2	All Criteria
14	28	35	39	49	47	12

The clusters were compared to the log-entry events by noting for events of the same sensor type whether the time the log-event was recorded occurred within one standard deviation of the mean start and end time for a cluster. For the combined data, 42.19% of the log-entry events occurred within one standard deviation of the 12 clusters meeting all the evaluation criteria. This percentage was 41.07% for the off-day clusters, 22.54% for the work-day clusters, and 59.15% when the off- and work-day clusters were combined. When clusters from all three data sets were used, the percentage was 66.29%. This percentage included all bedroom and front-door entries 75.0% of the office entries, 64.2% of the kitchen entries, 45.3% of the bathroom entries, 37.5% of the living room entries, and only 5.4% of the laundry room entries. These results help validate that the results of the mixture-model analysis represent many of the patterns of one event type that actually occurred.

The clusters were then compared to the activity log to determine which clusters may potentially represent an activity. Of the 12 clusters from the combined data, nine represented at least 25% of the occurrences for at least one activity type. The corresponding activities included sleep, waking-up, changing clothes, computer use, television use, and meal activities. For the off-day data, seven of the clusters corresponded to an activity with the activities including all of the previous types except television use plus toilet use. The work-day data had the lowest number of clusters corresponding to activities at six. The activities included sleep, changing clothes, toilet use, entering the house from the front door, and eating. These results show that most of the significant clusters correspond to actual event patterns.

VI. Conclusion

The results of this analysis demonstrate that a remote monitoring system with inexpensive motion sensors, such as the SmartHouse, can be used to detect behavioral patterns. The mixture-model results show that there are clusters that occur consistently over time with low classification uncertainty. Such clusters represent behavioral patterns of one event type. Comparison of these results to the user activity log shows that even though the data are from different time periods the significant clusters accurately represent many of the recorded events. Identifiable events include sleep behavior, changing clothes, bathroom/toilet use, leaving/returning home, and meal preparation, which constitute the majority of the activities of daily living that are used in functional assessments performed by healthcare professionals. Evaluating the work and off-days separately helps to reduce the uncertainty in classification of observations and provides a method for detecting event patterns that may be specific to work or off-days.

For future application, a SmartHouse-type system could be adapted to meet the needs of the individual user. Additional sensors could be used and placed where needed to monitor specific concerns. The analysis methodology described in this paper could then be utilized to provide a baseline characterization of the user's activity pattern, which could be refined through interaction with the user, and provide a reference for detecting behavioral changes. As such, the results described here provide the basis for turning the current SmartHouse technologies into ones with greater intelligence that can adapt the environment to user needs and conditions.

References

[1] A Profile of Older Americans (2001). [Online]. Available: www.aoa.gov/aoa/stats/profile

[2] J. D. Banfield and A. E. Raftery, "Model-based Gaussian and non-Gaussian clustering," *Biometrics*, vol. 43, pp. 2–18, 1992.

[3] G. Celeux and G. Govaert, "Gaussian parsimonious clustering models," *Pattern Recognit.*, vol. 28, pp. 781–793, 1995.

[4] B. G. Celler, E. D. Ilsar, and W. Earnshae, "Preliminary results of a Pilot project on remote monitoring of functional health status in the home," in *Proc. 18th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, 1996, pp. 63–64.

[5] M. Chan, C. Hariton, P. Ringeard, and E. Campo, *Smart House Automation Syst. Elderly Disabled*, pp. 1586–1589.

[6] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood for incomplete data via the EM algorithm," *J. Royal Stat. Soc., Series B*, vol. 39, pp. 1–38, 1977.

[7] C. Fraley and A. E. Raftery, "How many clusters? Which clustering method? Answers via model-based cluster analysis," *Comput. J.*, vol. 41, no. 8, pp. 578–588, 1998.

[8] ——, "MCLUST: Software for Model-Based Cluster and Discriminant Analysis,", Tech. Rep. 342, Nov. 25, 1998.

[9] G. L. Goodman, "Detection and classification for unattended ground sensors," in *Proc. Inform. Decision, Control*, Adelaide, Australia, 1999, pp. 419–424.

[10] J. A. Hartigan and M. A. Wong, "A K-means clustering algorithm," *Appl. Stat.*, vol. 28, pp. 100–108, 1979.

[11] T. Kohonen, "The self-organizing map," *Proc. IEEE*, vol. 78, no. 9, pp. 1464–1479, Sep. 1990.

[12] C. P. Mathews and K. Warwick, "Using cluster analysis techniques to improve machine reliability," in *IEE Colloquium Intell. Meas. Syst. Control Applicat.*, Apr. 4, 1995, pp. 4/1–4/3.

[13] M. Ogawa, S. Ochiai, K. Shoji, M. Nishihara, and T. Togawa, "An attempt of monitoring daily activities at home," in *Proc. 22nd Annu. EMBS Conf.*, Chicago, IL, Jul. 23–28, 2000, pp. 786–788.

[14] M. Ogawa and T. Togawa, "Monitoring daily activities and behaviors at home by using brief sensors," in *Proc. 1st Annu. Int. IEEE-EMBS Special Topic Conf. Microtechnol. Med. Biol.*, Lyon, France, Oct. 12–14, 2000, pp. 611–614.

[15] M. Ogawa, R. Suzuki, T. Izutsu, T. Iwaya, and T. Togawa, "Long term remote behavioral monitoring of elderly by using sensors installed in ordinary houses," in *Proc. 2nd Annu. Int. IEEE-EMBS Special Topic Conf. Microtechnol. Med. Biol.*, Madison, WI, May 2–4, 2002, pp. 322–325.

[16] G. Virone and N. Noury, "A system for automatic measurement of circadian activity deviations in telemedicine," *IEEE Trans. Bio-Med. Eng.*, vol. 49, no. 12, pp. 1463–1469, Dec. 2002.

[17] J. H. Ward, "Hierarchical grouping to optimize an objective function," *J. Amer. Stat. Assoc.*, vol. 58, pp. 236–244, 1963.

[18] A. Yamaguchi, M. Ogawa, T. Tamura, and T. Togawa, "Monitoring behavior in the home using positioning sensors," in *Proc. 20th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, vol. 20, 1998, pp. 1977–1979.

Tracy S. Barger received the B.S. and M.S. degrees in systems engineering from the University of Virginia, Charlottesville, in 2003.

In 2003, she began working as a Systems Engineer for Northrop Grumman, Chantilly, VA.

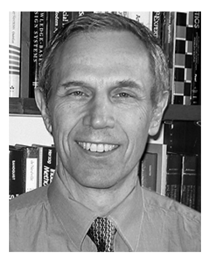

Donald E. Brown (A'86–SM'90–F'01) received the B.S. degree from the United States Military Academy, West Point, NY, in 1973, the M.S. and M.Eng. degrees in operations research from the University of California, Berkeley, in 1979, and the Ph.D. degree in industrial and operations engineering from the University of Michigan, Ann Arbor, in 1985.

He is a Professor and Chair of the Department of Systems Engineering, University of Virginia, Charlottesville. His research is statistical learning and predictive modeling.

Majd Alwan (S'94–M'96) received the B.S. degree in electrical engineering from Damascus University, Syria, in 1988, the M.S. degree in control engineering with distinction from Bradford University, Bradford, UK, in 1992, and the Ph.D. degree in intelligent robotics from the Imperial College of Science Technology and Medicine, University of London, London, UK, in 1997.

From 1988 to 1992, he was a Design Engineer with the Higher Institute of Applied Sciences and Technology (HIAST). After obtaining his Ph.D. in 1997, he returned to HIAST as a Lecturer and Researcher. In addition to lecturing on fuzzy control, control and logic systems, and digital circuits, he directed several projects and consulted with many departments, organizations, and companies on various instrumentation, automation, and computerization issues. In the beginning of 2002, he joined the Medical Automation Research Center, University of Virginia as a Manager of the Robotics and ElderCare Technologies Program.