30th Annual International IEEE EMBS Conference
Vancouver, British Columbia, Canada, August 20-24, 2008

Emergent Trend Detection in Diurnal Activity

Isaac J. Sledge, *Student Member, IEEE,* James M. Keller, *Fellow, IEEE,* Gregory L. Alexander

***Abstract*— When monitoring elders' daily routines, it is desirable to track significant deviations from a baseline pattern, as consecutive, aberrant days may foreshadow a need for medical attention. However, many traditional, unsupervised methods for pattern classification are ill-suited for this task, as they are incapable for adapting to additive datasets. To surmount this shortcoming, we establish a framework for recognizing temporal trends in feature data extracted from passive sensors.**

I. Introduction

"Aging in place" (AIP) revolves around the notion of independent, or partially assisted, living and the ability to continuously receive any necessary support for a growing gamut of needs. To further this style of care, the University of Missouri Sinclair School of Nursing and Americare Systems Inc. have collaborated to create the TigerPlace domicile complex [1]. As one of four state-approved AIP projects, TigerPlace has spawned a number of ongoing research ventures focused on improving and personalizing elder care through environment monitoring [2-3]. Out of these endeavors, one facet of the work is focused on crafting a hybrid, intelligent software-sensor system capable of providing caregivers additional information about the elders' eudemonia.

While placing a variety of devices in an individuals dwelling yields a wealth of activity information, a major issue is rooted in how to analyze the data to locate trends that correspond to states of wellbeing. Before embarking on data analysis, an important first step is the extraction of features that elucidate important information embedded within the data. Unlike the raw sensor signals, a matrix of computed features, $X = \{\vec{x}_1, \ldots, \vec{x}_n\} \subset \mathbb{R}^s$, may contain an arbitrary number of data dense regions that correspond to distinct diurnal patterns. By utilizing exploratory data analysis techniques [4], X can be organized into c self-similar subsets, called clusters, based upon an underlying similarity measure. But, since the activity of the monitored individuals, and thus the collected sensor data, can vary on a daily basis, many of these methods are insufficient for unsupervised temporal classification. To ameliorate the data exploration process, we draw upon recent work, growing neural gas clustering (GNGC) [5], from the realm of temporal clustering. In the subsequent sections, we not only show that the technique can recognize emergent activity trends in the computed sensor data features, but also provide evidence that the feature clusters align well with clinical records.

Manuscript received June 27, 2008. This work was supported in part by the National Science Foundation under ITR grant number IIS-0428420.

Isaac Sledge and James Keller are with the Electrical and Computer Engineering Department, University of Missouri, Columbia, MO 65211, USA (corresponding author: phone: +1.573.882.6387, fax: +1.573.882.0397, email: ijs3h6@missouri.edu, kellerj@missouri.edu).

Gregory Alexander is with the Sinclair School of Nursing, University of Missouri, Columbia, MO 65211, USA (email: alexanderg@missouri.edu).

II. Sensor Network & Features

As a method to gain insight into any type of data, feature calculation depends heavily on the quality and type of recorded signals. To ensure that a broad spectrum of activities is preserved for unearthing this information, various apartments at TigerPlace are outfitted with a suite of sensor elements developed by collaborators at the University of Virginia [6]. From each of the deployed networks, which include passive, infrared motion and pneumatic bed sensors, quantized room location and bed restlessness, pulse and respiration data are gathered, time-stamped and logged. Though some minutiae are lost in the coarse quantization process, and others due to communication errors, a bulk of the essential information often remains intact or can be inferred from other sensors.

Since new sensor data are constantly collected, for each participant, feature computation begins with the segmentation of the time-delimited data into 24-hour intervals. Utilizing these daily snapshots, both motion firings and bed information are fused together to generate activity density plots.

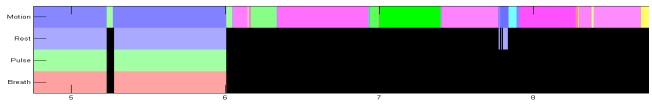

(a) Zoomed activity plot, of a 4-hour period, for an arbitrary day

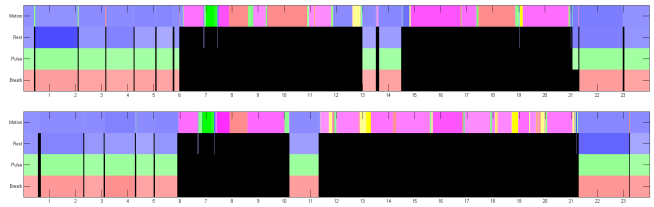

(b) Activity plots for two consecutive days

Fig. 1: Sensor density plots of the participants room location and bed restlessness, pulse and respiration, as a function of time.

In the above plots, the elder's room location, for any given moment, is encoded on the first abscissa using the color scheme: blue, green, cyan, yellow, magenta and red denote, respectively, a presence in the bedroom, bathroom, closet, kitchen, living room and entryway. Density information is then coalesced by varying the saturation of the colors based upon the aggregated motion firings; the more vivid the hue, the more sensor hits. Similarly, the remaining three axes display bed restlessness, pulse and respiration firing densities in blue, green and red; for these, black denotes bed vacancy.

Creating density images, like the ones in Fig. 1, serve a dual purpose: it not only helps to visualize trends over long time spans, but also aids in filtering out erroneous data. Furthermore, pertinent features, such as the total time spent in bed or the number of nightly bathroom visits, become easily discernible. Other attributes can also be reaped from the ac-

978-1-4244-1815-2/08/$25.00 ©2008 IEEE.

Authorized licensed use limited to: University of Missouri. Downloaded on January 15, 2009 at 17:00 from IEEE Xplore. Restrictions apply.

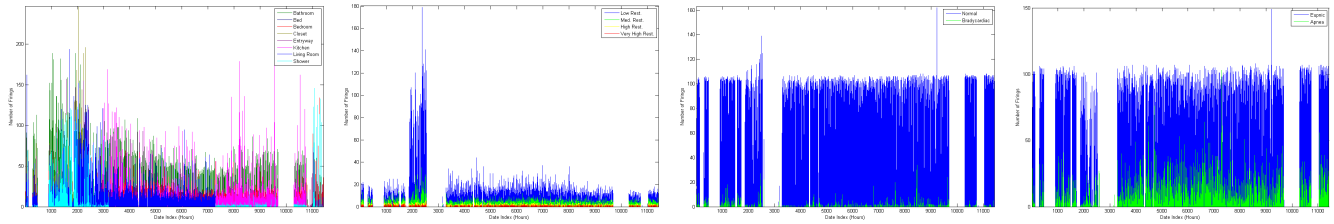

(a) Plot of motion data (b) Plot of bed restlessness data (c) Plot of bed pulse data (d) Plot of bed respiration data

Fig. 2: Plots of the raw sensor data, in hourly units, for Participant I from 10/10/2005 to 01/29/2007. In (a), the hourly-aggregated firings for eight different areas: bathroom (green), bed (dark blue), bedroom (red), closet (gold), entryway (purple), kitchen (magenta), living room (blue) and shower (cyan), are shown. In (b), the hourly-summed low (blue), medium (green), high (yellow) and very high (red) restlessness are plotted as a function of time. Similarly, (c) displays the number of normal (blue) and bradycardiac (green) firings, while (d) shows the number of eupnic (blue) and apneic (green) sensor hits.

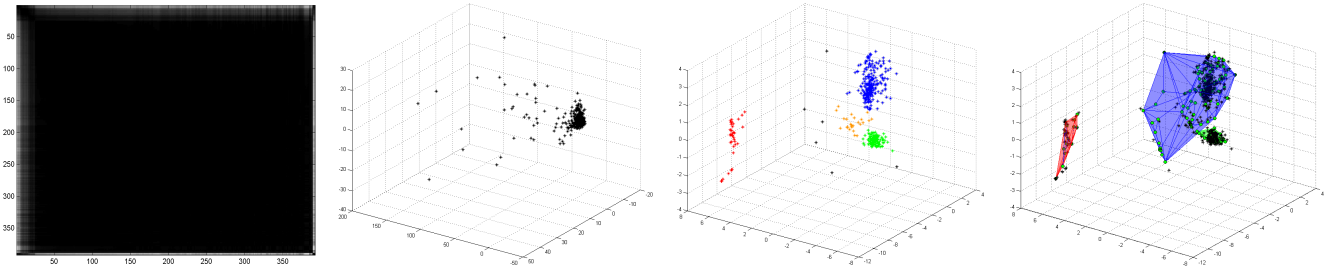

(a) VAT of features, $X_o^t \subset \mathbb{R}^8$ (b) Original features, $X_o^t \subset \mathbb{R}^8 \downarrow \mathbb{R}^3$ (c) Revised features, $X_r^t \subset \mathbb{R}^{32} \downarrow \mathbb{R}^3$ (d) X_r^t with $\vec{w}_k$'s and convex hulls

Fig. 3: Plots of the feature data, in daily units, for Participant I from 10/10/2005 to 01/29/2007.

tivity graphs, and include the following: time woke up/went to bed, number of times out of bed in the morning/night, total time spent out of bed during those trips, total amount of nap time, aggregated bed restlessness, pulse and respiration sensor firings, number of daily bathroom visits, time spent in each room, aggregated motion sensor firings, and the number of room changes (puttering index). In total, 32 different activity characteristics, are currently measured for each 24-hour period; others, such as visitor information and time spent out of the apartment, are being explored for future inclusion.

III. Temporal Clustering

As new feature vectors are appended to X, the possibility exists for new data dense regions to form or even amalgamate over time. Due to these transient changes in topology, conventional pattern recognition techniques, such as clustering, are unsuitable for this task as they are reliant on a specified class count, while the classification results are localized in time. However, at least one method, growing neural gas clustering, is able to accurately capture the temporal activity distributions formed by the features.

Unlike conventional clustering, which normally involves optimizing the relationship between the fixed prototype set, $V = \{\vec{v}_1, ..., \vec{v}_c\} \subset \mathbb{R}^s$, and the set of objects, GNGC draws upon concepts from a number of different fields to locate cluster prototypes. Foremost, vector quantization is used to encode each manifold, $M \subseteq \mathbb{R}^s$, of signals using a set of reference vectors, $W = \{\vec{w}_1, ..., \vec{w}_m\} \subset \mathbb{R}^s$. To accommodate the inclusion of new feature vectors, the size of W is allowed to grow as a function of the number of added attributes. A hybrid growing neural gas [7] and adaptive resonant theory [8] scheme is then utilized to update the best-matching $\vec{w}_k$, and its connected neighbors on a dynamic lattice structure, for each input stimulus. Since there are no explicit constraints on the lattices topological arrangement, new connections can be forged, between arbitrary, non-connected $\vec{w}_k$'s, based on the induced magnitude response of each $\vec{w}_k$'s receptive fields. In addition, obsolete connections are allowed to die out, due to an 'aging' factor. By exploiting this behavior, the number of clusters, at a given time instant, can be determined by isolating non-connected lattices and finding the number of unique, non-Hamiltonian graph paths. Combining this with computational geometry concepts, such as convex hull computation and point-in-polygon tests, cluster shapes and centers can be determined. Crisp, fuzzy, or possibilistic membership values are then assigned to each datum in X using the distance metric that corresponds to each clusters shape. Provided that there is a constant stream of data, these membership partitions are updated ad infinitum.

For additional details about this algorithm, with synthetic dataset experiments, see Sledge and Keller [5].

IV. Activity Analysis

With the dearth of any parameter decay, the ability to iteratively add new neuronal reference vectors to the network, and an incremental style of learning, GNGC is particularly attractive for temporal activity analysis. As such, we test its effectiveness, along with the quality of the computed features, in a series of case studies using data collected by the TigerPlace sensor network system. To aid in annotating the exposed trends, we make use of medical records and assessments of the participants' wellbeing collected by registered nurses during clinical interviews.

A. Case Study - Participant I

Over the course of multiple iterations, the current feature set, presented in Sec. II, has evolved from a much earlier subset of 8 characteristics. For each generation of attributes,

Authorized licensed use limited to: University of Missouri. Downloaded on January 15, 2009 at 17:00 from IEEE Xplore. Restrictions apply.

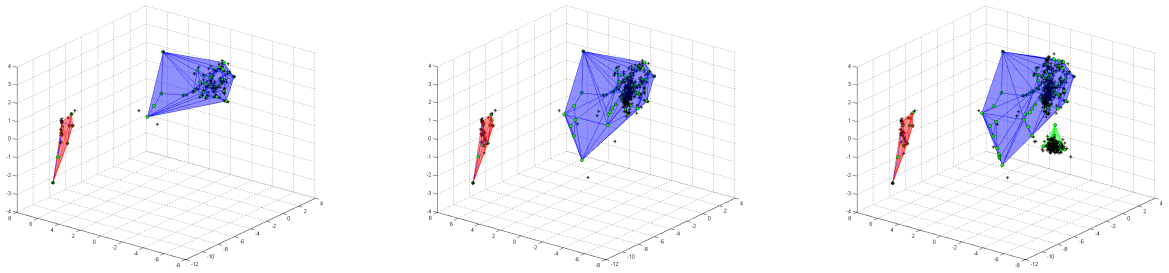

(a) $X_r^t \subset \mathbb{R}^{32} \downarrow \mathbb{R}^3$ from 10/10/2005 to 01/18/2006

(b) $X_r^t \subset \mathbb{R}^{32} \downarrow \mathbb{R}^3$ from 10/10/2005 to 06/17/2006

(c) $X_r^t \subset \mathbb{R}^{32} \downarrow \mathbb{R}^3$ from 10/10/2005 to 11/14/2006

Fig. 4: Temporal plots of the GNGC clustering results, for Participant I, with 40, 60, and 80 neuronal references, respectively. As more data is added to X_r^t, such as in (c) and (d), the $\vec{w}_k$'s, shown using filled green spheres, adapt their location and neighborhood connectivity to better model the data.

the quality benchmark has been both the recognition of trends in stored data along with any future patterns that may arise. To help probe for these tendencies, both activity density plots and physiological graphs, such as those in Fig. 2(a)-(d), are used. Viewing the first two plots in Fig. 2, a number of conspicuous, distinct patterns immediately emerge: a large, abnormal spike in bed restlessness, which occurred after an ER visit, a slightly decreasing, multimodal distribution after the spike, and an overall decrease in motion firings over time, possibly from the elder spending more time out of the apartment. Though there is some correlation between the large restlessness peak and the pulse and respiration data in Fig. 2(c)-(d), these two plots did not play a major role in this example.

At the conception of the feature extraction process, it was uncertain what type of activity clusters would egress from the physiological data. To aid in visualizing the attributes, we used both VAT [9], which yielded the reordered pairwise dissimilarity image in Fig. 3(a), and principle component analysis (PCA), which produced the projection plot shown in Fig. 3(b). Viewing these two images, we found that many of the daily attributes clumped together in a single region. With only a scarce number of outliers, denoting days during the large restlessness peak, it was apparent that these original features were insufficient in emphasizing all of the visually perceptible trends from Fig. 2(a)-(b).

Spurring the evolution of the feature set, the number, and types, of measured characteristics was updated to address these shortcomings. Upon projecting the revised feature set, which was made up of the features listed in Sec. II, into $\mathbb{R}^3$ using PCA, we found that a number of data dense regions formed. One of these clusters, the elongated cluster, shown in Fig. 3(c) using blue, formed in the beginning and was indicative of the elder's "normal" baseline. A second cluster, the wispy strand of red points, denoted heavily abnormal behavior, which was a culmination of both the large restlessness spike and the period of bed inactivity, yet motion activity, that ensuingly occurred. Similarly, a third activity cluster, highlighted in green, captured the decrease in motion firings. This new cluster became the dominant baseline, for a time, until near the end of the captured data. At this point, an amber distribution arose, which coincided with hospice caregivers entering and leaving the apartment.

Given the vast improvement for disinterring a variety of activity trends, the additive, PCA-reduced data, $X_r^t \subset \mathbb{R}^{32} \downarrow \mathbb{R}^3$, was fed, a day at a time, to the GNGC algorithm. Over multiple iterations, as shown in Fig. 4(a)-(c), GNGC updated the spatial location of the reference vectors and found that 3 clusters, shown in Fig. 3(d), emerged. Unfortunately, the amber distribution, in Fig. 3(c), eluded detection due to the small number of data points and its sparse nature. However, as a testament to the algorithms capabilities, when the non-projected data was used, we found that the hospice worker cluster and several outlier points formed a number of isolated groups, thus driving up the cluster count. The only drawback of clustering up in 32-space, with GNGC, is that we lose the ability to directly display the learned distributions; recently, however, we have contrived a suitable visualization [10].

B. Case Study - Participant II

In contrast to trends found in the previous study, those in Fig. 5(a)-(d) are less pronounced. Delving through medical records, we found that the elder had a total knee replacement a few months after the sensor data collection began. Com-paring the event with the plots, we see that, in the beginning, there are a large number of bed restlessness sensor firings, due to the resident constantly readjusting; however, as ex-pected, the levels slowly died down. After the surgery, there is also a surge in bradycardiac (1-30 BPM) pulse firings, in Fig. 5(c), that decreases much later in the data collection process. This initial increase in bradycardia is unexpected, as the individual's heart rate is likely to increase episodically, after the surgery, due to increased pain from bed movement. Viewing Fig. 5(d), a decaying trend is also apparent, as the decrease of eupnic and apneic sensor firings coincides with the drop in low pulse firings.

As with the previous example, we consider the amount of information captured in the updated features versus the original set. Projecting the original attributes into 3-space using PCA, the plot in Fig. 6(a) portrays a distribution that varies greatly in comparison to the one shown in the previous study. Nonetheless, since many of the points are strewn in the reduced feature space, and none of the outlined trends are visible, we again conclude that these features are insufficient.

Turning now to the revised feature set in Fig. 6(b), we see that there are three major activity distributions present in the data: a long, thin strand of red-highlighted points, a large blue cluster, and a sparse group of green data. Much like the red distribution in Fig. 3(c), the one in Fig. 6(b) corresponds to

Authorized licensed use limited to: University of Missouri. Downloaded on January 15, 2009 at 17:00 from IEEE Xplore. Restrictions apply.

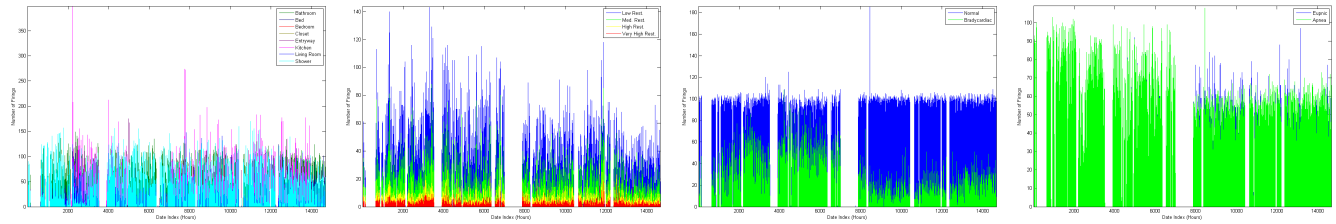

(a) Plot of motion data (b) Plot of bed restlessness data (c) Plot of bed pulse data (d) Plot of bed respiration data

Fig. 5: Plots of the raw sensor data, in hourly units, for Participant II from 11/11/2005 to 07/16/2007.

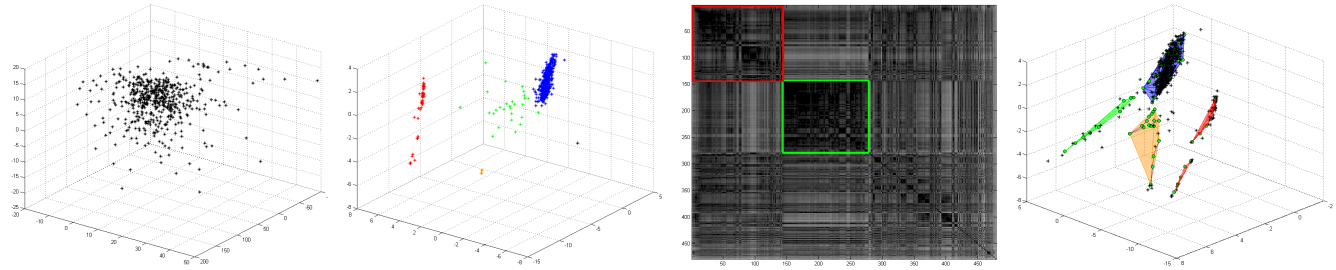

(a) Original features, $X_o^t \subset \mathbb{R}^8 \downarrow \mathbb{R}^3$ (b) Revised features, $X_r^t \subset \mathbb{R}^{32} \downarrow \mathbb{R}^3$ (c) VAT of blue cluster from (b) (d) X_r^t with $\vec{w}_k$'s and convex hulls

Fig. 6: Plots of the feature data, in daily units, for Participant II from 10/10/2005 to 07/16/2007.

the days where there was little, to no, restlessness, pulse or respiration data and only motion firings; this may indicate that the elder is sleeping on a couch or in a reclining chair. Beyond the red cluster, the sparse green and amber distributions are indicative of aberrant days. Similarly, the outstretched blue distribution is related to the individuals baseline; however, contained inside this single cluster are actually two groups that form over time. The first, which is located in the lower half of the distribution, is associated with the high restlessness, pulse and respiration firings, while the second formed near the upper half after the firing levels dropped. To see if these temporal clusters are also closely related in $\mathbb{R}^{32}$, we produced the VAT image, shown in Fig. 6(c). By concentrating only on the data in the blue distribution, two dark blocks, highlighted in red and green, formed in the VAT image and exposed the aforementioned inter-cluster structure. Given that the motion attributes dominate the feature set, in terms of the total number of motion characteristics, this may explain why the two clusters are present in $\mathbb{R}^{32}$, but not visually separated in the projected plot.

Despite the potential loss in cluster structure when decorrelating the data, we again fed the projected dataset, $X_r^t \subset \mathbb{R}^{32} \downarrow \mathbb{R}^3$, to GNGC, one datum at a time. After the entire dataset had been presented, and the neuronal references stabilized, the algorithm reported that five clusters, shown in Fig. 6(d), materialized. Though we originally surmised that only four distributions existed in the data, the fifth temporal cluster arose when the neural connections in the red group were severed to create two abnormal behavior classes. Since these results mesh well with the expected ones, we also clustered the non-projected data and found that, as with the previous study, a number of new clusters were found in $\mathbb{R}^{32}$. Projecting the neuronal references down into $\mathbb{R}^3$, we found that the blue distribution, in Fig. 6(b), split into two clusters. As such, the detection of these emergent groups provides evidence that GNGC can function well for locating trends.

V. Future Research

With the ability to discover forming distributions, growing neural gas clustering is an integral part of the adaptive, intelligent software-sensor system currently under construction. To aid in assessing well-being from the exposed trends, we plan on coalescing the cluster partitions with a fuzzy logic classifier. Unlike previous research in elder environment monitoring [11-13], the completed system will be able to detect baseline changes and ascribe linguistic descriptions to each day. Furthermore, it will only require a minimal amount of human intervention and data interpretation, making it rather attractive for simultaneously tracking the wellbeing of multiple individuals.

References

[1] M. Rantz, et al., "TigerPlace: A New Future for Older Adults", *Journal of Nursing Care Quality*, vol. 20, no. 1, pp. 1-4, 2005.

[2] D. Anderson, et al., "Modeling Human Activity From Voxel Person Using Fuzzy Logic", *IEEE Trans. Fuzzy Systems*, submitted, 2007.

[3] S. Wang and M. Skubic, "Density Map Visualization from Motion Sensors for Monitoring Activity Level", *IET Proc., IE*, 2008.

[4] J. Bezdek, J. Keller, R. Krishnapuram, N. Pal, *Fuzzy Models and Algorithms for Pattern Recognition and Image Processing*, Kluwer, 1999.

[5] I. Sledge and J. Keller, "Growing Neural Gas for Temporal Clustering", *IEEE Proc., ICPR*, submitted, 2008.

[6] D. Mack, et al., "A Passive and Portable System for Monitoring Heart Rate and Detecting Sleep Apnea and Arousals: Preliminary Validation", *IEEE Proc., D2H2*, 2006.

[7] B. Fritzke, "A Growing Neural Gas Network Learns Topologies", *Advances in Neural Inform. Processing Sys.*, vol. 7, pp. 625-632, 1994.

[8] G. Carpenter and S. Grossberg, "Adaptive Resonance Theory", *The Handbook of Brain Theory and Neural Networks*, MIT Press, 2003.

[9] J. Bezdek and R. Hathaway, "VAT: A tool for visual assessment of (cluster) tendency", *IEEE Proc., IJCNN*, 2002.

[10] I. Sledge, et al., "Temporal Activity Analysis", *Proc., AAAI*, submitted, 2008.

[11] M. Ogawa, et al., "Long-Term Remote Behavioral Monitoring of the Elderly using Sensors Installed in Domestic Houses", *IEEE Proc., EMBS/BMES*, 2002.

[12] N. Barnes, et al., "Lifestyle Monitoring: Technology for Supporting Independence", *Computing and Control Eng. J.*, pp. 169-174, 1998.

[13] S. Brown, et al., "Developing a Well-being Monitoring System Modeling and Data Analysis Techniques", IOS Press, 2006.

Authorized licensed use limited to: University of Missouri. Downloaded on January 15, 2009 at 17:00 from IEEE Xplore. Restrictions apply.