Temporal Activity Analysis

Isaac J. Sledge[1], **James M. Keller**[1], **Timothy C. Havens**[1], **Gregory L. Alexander**[2], **Marge Skubic**[1]

[1] University of Missouri, Electrical and Computer Engineering Department
[2] University of Missouri, Sinclair School of Nursing
{ijs3h6, kellerj, tch9q3, alexanderg, skubicm}@missouri.edu

Abstract

When monitoring elders daily routines, it is desirable to detect aberrant activity trends, as they may foreshadow a need for medical attention. But, traditional, unsupervised pattern classification techniques are ill-suited for this task, because the data distributions formed by the captured patterns are temporal in nature. To overcome this algorithmic deficit, we craft a framework for analyzing and displaying additive trends in feature data extracted from passive sensors.

Introduction

By 2030, the elderly population is expected to double. Without a firm, existing caregiving framework to accommodate the influx of individuals, faculty and researchers at the University of Missouri Sinclair School of Nursing have adopted an Aging in Place (AIP) style of care (Marek, Rantz and Porter 2004). As an alternative to institutionalized, long-term oversight, AIP revolves around the notion of independent, at home living and the ability for elders to continuously receive support for their growing gamut of needs. But, the model alone is not sufficient to combat the growing financial burden for caregiving services (Knickman and Snell 2002) along with staving off a degenerate functional decline. Instead, automated, intelligent systems could be coupled with the approach to reduce costs and improve the quality of life.

To help promote this technologically enhanced, independent living model, the MU Sinclair School of Nursing and Americare Systems Inc. have collaborated to create the TigerPlace domicile complex (Rantz *et al*. 2005). 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 (Anderson *et al*. 2007; Luke *et al*. 2007; Sledge, Keller and Alexander 2008; Wang and Skubic 2008). Out of these endeavors, one facet of the work is focused on crafting a hybrid software-sensor system capable of providing caregivers additional information about the elders' wellbeing.

While placing a variety of sensors 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. However, 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, $\mathbf{X} = \{\vec{x}_1, \ldots, \vec{x}_n\} \subset \mathbb{R}^s$, may contain an arbitrary number of data dense regions that correspond to distinct diurnal patterns. By utilizing exploratory data analysis techniques (Bezdek *et al*. 1999), $\mathbf{X}$ 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), from the realm of temporal clustering (Sledge and Keller 2008). In the subsequent sections, we not only show that the technique can locate emergent activity trends in the computed sensor data features, but also provide a mechanism for viewing the learned groups in $\mathbb{R}^s$.

Feature Extraction

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 (Mack *et al*. 2006; Alwan *et al*. 2003). In each of the deployed networks, infrared motion sensors record both quantized activity information and room location. Pneumatic strips, placed under the bed linens, are also used to gather a variety of bed-related data, such as restlessness, pulse and respiration. However, unlike the binary motion sensors, the bed sensor returns level-thresholded, physiological data. For example, if the resident is readjusting, the sensor will initially report low restlessness. Provided that the device continues to detect movement, and that a sufficient amount of time has elapsed, medium, high or very high firings are generated. As well, respiration and pulse related data are captured, at regular intervals; but, instead of escalating firing levels, rough estimates of the breaths and heart beats per minute are 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 segmenta-

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

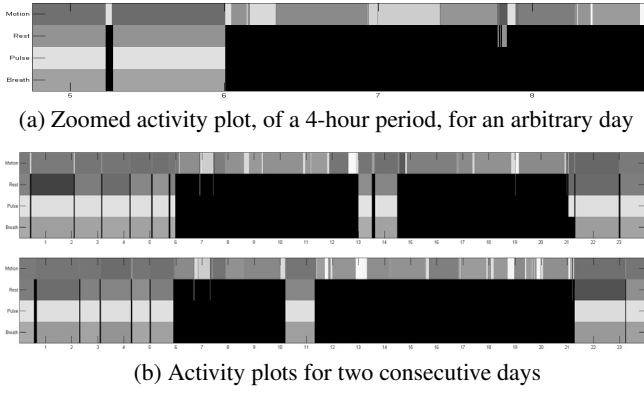(b) Activity plots for two consecutive days

Figure 1: Sensor density plot of the participants color-coded room location (first row) and bed restlessness (second row), pulse (third row) and respiration (fourth row), as a function of time. The higher the saturation, the more sensor firings that occurred in a particular area before the elder moved to another; black denotes bed vacancy. To see image details, please zoom in.

tion of the time-delimited data into 24-hour intervals. Using these daily snapshots, both motion firings and bed information are fused together to generate activity density plots, such as those in Fig. 1. In Fig. 1(a)-(b), the individuals room location, for any given mo-ment, is encoded on the first ordinate using the color scheme: blue, green, cyan, yellow, magenta and red denote presence in the bedroom, bathroom, closet, kitchen, living room and entryway, respectively. Density information is then coa-lesced 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 ordinates display bed restlessness, pulse and respiration firing densities in blue, green and red. For each of these three axes, black is used for periods of inactivity and is an indication of bed vacancy.

Creating density images, like those in Fig. 1(b), 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 activity graphs, and a complete list is given in Table 1. In total, 32 different 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.

Feature Extraction

As new feature vectors are appended to $\mathbf{X}$, 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 rely on a specified class count, c, while the classification results are localized in time. Although an estimate of the number of coherent groups, c_{est}, can be obtained using tendency assessment approaches (Bezdek *et al.* 1999; Bezdek and Hathaway 2002;

Table 1: Current set of daily computed features

Feature Type	Measured Quantity
Motion	Number of nightly bathroom visits, time woke up/went to bed, number of times out of bed in the morning/night, time out of bed during those trips, number of daily bathroom visits, time spent in each room number of room changed, aggregated motion firings (8 areas)
Restlessness	Total time in bed, total amount of nap time, aggregated restlessness firings (4 levels)
Pulse	Aggregated pulse firings (2 levels)
Respiration	Aggregated breath firings (2 levels)

Havens *et al.* 2008; Sledge *et al.* 2008), and the classification method re-executed, it is desirable to not only glean c_{est} automatically but also reuse previous clustering results. To realize these algorithmic desires, Sledge and Keller (2008) proposed growing neural gas clustering, which is capable of capturing temporal distributions formed by additive datasets.

Clustering Preliminaries

During traditional, exploratory data analysis, clustering algorithms attempt to optimize the spatial location of a prototype set, $\mathbf{V} = \{\vec{v}_1, ..., \vec{v}_c\} \subset \mathbb{R}^s$, w.r.t. the set of unlabeled objects, $\mathbf{X}$. Once the $\vec{v}_j$'s have captured a compact representation of the clusters' structure, a membership partition, $\mathbf{U} = [u_{j,i}]_{(c \times n)}$, is returned, which succinctly describes the commitment of every point to each of the c groups. For classical (hard) clustering, a vector is required to have full commitment, or membership, to a single cluster. However, in fuzzy clustering, a datum can belong to any number of groups, provided that the sum of the memberships, in each cluster, adds up to one. Possibilistic algorithms further build upon the idea of shared belongingness by relaxing the fuzzy summation constraint; this allows for a vector's summed commitment to be greater than one. Since the c-partitions of $\mathbf{X}$ can be either soft or hard, the nested sets of all non-degenerate c-partitions are represented as:

$$\mathbf{M}_{\text{pcn}} = \left\{ \mathbf{U} \in \mathbb{R}^{c \times n} \; \middle| \; \begin{array}{c} u_{j,i} \in [0,1], \forall j, i; \sum_{i=1}^{n} u_{j,i} \leq n, \\ \forall j; \max_j u_{j,i} > 0, \forall i \end{array} \right\}$$

$$\mathbf{M}_{\text{fcn}} = \left\{ \mathbf{U} \in \mathbf{M}_{\text{pcn}} \; \middle| \; \begin{array}{c} u_{j,i} \in [0,1], \forall j, i; \sum_{i=1}^{n} u_{j,i} < n, \\ \forall j; \sum_{j=1}^{c} u_{j,i} = 1, \forall i \end{array} \right\}$$

$$\mathbf{M}_{\text{hcn}} = \left\{ \mathbf{U} \in \mathbf{M}_{\text{fcn}} \; \middle| \; \begin{array}{c} u_{j,i} \in \{0,1\}, \forall j, i; \sum_{i=1}^{n} u_{j,i} < n, \\ \forall j; \sum_{j=1}^{c} u_{j,i} = 1, \forall i \end{array} \right\}$$

where $\mathbf{M}_{\text{pcn}}$, $\mathbf{M}_{\text{fcn}}$, and $\mathbf{M}_{\text{hcn}}$ are possibilistic, fuzzy and hard, respectively (Bezdek *et al.* 1999).

Temporal Clustering

Unlike the popular c-means family of methods, which use a least squares optimization approach to update the positions of the prototypes, GNGC draws upon concepts from a number of different fields to locate cluster centroids. Foremost,

learning vector quantization is used to encode each manifold, $\mathcal{M} \subseteq \mathbb{R}^s$, of signals using a finite set of reference vectors, $\mathbf{W} = \{\vec{w}_1, ...\} \subset \mathbb{R}^s$. To accommodate the inclusion of new feature vectors, the size of $\mathbf{W}$ is allowed to grow as a function of the number of added data points. A hybrid growing neural gas (Fritzke 1994) and adaptive resonant theory (Carpenter and Grossberg 2003) 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. Provided that there is a constant stream of data, the $\vec{w}_k$'s are continuously updated.

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 graph paths. Combining this with computational geometry concepts, such as convex hull computation and point-in-polygon tests, cluster centers can be determined by: $\vec{v}_j = \sum_{i=1}^{k} \vec{x}_i^j / k$ where $\vec{x}_i^j$ is the i-th object contained inside the j-th convex polytope, $\mathcal{P}_j \in \mathbb{R}^s$, and k is the total number of points inside $\mathcal{P}_j$. If, however, there is not a sufficient amount of references to create a convex polytope, e.g. the number of $\vec{w}_k$'s in the j-th graph, $\mathcal{G}_j$, is less than $s + 1$, then the centroid is found to be: $\vec{v}_j = \sum_{i=1}^{k} \vec{w}_i^j / k$, where k is the number of neuronal references, or $\vec{w}_k$'s, in $\mathcal{G}_j$.

Since GNGC is not constrained to using a specific prototypical shape, such as a point, hyerplane, hypershell, etc., the shape of each cluster can be determined through a series of tests, provided $\mathcal{P}_j$ can be found. Utilizing this information, possibilistic, (1), or fuzzy, (2), membership values can then be assigned to each datum in $\mathbf{X}$, subject to the constraints outlined in $\mathbf{M}_{\text{pcn}}$ and $\mathbf{M}_{\text{fcn}}$:

$$u_{j,i} = \left(1 + \left(d_{j,i}^2/\eta_j\right)^{1/m-1}\right)^{-1}, \forall j, i \qquad (1)$$

$$u_{j,i} = \left(\sum_{q=1}^{c} \left(d_{j,i}/d_{q,j}\right)^{2/m-1}\right)^{-1}, \forall j, i \qquad (2)$$

where $d_{j,i}$ is the distance from $\vec{x}_i$ to the j-th prototype, $\vec{v}_j$, $m \geq 1$ a degree of fuzzification and η_j the distance at which the membership of a point becomes 0.5. For hypershellular clusters, once the centroid of the cluster is found, $\mathbf{U}$ is formed using (1) or (2) with radial distance measure:

$$d_{j,i}^2 = \left(\left(||\vec{x}_i - \vec{v}_j||_{\mathbf{A}_j} - 1\right)^2 ||\vec{x}_i - \vec{v}_j||^2\right) \Big/ ||\vec{x}_i - \vec{v}_j||_{\mathbf{A}_j}^2$$
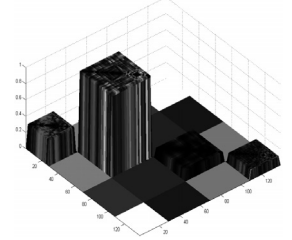
where $\mathbf{A}_j$ is a positive, definite, symmetric matrix accounting for the eccentricity and orientation of the hypershell (Bezdek et al. 1999). The elements of $\mathbf{A}_j$, for this equation, are approximated by finding the Löwner hyperellipsoid, of the convex polytope, using the Khachiyan algorithm (Khachiyan and Todd 1993). Likewise, points in hyperplanar structures are assigned membership using (1) or (2), where:
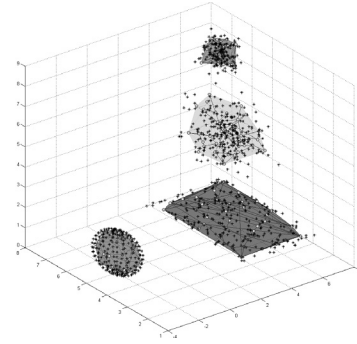
(a) Synthetic data, $\mathbf{X} \subset \mathbb{R}^3$, with learned $\vec{w}_k$'s, denoted using filled green spheres, and colored convex hulls
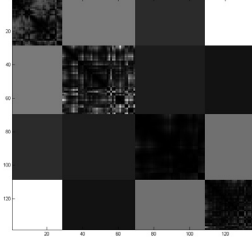
(b) Corresponding neuronal dissimilarity image for (a)   (c) NerDI with normalized polytope volume information

Figure 2: GNGC results showing four clusters, along the main diagonal of (b) and (c), and the spatial similarities between the learned distributions, using gray-scale values. In (b) and (c), the purple cluster is near the green cluster, somewhat close to the blue planar cluster, and far from the red spherical cluster. In (c), the purple cluster has the smallest convex polytope volume. To see image details, please zoom in.

$$d_{j,i}^2 = ||\vec{x}_i - \vec{v}_j||_{\mathbf{A}_j}^2 - \sum_{k=1}^{2} \langle \vec{x}_i - \vec{v}_j, \vec{b}_{j,k} \rangle_{\mathbf{A}_j}^2$$

is the orthogonal distance from $\vec{x}_i$ to the j-th variety, in $\mathbb{R}^s$, and $\mathbf{A}_j$ is an arbitrarily defined positive, definite, symmetric weight matrix (Bezdek et al. 1999). In this equation, $\vec{b}_{j,k}$ is estimated by finding the eigenvectors of the neuronal references that model the linear variety. Finally, for point-cloud clusters, $d_{j,i}$ is found as the distance between $\vec{x}_i$ and the j-th point centroid; this cluster type is the default case when $\mathcal{P}_j$ cannot be computed.

High Dimensional Visualization

When working with low dimensional data, say $\mathbf{X} \subset \mathbb{R}^3$, it is easy to display intermittent GNGC results, such as those in Fig. 2(a). However, if the dimensionality of the dataset grows beyond $\mathbb{R}^3$, capturing the same spatial information and visualizing the learned distributions is problematic. To rectify this issue, several conventions were borrowed from the visual assessment of [cluster] tendency (VAT) algorithm (Bezdek and Hathaway 2002). In VAT, a matrix, $\mathbf{R} = [r_{j,i}]_{n \times n}$, of normalized, pair-wise dissimilarity values, are ordered using a modified variant of Prim's

algorithm for finding a minimal spanning tree (MST). After the sorting process, if the matrix is displayed as an intensity image then cluster structure is indicated by the presence of dark blocks along the main diagonal.

By modifying the VAT approach, to instead use information obtained from GNGC, we can create images like those in Figs. 2(b) and 2(c). These plots, which we call neuronal dissimilarity images (NerDI), are generated by first computing the normalized, pair-wise dissimilarity of the neurons in each isolated graph. Utilizing Prim's modified MST process, the dissimilarities are rearranged, for each $\mathcal{G}$, which aids in looking for dense, intra-cluster distributions of neurons. Each of these sub-matrices are then colored and placed along the main diagonal of the NerDI. Inter-cluster spatial relationships are also incorporated in the image by finding the minimal distance between the neurons in each $\mathcal{G}$; these values are then normalized so that white denotes the largest distance between two clusters, in $\mathbb{R}^s$. Finally, volume or clusterness (Keller and Sledge 2007) information can be added, as a third dimension, to provide additional insight about the approximated manifolds.

Temporal Activity Analysis

With the ability to iteratively add new neuronal reference vectors, and an incremental style of learning, GNGC is particularly attractive for temporal analysis. As such, we tested its effectiveness in locating both gradually changing and sud-den, fluctuating activity changes through a series of case studies using data collected by the TigerPlace sensor network system. To aid in annotating the exposed trends, we made use of medical records and assessments of the participants' wellbeing collected by registered nurses and social workers during clinical interviews.

Case Study - Participant I

Over the course of multiple iterations, the current feature set, outlined in the second section, has evolved from a much earlier subset of characteristics. For each generation of attributes, 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, activity density plots and physiological graphs, such as those in Fig. 3(a)-(d), are used. Viewing the first two plots in Fig. 3, a number of conspicuous patterns emerge: 1) a large, abnormal spike in bed restlessness, which occurred after an ER visit, 2) a slightly decreasing, multimodal distribution after the spike, 3) 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. 3(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. By conducting a series of studies (Sledge, Keller and Alexander 2008), we found that the current features, listed in Table 1, highlighted several activity trends that were not present in previous sets. To visualize the

(a) Original feature set, $\mathbf{X}_o \subset \mathbb{R}^8 \downarrow \mathbb{R}^3$

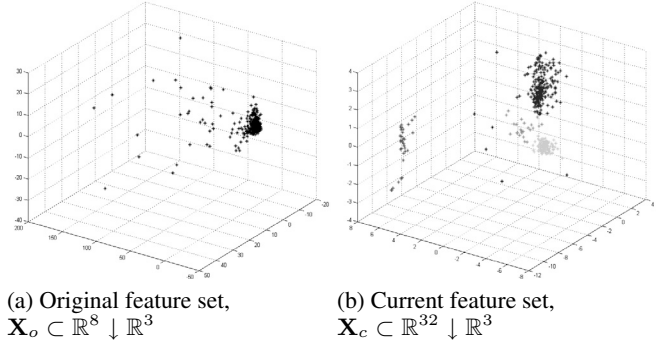(b) Current feature set, $\mathbf{X}_c \subset \mathbb{R}^{32} \downarrow \mathbb{R}^3$

Figure 4: Plots of the first three principle components from the feature data for Participant I from 10/10/2005 to 01/29/2007. For both plots, each of the points corresponds to a single 24-hour period. In (b), the current features, $\mathbf{X}_c$, are those listed in Table 1. To see image details, please zoom in.

differences, principle component analysis (PCA) was used, which produced the projection plots shown in Fig. 4. Viewing the first image, Fig. 4(a), we discovered 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. 3(a)-(b). But, upon projecting the current feature set, $\mathbf{X}_c$, into $\mathbb{R}^3$, we found that a number of data dense regions formed. One of these clusters, the elongated blue cluster, shown in Fig. 4(b), 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 from Fig. 3(a). This new cluster became the dominant baseline, for a time, until near the end of the recorded data. At this point, the amber distribution arose, which coincided with hospice caregivers entering and leaving the apartment.

Given the vast improvement for locating a variety of activity trends, we first fed the additive, PCA-reduced features, $\mathbf{X}_c \subset \mathbb{R}^{32} \downarrow \mathbb{R}^3$, a day at a time, to the GNGC algorithm. Over multiple iterations, as exhibited in Fig. 5(a)-(d), GNGC updated the spatial location of the reference vectors and found that 3 clusters, shown in Fig. 5(d), emerged. Unfortunately, the amber distribution, in Fig. 4(b), eluded detection due to the small number of data points and its sparse nature.

Though these results are consistent with our expectations, given that information is lost in the reduction process, they do not highlight the true abilities of temporal clustering. To fully measure the algorithms capabilities, the non-projected data, $\mathbf{X}_c \subset \mathbb{R}^{32}$, was iteratively introduced to GNGC. Once the algorithm stabilized, for each batch of new features, a series of NerDI plots, in Fig. 6(a)-(d), were generated. Comparing these clustering results with those in Fig. 5, it is ev-

(a) Plot of motion data   (b) Plot of bed restlessness data   (c) Plot of bed pulse data   (d) Plot of bed respiration data
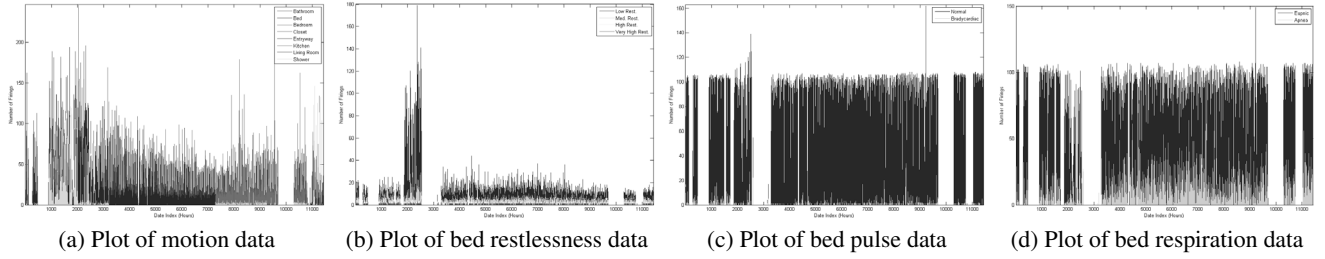
Figure 3: Plots of the aggregated, raw sensor data, in hourly units, for Participant I from 10/10/2005 to 01/29/2007. In (a), the hourly-aggregated firings for eight different areas: bathroom (green), bed (dark blue), bedroom (red), closet (gold), entryway (purple), kitchen (magenta), living room (blue) 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. To see image details, please zoom in.
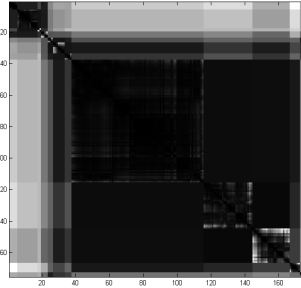

Figure 7: VAT reordered NerDI from Fig. 6(d). Here, the $\mathcal{G}_j$'s in the upper, left corner of the NerDI correspond to the abnormal (red) distribution in Fig. 4(b). The next group of $\mathcal{G}_j$'s captured some of the far-removed outlier points. Finally, the red, tan and yellow colored $\mathcal{G}_j$'s correspond to the blue, green and amber distributions, respectively, in Fig. 4(b). To see image details, please zoom in.

ident that a number of previously undetected activity distributions were captured. Reordering the spatial similarities of the NerDI plot in Fig. 6(d), to produce Fig. 7, and projecting the connected neurons down to $\mathbb{R}^3$ we found that all of the highlighted distributions outlined in Fig. 4(b) were successfully learned. However, some of the activity trends, such as the colored blocks with only a small number of neuronal references, are products of over-clustering. While detecting these outlying feature points is paramount, we are currently investigating ways rely on membership values to flag isolated, aberrant days; by doing so, a number of $\vec{w}_k$'s that only model singletons could be freed and used in learning more complex distributions.

Case Study - Participant II

In contrast to trends found in the previous study, those in Fig. 8(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. Comparing 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 expected, the levels slowly died down. After the surgery, there is also a surge in bradycardiac (1-30 BPM) pulse fir-

(a) Current feature set, $\mathbf{X}_c \subset \mathbb{R}^{32} \downarrow \mathbb{R}^3$ (b) VAT image of the blue distribution in (a)
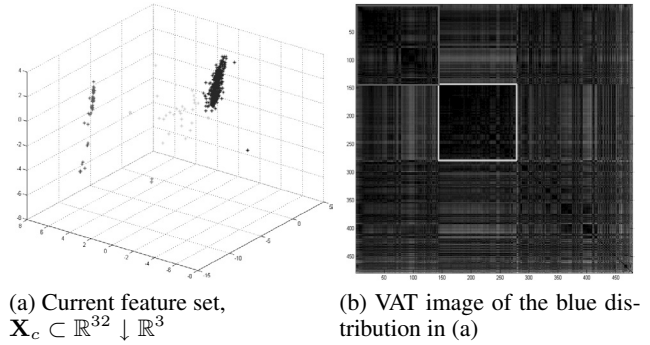
Figure 9: Plots of feature data for Participant II from 11/11/2005 to 05/18/2008. In (b), the two dense regions in the blue cluster are highlighted. Beyond the vectors in the red and green boxes, the remaining data points have a minimal cluster structure. This may imply that the resident's activities change over time. To see image details, please zoom in.

ings, in Fig. 8(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. 8(d), a decaying trend is also apparent while the decrease of eupnic and apneic sensor firings coincides with the drop in low pulse firings.

Turning now to the feature set in Fig. 9(a), 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. 4(b), the one in Fig. 9(a) corresponds to 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 associ-
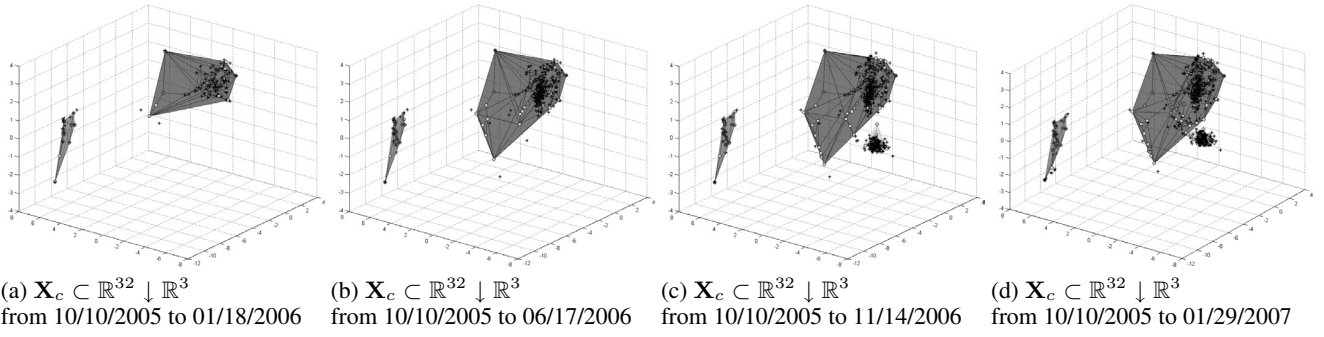
(a) $\mathbf{X}_c \subset \mathbb{R}^{32} \downarrow \mathbb{R}^3$ from 10/10/2005 to 01/18/2006

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

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

(d) $\mathbf{X}_c \subset \mathbb{R}^{32} \downarrow \mathbb{R}^3$ from 10/10/2005 to 01/29/2007

Figure 5: Temporal plots of the GNGC clustering results, for Participant I, with 40, 60, and 80 neuronal references, respectively. As more data is added to $\mathbf{X}_c$, such as in (b), (c) and (d), the $\vec{w}_k$'s, shown using filled green spheres, adapt their location and neighborhood connectivity to better model the data. To see image details, please zoom in.

(a) NerDI of $\mathbf{X}_c \subset \mathbb{R}^{32}$ from 10/10/2005 to 01/18/2006

(b) NerDI of $\mathbf{X}_c \subset \mathbb{R}^{32}$ from 10/10/2005 to 06/17/2006

(c) NerDI of $\mathbf{X}_c \subset \mathbb{R}^{32}$ from 10/10/2005 to 11/14/2006

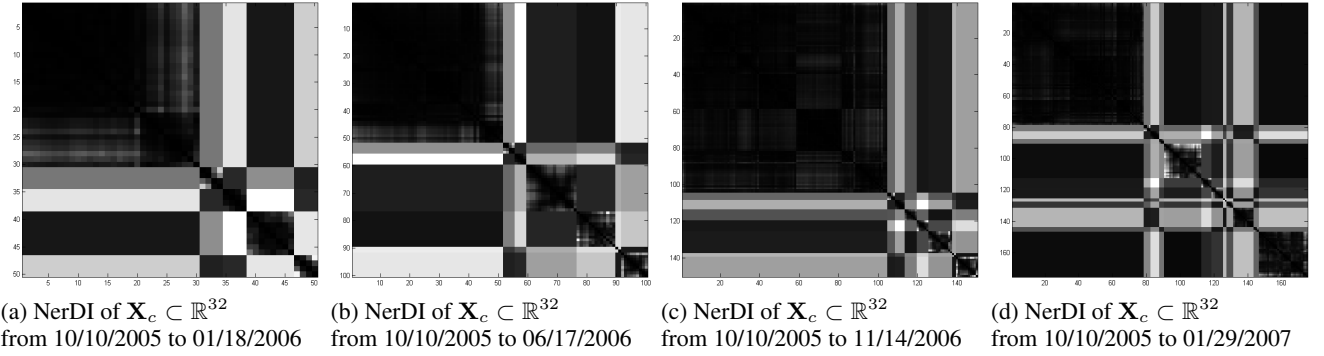(d) NerDI of $\mathbf{X}_c \subset \mathbb{R}^{32}$ from 10/10/2005 to 01/29/2007

Figure 6: Temporal plots of the GNGC clustering results, for Participant I, with 50, 100, 150, and 175 neuronal references, respectively. As with Fig. 5, when more data is appended, GNGC updates the global neuronal topology to account for the emerging distributions. In (a)-(d), pockets of related neurons, predominantly in the red colored block, indicate that there might be multiple activity trends in the baseline cluster from Fig. 4(b). To see image details, please zoom in.

(a) Plot of motion data

(b) Plot of bed restlessness data

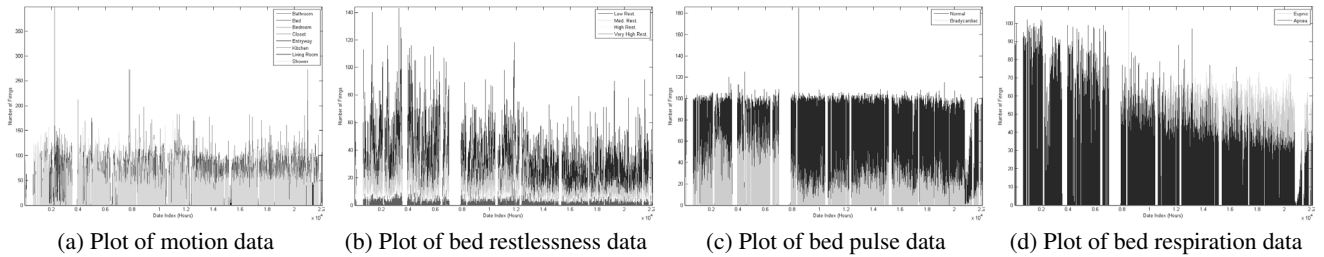(c) Plot of bed pulse data

(d) Plot of bed respiration data

Figure 8: Plots of the aggregated, raw sensor data, in hourly units, for Participant II from 11/11/2005 to 05/18/2008. 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. To see image details, please zoom in.
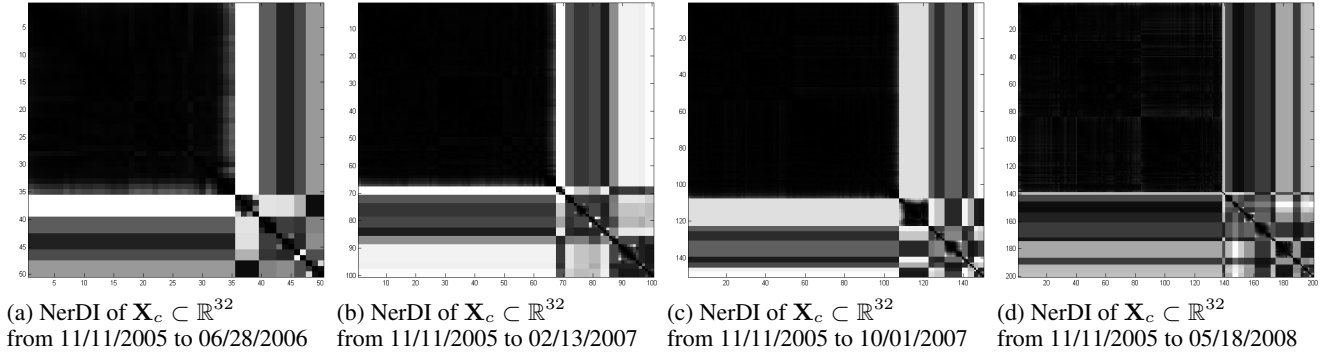
(a) NerDI of $\mathbf{X}_c \subset \mathbb{R}^{32}$
from 11/11/2005 to 06/28/2006

(b) NerDI of $\mathbf{X}_c \subset \mathbb{R}^{32}$
from 11/11/2005 to 02/13/2007

(c) NerDI of $\mathbf{X}_c \subset \mathbb{R}^{32}$
from 11/11/2005 to 10/01/2007

(d) NerDI of $\mathbf{X}_c \subset \mathbb{R}^{32}$
from 11/11/2005 to 05/18/2008

Figure 10: Temporal plots of the GNGC clustering results, for Participant II, with 50, 100, 150, and 200 neuronal references, respectively. As with Fig. 6, groups of related neurons, in each of the colored blocks in (a)-(d), highlight different activity trends within each cluster. To see image details, please zoom in.

ated 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. 9(b). 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 the potential loss in cluster structure when using PCA, we presented the non-projected dataset, $\mathbf{X}_c \subset \mathbb{R}^{32}$ to GNGC, one datum at a time. After the entire dataset had been presented, and the neuronal references stabilized, the algorithm reported that 13 clusters, shown in Fig. 10(d), materialized. Though we originally surmised that only four major distributions existed in the data, a number of isolated points drove up the cluster count. As with the previous case study, the large number of non-connected graphs was attributed to over-clustering. Somewhat disconcerting, however, was the failure to segment the blue distribution, in Fig. 9(a), into two separate clusters. While the two dark blocks, in the red dissimilarity matrix of Fig. 10(d), do capture the shift in both pulse and respiration firings, additional bed-related features or temporal information could help in separating these distributions. Techniques employed by Havens *et al*. (2008) or Sledge *et al*. (2008) would also prove useful for segmenting regions of dissimilar neurons.

Case Study - Participant III

Like the last two case studies, the data plots outlined in Fig. 11, show a number of different activity patterns. Considering the bed restlessness data, in Fig. 11(a), we see that there are two prominent trends: a multimodal distribution of firings over time and an increase in firings near the end of the currently logged data. A similar multimodal distribution is also present in the pulse and motion data, along with a period of a low number of firings in all four data plots. As well, in Fig. 11(b), there appears to be a lack of kitchen sensor firings for a period of around 5 months. While this blackout of firings may have been due to a faulty motion sensor, there are now safeguards in place to minimize data collection downtime.

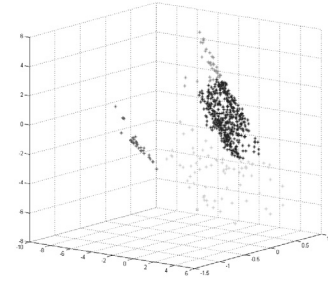Despite the visually perceptible trends in the raw sensor

Figure 12: Plots of the first three principle components of the feature data for Participant III from 08/03/2006 to 05/18/2008. To see image details, please zoom in.

data, a majority of the activity patterns were found by sifting through the PCA-reduced features and activity density plots, in Fig. 13. In both, a common theme was the large amount of time that the resident spent in bed, per day. This trend is not only present in the baseline (blue) cluster, in Fig. 12, but also the high restlessness (green) distribution. The few days that did not follow suit typically had high levels of activity, such as the amber region in Fig. 12, or were days when there was a low number of sensor firings, which clumped together to form the red distribution in Fig. 12.

Despite the extremes in activity levels, the features were introduced, a day at a time, to the GNGC algorithm. After the neuronal references stabilized, in each of the NerDI plots shown in Fig. 14, we projected the $\vec{w}_k$'s down to $\mathbb{R}^3$ and determined that the algorithm was successful in locating both the predominant distributions and the outlier points. Interestingly, some of the over-clustering issues, that plagued results in the earlier case studies, were automatically corrected during run-time.

Current & Future Research

With the ability to discover forming distributions, GNGC is an integral part of the adaptive, intelligent software-sensor system currently under construction. To aid in assessing eudemonia from the exposed trends, we plan on coalescing the cluster partitions with a fuzzy classifier. Unlike previous research in elder environment monitoring (Barnes *et al*. 1998;

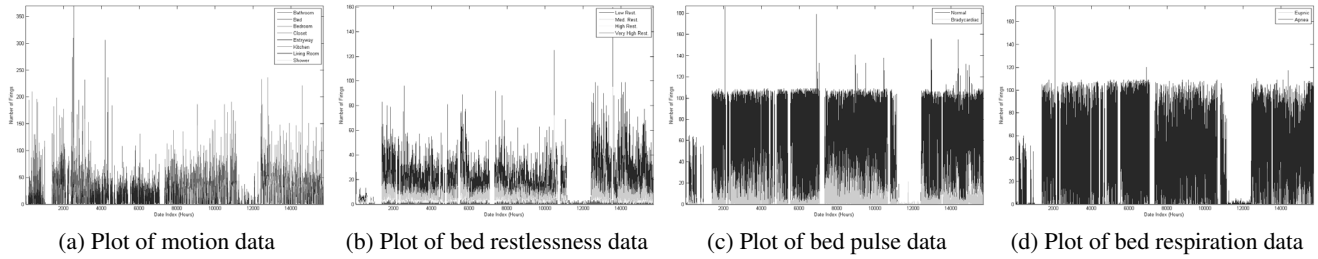| (a) Plot of motion data | (b) Plot of bed restlessness data | (c) Plot of bed pulse data | (d) Plot of bed respiration data |

Figure 11: Plots of the aggregated, raw sensor data, in hourly units, for Participant III from 08/03/2006 to 05/18/2008. 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. To see image details, please zoom in.

Brown *et al.* 2006; Ogawa *et al.* 2002), 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.

Although the features we have listed are useful for unearthing a number of activity distributions, additional ones may help differentiate between trends. Given that the motion attributes dominate the current feature set, in terms of the total number of motion characteristics, both bed respiration and pulse features are prime candidates for inclusion. In addition, by incorporating high-level information, like the number of visitors in a given period and time spent out of the apartment, the system may, one day, be able to detect the early onset of maladies such as depression, Alzheimer's disease, and so forth. Participant specific features may also play a role in future iterations of the system.

Beyond crafting the hybrid software system, we are also exploring methods such as multi-universe clustering (Wiswedel and Berthold 2007) for analyzing trends across multiple participants. Provided that there is a correlation, this infor-mation could help establish a baseline for incoming participants and predict future events for both new and existing residents.

While still in its infancy, growing neural gas clustering has been largely successful in locating emerging manifolds in both synthetic and real-world data. But, despite the temporal clustering designation, GNGC does not currently make use of any time-related information. To address this issue, we are currently experimenting with ways to fuse this additional quantity in both the algorithm and the visualizations (Sledge, Keller and Havens 2009). As well, we are presently developing other temporal clustering and classification methods that can natively handle temporal labels (Sledge and Keller 2009).
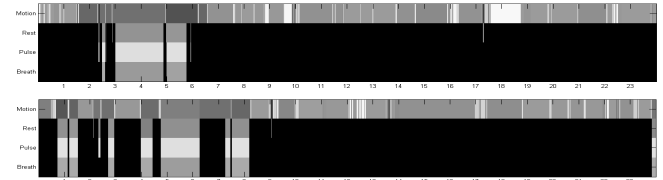
Acknowledgement
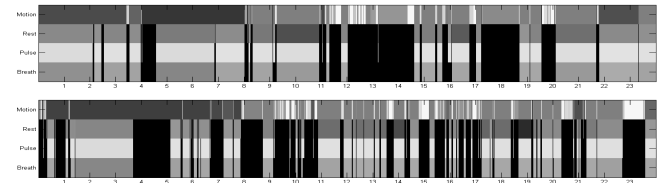
(a) Plots of the resident's "baseline" from the blue cluster. From these images, it's clear that the resident spends a majority of time in bed.
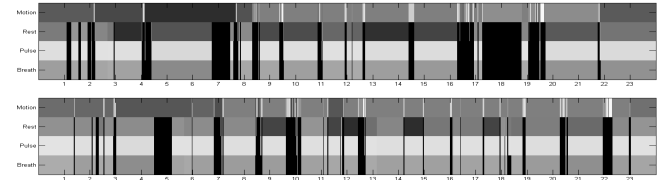
(b) Plots showing lethargy and high restlessness from the green cluster. As with (a), the participant spends most of the day laying in bed.

(c) Plots highlighting constant vivacity from the amber distribution. These plots may indicate that the resident is awake for a majority of the day. Alternatively, it could imply that the elder is resting on a couch or chair, given the time spent in the living room (shown using magenta).

Figure 13: Activity density images of various patterns for Participant III. To see image details, please zoom in.

References

Alwan, M., *et al.*, 2003. In-Home Monitoring System and Objective ADL Assessment: Validation Study. In Proc., ICADI.

Anderson, D., *et al.*, 2007. Linguistic Summarization of Activities from Video for Fall Detection Using Voxel Person and Fuzzy Logic. Computer Vision and Image Understanding. Submitted.

Barnes, N., Edwards, N., Rose, D., Garner, P., 1998. Lifestyle Monitoring: Technology for Supporting Independence. Computing and Control Engineering Journal, vol. 9, pp. 163-174.

(a) NerDI of $\mathbf{X}_c \subset \mathbb{R}^{32}$ from 11/11/2005 to 06/28/2006

(b) NerDI of $\mathbf{X}_c \subset \mathbb{R}^{32}$ from 11/11/2005 to 02/13/2007

(c) NerDI of $\mathbf{X}_c \subset \mathbb{R}^{32}$ from 11/11/2005 to 10/01/2007

(d) NerDI of $\mathbf{X}_c \subset \mathbb{R}^{32}$ from 11/11/2005 to 05/18/2008
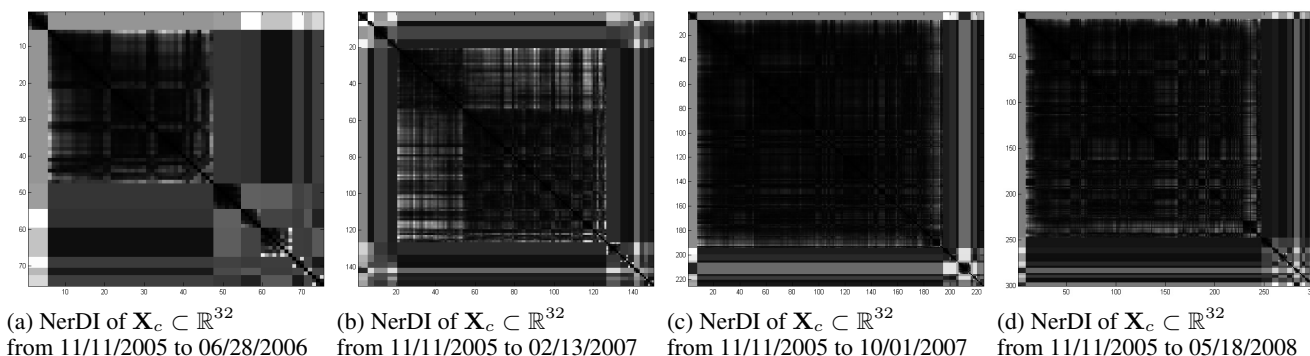
Figure 14: Temporal plots of the GNGC clustering results, for Participant III, with 75, 150, 225, and 300 neuronal references, respectively. To see image details, please zoom in.

Bezdek, J., and Hathaway, R., 2002. VAT: A Tool for Visual Assessment of (Cluster) Tendency. In IEEE Proc., IJCNN, pp. 2225-2230.

Bezdek, J., Keller, J., Krishnapuram, R., Pal, N., 1999, Fuzzy Models and Algorithms for Pattern Recognition, New York: Springer.

Brown, S., Majeed, B., Clarke, N., Lee, B., 2006. Developing a Well-Being Monitoring System Modeling and Data Analysis Techniques. Promoting Independence for Older Persons with Disabilities: IOS Press.

Carpenter, G., and Grossberg, S., 2003. Adaptive Resonance Theory. The Handbook of Brain Theory and Neural Networks, Massachusetts: MIT Press.

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

Havens, T., *et al.*, 2008. Clustering in Ordered Dissimilarity Data. In Pattern Recognition. Submitted.

Khachiyan, L., and Todd, M., 1993. On the Complexity of Approximating the Maximal Inscribed Ellipsoid for a Polytope. In Math Programming, vol. 61, pp. 137-159.

Keller, J., and Sledge, I., 2007. A Cluster By Any Other Name. In IEEE Proc., NAFIPS, pp. 427-432.

Knickman, J., and Snell, E., 2002. The 2030 Problem: Caring for Aging Baby Boomers. In Heath Services Research, vol. 7, pp. 849-883.

Luke, R., Anderson, D., Keller, J., and Skubic, M., 2007. Moving Object Segmentation from Video using Fused Color and Texture Features. IEEE Trans. Pattern Analysis and Machine Intel. Submitted.

Mack, D., *et al.*, 2006. A Passive and Portable System for Monitoring Heart Rate and Detecting Sleep Apnea and Arousals: Preliminary Validation. In IEEE Proc., D2H2, pp. 51-54.

Martinez, T., and Schulten, K., 1991. A Neural Gas Learns Topologies. Artificial Neural Networks, pp. 397-402.

Marek, K., Rantz, M., and Porter, R., 2004. Senior Care: Making a Difference in Long-Term Care of Older Adults. In Jo. of Nursing Education, vol. 43, pp. 81-83.

Ogawa, M., *et al.*, 2002. Long-Term Remote Behavioral Monitoring of the Elderly using Sensors Installed in Domestic Houses. In IEEE Proc., EMBS/BMES, pp. 1853-1854.

Rantz, M., *et al.*, 2005. TigerPlace: A New Future for Older Adults. Journal of Nursing Care Quality, vol. 20, pp. 1-4.

Sledge, I., *et al.*, 2008. Partitioning Ordered Dissimilarity Data. IEEE Trans. Knowledge and Data Engineering. Submitted.

Sledge, I., and Keller, J., 2008. Growing Neural Gas for Temporal Clustering. In IEEE Proc., ICPR.

Sledge, I., Keller, J., and Alexander, G., 2008. Emergent Trend Detection in Diurnal Activity. In IEEE Proc., EMBS/BMES.

Sledge, I., Keller, J., and Havens, T., 2009. Temporal Neuronal Clustering. IEEE Trans., Neural Networks. In preparation.

Sledge, I., and Keller, J., 2009. Swarming Agents for Temporal Exploratory Data Analysis. In IEEE Proc., SIS. Submitted.

Wang, S., and Skubic, M., 2008. Density Map Visualization from Motion Sensors for Monitoring Activity Level. In IET Proc., IE.

Wiswedel, B., and Berthold, M., 2007. Fuzzy Clustering in Parallel Universes. In Jo. of Approx. Reasoning, vol. 45, pp. 439-454.