Long-term behavioural change detection through pervasive sensing

Kemp, J. , Gaura, E. , Rednic, R. and Brusey, J.

Long-term Behavioural Change Detection Through Pervasive Sensing

John Kemp, Elena Gaura, Ramona Rednic, James Brusey
Coventry University, Coventry, UK
{aa9384, e.gaura, aa6418, j.brusey}@coventry.ac.uk

***Abstract*—The paper proposes an information generation and summarisation algorithm to detect behavioural change in applications such as long-term monitoring of vulnerable people. The algorithm learns the monitored subject's behaviour autonomously post-deployment and provides time-suppressed summaries of the activity types engaged with by the subject over the course of their day to day life. It transmits updates to external observers only when the summary changes by more than a defined threshold. This technique substantially reduces the number of transmission required by a wearable monitoring system, both through summarisation of the raw data into useful information and by preventing transmission of duplicated or predictable data and information. Based on evaluation using simulated activity data, the proposed algorithm results in an average of one transmission per month following an initial convergence period (reaching less than 1 transmission per day after only three days) and detects a change in behaviour after an average of 1.1 days.**

***Index Terms*—body sensor networks, pervasive sensing, behavioural change detection**

I. Introduction

The ability to monitor a subject's daily activities is a key requirement for a series of emerging societal trends, such as assistive living scenarios and interventions driving obesity-related behavioural changes in adults and young children. Given recent advances in body sensor networks (hardware and associated signal processing), many common daily tasks performed by subjects within and around the home can be monitored with a single wearable device (such as a mobile phone or a dedicated sensing node). Challenges remain however in i) effectively deploying such systems in real-life scenarios, ii) ensuring a system battery life commensurate with the requirements of the application and iii) adequately presenting "information" rather than "data" to the stakeholders (both the observed subjects as well as the medical stakeholders). The paper here relates to the last two challenges and demonstrates that wearable activity monitoring systems can be effectively designed to be long lived if the informational outputs are clearly defined. In turn, by designing systems to deliver information rather than data, providing visualization to enable decision-making becomes a straightforward design task.

The authors' core insight is the relationship between data and information in activity monitoring systems: *Wearable devices for activity monitoring usually contain high data rate sensors such as accelerometers and gyroscopes (with sampling rates from tens to hundreds of Hz per sensor). However, the resulting extracted activity information is generally much smaller than the data required to generate it.*

One of the biggest challenges in monitoring system design is the common need for wireless communication, with radio transcievers often being the largest consumers of available power on a node. The Shimmer devices, for example, typically consume 0.1 mA for the CPU, 1 mA during ADC conversions, and up to 20 mA for the 802.15.4 radio [1]. Processing the gathered data to generate the information that the end-user requires may allow large savings, particularly in applications such as activity monitoring where high-rate source data is used to generate smaller quantities of information. Particularly, this makes it possible to use the radio's sleep mode for longer periods; in the example of the Shimmer devices this reduces radio power consumption to less than 0.03 mA. The implication therefore is that generating this information closer (in terms of communication hops) to the sensing devices themselves will allow a significant reduction in transmissions and power usage.

This paper proposes the use of an information generation and summarisation algorithm named Bare Necessities (BN) [2] to allow for detection of behavioural change in applications such as long-term monitoring of vulnerable people. This application is used as an example to demonstrate the utility of the algorithm in transmission reduction and automated behavioural change detection, but is not the sole focus of the work. The algorithm supports the development of long-lived "deploy and forget" systems targetted at detecting behavioural change by transmitting behavioural information updates only when required to show that a behavioural change has occured. Such summarisation then allows for changes in behaviour over a period of time to be detected. Furthermore, the level of supervision required by human observers is reduced as redundant information is not transmitted. These advantages are particularly beneficial where multiple subjects are monitored.

The advance described by this paper, therefore, is a method of activity summarisation that allows a significant reduction in wireless communication from an on-body activity classification system. This method further reduces transmissions by eliminating redundant information, and as a side-effect autonomously highlights changes in behaviour.

This paper is structured as follows: Section II discusses existing research related to the work here, particularly in the area of activity monitoring. Section III describes the BN algorithm, while Section IV discusses its application to behavioural change detection. Section V gives the application example used for evaluation of the BN-based algorithm presented in this paper and describes the simulator developed to generate activity information for evaluation. Section VI presents the results of evaluation using the simulator. Finally, Section VII concludes on the work presented.

II. Related work

Monitoring of day-to-day activities is a common target application in the posture and activity classification domain and has prompted the development of a variety of sensing systems, ranging from wearable systems (based on inertial and other sensors) to ambient systems (which use video cameras, motion sensors, or open/close sensors on doors and cupboards, or example). A brief survey of the state of the art is given here to demonstrate the types of activities that have been commonly targetted for classification using these systems. Generally, current work based around one of three main activity groups: activities within the home, leisure/fitness activities outside of the home, and office activities.

Bao and Intille [3], for example, developed a classifier targeting twenty different home and leisure activities including walking, sitting while folding laundry, bicycling, and vacuuming. Huynh *et al.* [4] also focused on home activities, drawing a difference between low-level activities (such as walking, sitting, standing, eating, and washing dishes—usually lasting up to several minutes) and high-level activities (such as cleaning the house—composed of multiple low-level activities and lasting as long as a few hours). Pansiot *et al*. [5] present a system integrating an ear-worn activity recognition sensor (e-AR, which senses tilt and movement frequency spectrum) and ambient blob sensors that process a video signal to identify blobs or silhouettes and their motion based on optical flow. The system is capable of differentiating between sitting, sitting (sofa), standing, standing (head tilted), reading, eating, lounging, walking, and lying down. Ermes [6] targeted a mixture of indoor activities (such as lying, working on a computer, and standing reading a paper) and outdoor activities (such as playing football, running, rowing, and cycling). Laerhoven *et al.* [7] expanded on classification of daily activities by introducing a rhythm model that captures the user's normal daily pattern of behaviour. Activities included having breakfast, relaxing in the sauna, and watching TV. The rhythm model allows the system to perform classification of otherwise ambiguous sensor data. Gyllensten and Bonomi [8] used a single Tracmor device mounted on the subject's waist to classify activity, along with an IDEEA device that was considered to provide a high enough accuracy that it could be used to determine the "ground truth" for activity.

Further to monitoring of activities, this paper presents a solution to detecting changes in activity routines. Change in daily activity and routine is an important area of research in relation to the elderly, ill, disabled, or those with other impairments to their day-to-day activities. Campbell *et al.* [9], for example, studied the activity patterns of residents at a retirement community during a quarantine. PIR sensors were used to determine the number of transitions from room to room that occur during a day. They demonstrate that there was significant difference in room transitions between those that remained healthy and those that fell ill. They show that, while room transitions increased during the quarantine, the time-adjusted count remained similar to pre-quarantine levels. Furthermore, they show that the time-adjusted count was significantly different between those that remained healthy and those that reported feeling ill. Shin *et al.* [10] used IR sensors in each room of a house to identify abnormal activity patterns using the Support Vector Data Description technique. This resulted in a positive predictive value of 90.5% when tested on real data gathered in the home. Unlike the method described in this paper, the technique used by Shin *et al.* requires training prior to deployment, whereas the algorithm here forms a model of normal activity autonomously after deployment. Madan *et al*. [11] used data gathered using mobile phones to investigate the effect of illness on behavioural features such as communication and movement patterns, as well as the ability to use these features to predict illness.

Common to most works mentioned is that the scope of monitoring is to leverage condensed knowledge about the subject's behaviour, changes in behaviour, habits and routines. While the works propose that these key domain questions are answered through a series of post-processing, non-automated, human-in-the-loop steps, the authors here promote a design approach for the monitoring systems which stems from and revolves around the domain questions. This approach to design ensures highly optimized systems which deliver information and knowledge to answer set questions, enable long-lived monitoring, and promote simple visualizations of the systems' output.

For the purpose of the work presented in this paper, the monitoring scenario and associated high level question are as follows: a wearable activity monitoring system is able to classify, in real-time, daily activities within and around the home; from an external (medical) observer viewpoint, the system should report on changes in the subject's behaviour indicative of a decrease in energy intensive activities (such as for example gardening) and an increase in sedentary activities (such as laying in bed). With this in view, the next section describes a summarization algorithm which, after a convergence period (within which the system autonomously learns the normal subject behaviour) issues transmissions only when a behaviour change takes place. The algorithm is named *Bare Necessities*, reflecting its nature: transmitting only information relevant to the high level question.

III. The Bare Necessities algorithm

This section provides an overview of the concept and operation of BN. For a more detailed discussion see Gaura *et al.* [2]. The BN algorithm provides a means by which transmission by a monitoring system such as a Wireless Sensor Network (WSN) can be dramatically reduced while still providing the end-user with the information they require. The reasoning behind BN is that once a monitoring system such as a WSN

Algorithm 1 Online BN algorithm for estimating exposure band distribution B.

1) (*update band count*)
$B^{-}(i) \leftarrow \gamma B^{-}(i) + b(i,k)$,
for each measurand and for all i.
The predicate function $b(i,k)$ gives 1 if the current reading k is in band i and zero otherwise. The update decays the current count estimate by decay constant γ and then increments the currently active band. The decay half-life is $t_{1/2} = T \ln 2/(1-\gamma)$ where T is the sensing period.
2) (*update distribution*)
$B(i) \leftarrow B^{-}(i) / \sum_i B^{-}(i)$,
for each measurand and for all i.
This converts the counts to a distribution that sums to 1.
3) (*event detect*)
if, for any i, $|B(i) - B'(i)| > \varepsilon$ then
   a) transmit B and
   b) update last transmitted state $B' \leftarrow B$

is validated and installed, it is no longer necessary to transmit the raw gathered data. Instead, the information required by the end-user can be generated locally (as close as possible to the point of sensing). Furthermore, the information need only be transmitted when it changes in a way that the end-user would find interesting or important (this operational principle is shared with the Spanish Inquisition Protocol [12]). BN therefore allows for summarisation of the relative time spent in given states (exposure bands) and transmission of updates only when these bands change by more than a defined threshold. Algorithm 1 summarises the BN algorithm.

Gaura *et al.* [2] evaluated BN for the case study of a social housing monitoring system which allows landlords to assess home comfort against energy use. Here the proportion of time spent in various temperature ranges, for example, was more important than the measured temperature at any given time instant. The use of BN was found to give a 6800 fold reduction in transmissions of the course of a year compared to continuous transmission of the gathered data.

IV. The use of BN in behavioural change detection

The application of BN to behavioural change detection is a natural evolution of its concept of exposure bands, wherein each band can be assigned to a particular activity or type of activity. Over several days or weeks, the algorithm will build up an overview of the typical activities, habits, and behaviours of the monitored subject. During these times the number of transmissions to provide band updates will gradually decrease as the ratios between activities stabilise or converge (this is demonstrated in Section VI). The number of transmissions beyond the initial convergence period will be minimal, providing a long wearable system lifetime and reducing the storage requirements if local storage of historical band information is required. If there is a significant change in the relative ratios of monitored activities after this point, an update transmission will be triggered based on the band ratios changing by more than the pre-defined threshold. In some applications (such as the behaviour change detection discussed here), the fact that a transmission occured at all is equally as important as the actual content of the transmission.

As BN is intended as a general-purpose transmission reduction technique, it does not require that the monitored states are of a specific type. For the class of applications here, for example, the states could be based on activities such as sitting and walking, higher-level activities such as cleaning and cooking, triggered events based on cupboard door and appliance sensors, or the room occupied at each time instant. This flexibility means that a particular application and data gathering system must be specified in order for meaningful evaluation to be performed. The next section describes the application considered here.

V. Application example and activity simulator

A. Application

Activity classifications for the subject being monitored are assumed to be provided, in the case study evaluated here, by a wearable single-node Body Sensor Network (BSN) system capable of classifying and reporting activity in real-time (such as the systems described by Khan *et al.* [13] or Gyllensten and Bonomi [8]). The purpose of activity monitoring in this application is to detect behavioural change in the monitored subject that may indicate they are unwell or have suffered a decrease in mobility. In the application here, the decrease in mobility is assumed to be represented by spending more time in bed and no longer leaving the house to perform activities such as gardening. Madan *et al.* [11] demonstrated a correlation between illness and time spent in the home for university students. It can be seen that a system such as the one proposed here would be very useful deployed alongside other (perhaps more traditional) assistive technologies.

B. A simple activity simulator

In order to evaluate the effectiveness of BN for behavioural change detection, a simple activity simulator was implemented. This simulator allows for a variety of activities and events (described for this application in Section V-C) and performs the following process for each simulated day:

1) Select a waking and sleeping time based on the application specification.
2) Reserve time periods for application-defined events (such as meals).
3) Select application-defined activities to occur between each event. For the purpose of the simple simulator here, it was assumed that a single activity would occur between each pair of events.
4) Generate appropriate postural classifications through the day based on the generated routine (at a rate of one classification per minute). Where an activity is specified as being made up of multiple possible postures, the ratio between the possible postures was randomly chosen without bias.

In addition to defining the possible events and activities, the application also defines the possible times at which they may

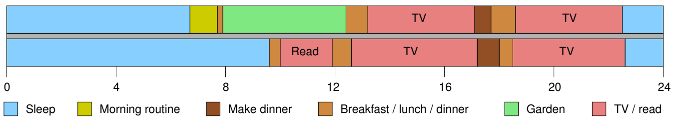

Figure 1. Example of simulated days. Top: Routine 1; Bottom: Routine 2.

Table I
DEFAULT PARAMETERS FOR ROUTINE GENERATION. (NOTE 1) RANDOM INTERVAL FOLLOWING WAKING USING A GAUSSIAN DISTRIBUTION WITH MEAN 60 MINS AND S.D. 30 MINS. (NOTE 2) IMMEDIATELY PRECEDES DINNER TIME.

	Start (Gaussian)		Duration	
Event	Mean	SD	Min	Max
Wake up	6:30am	15 mins	-	-
Breakfast	(Note 1)		20	45
Lunch	12:00pm	15 mins	20	45
Make dinner	(Note 2)		30	60
Dinner	6:00pm	30	30	60
Sleep	11:00pm	30 mins	-	-

occur. For example, a particular activity may have a start time with a Gaussian distribution around a given mean and standard deviation. While the above process results in a relatively simplistic view of activity, it serves to demonstrate the summarisation and behavioural change detection capabilities of BN.

For the purpose of simulating a change in behaviour, two routines specifications are defined, named R1 and R2. The simulator generates several weeks of R1, assumed to be the "normal" routine for the subject, followed by several weeks of R2. Ideally, this would cause BN to generate a behavioural update transmission soon after the change.

C. Application-based simulation parameters

For the purpose of the application here, the events considered were meals—breakfast, lunch, and dinner (with additional time reserved for making dinner)—and the activities were morning routine (washing, brushing teeth, etc), reading, watching TV, and working in the garden. The activity classifications during each event and activity were limited to lying (for sleeping), standing (for the morning routine and making dinner), walking (for gardening and making dinner), sitting (for eating, reading, and watching TV), and kneeling (for gardening).

The times for sleeping, waking, and meals within each day are determined as described in Table I. The available activities between breakfast and lunch and between lunch and dinner are reading, watching TV, and working in the garden. The available activities after dinner are reading and watching TV. Figure 1 shows an example of a single day generated using this method. The changes when switching to R2 are that the subject begins to wake several hours later than usual (the mean shifts to 9:30am), sleeps slightly earlier than usual (the mean shifts to 10:00pm), and no longer works in the garden—perhaps indicating that they are unwell or have suffered a decrease in mobility.

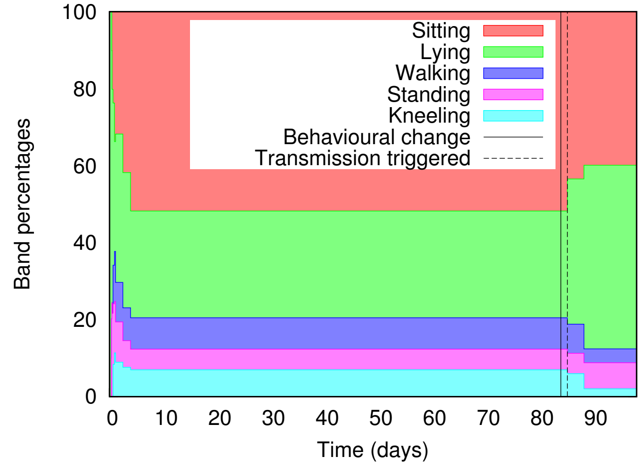

Figure 2. Example of band percentages during simulated activities over several weeks. A behavioural change occurs after 84 days. It can be seen that two transmissions are triggered during the "settling time" following the change.

VI. RESULTS FROM SIMULATION-BASED EVALUATION

Figure 2 shows an example of 98 days of activity simulated via the method described in Section V-B, with 12 weeks of R1 followed by 2 weeks of R2 (as described in Section V-B). BN was configured with a half-life of 2 days[1] and a change threshold of 10%[2]. It can be seen that the change in behaviour between the two routines causes a band update transmission to occur on the day following the change, and that day-to-day variations otherwise cause no tranmissions to occur.

The ideal behaviour for the system would be to minimise transmissions and only transmit following a behavioural change. Thus the metric used for evaluation in this work (based on the output of the simulator) is in two parts:

- The number of transmissions prior to the change should be as close to zero as possible.
- The change-triggered transmission should occur as soon after the change as possible.

To determine the optimal parameters for BN in the application here, the half-life and threshold parameters were varied (1–7 days and 5–20% respectively) and the metrics described previously were logged for each combination. Figures 3 and 4 summarise the results based on 100 iterations of each combination of parameters. Based on the results shown, it appears that the optimal parameters are a halflife of 2 days and a threshold of 10%. These gave an average of 14.2 tranmissions prior to the change (a reduction of 8518× compared to continuous transmission) and detected the change an average of 1.1 days after it occured. For thresholds of 10% to 20%, increasing the halflife beyond 2 days does not provide any significant improvement in terms of tranmissions prior to the change, but does increase the time taken to detect the change. The same

[1]Halflife is the time after which a given sample will be weighted with half the importance of a new sample.

[2]A transmission will be triggered if the size of any band changes by more than 10%.

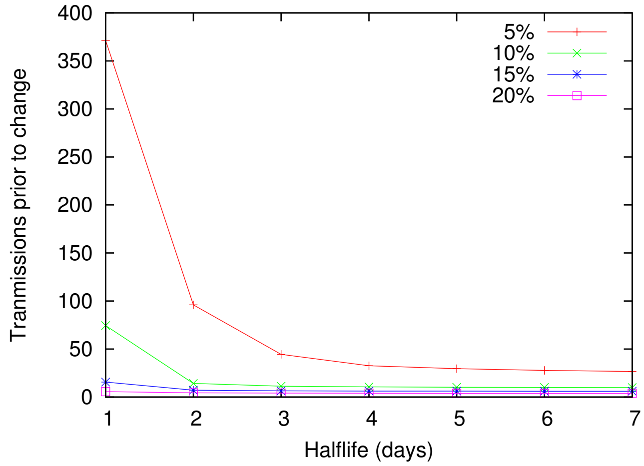

Figure 3. Impact of BN parameters on number of transmissions made prior to behavioural change occurring (84 day period).

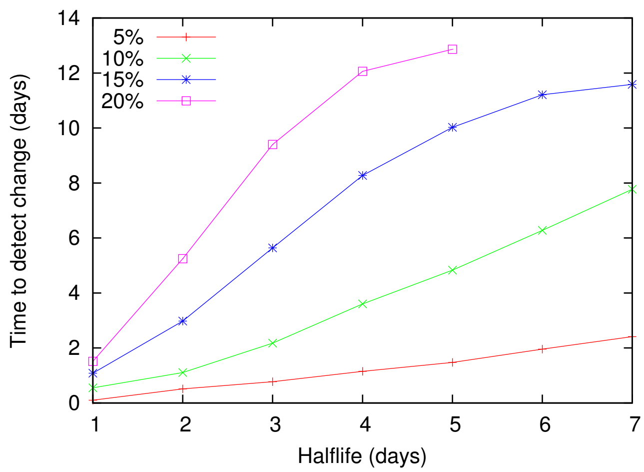

Figure 4. Impact of BN parameters on time taken to detect behavioural change. Note that the combination of a halflife of either 6 days or 7 days and a threshold of 20% failed to detect the change within the two week period allowed.

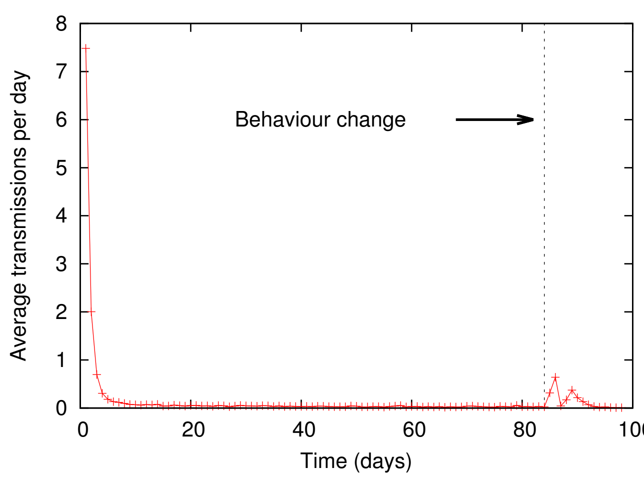

Figure 5. Average number of transmissions per day across 500 simulation iterations..

behaviour can be seen when increasing the threshold above 10%—no significant impact on tranmissions, but an increase in time taken to detect the change. Using the discovered parameters of halflife=2 days and threshold=10%, Figure 5 shows the average number of tranmissions per day across 500 iterations of the simulation. It can clearly be seen that there is an initial series of transmissions as the behaviour pattern is discovered, dropping off until the behaviour change occurs. The algorithm causes an average of 0.03 transmissions per day (a 48000× reduction in transmissions compared to continuous monitoring, roughly equivalent to one transmission per month) following an initial settling period. The transmission rate drops below 1 transmission per day after only three days.

VII. Conclusions

This paper demonstrated the use of the Bare Neccessities (BN) algorithm to reduce transmission requirements compared to continuous monitoring for long-term behavour monitoring. The need for human supervision of the incoming data is also reduced—only "interesting" data is reported. The example is given of monitoring an elderly person at home in order to detect behavioural change that might indicate that assistance is required. The application is based around a wearable device that can provide postural information based on, for example, acclerometer data. The algorithm, though, is generic to a variety of activity monitoring applications and data sources. The optimal BN parameters were found to be a halflife of 2 days and a change threshold of 10%, resulting in an average of 14.2 transmissions in the 84 days prior to a simulated behaviour change occuring and transmission of an update an average of 1.1 days after the change occured. The number of transmissions made is equivalent to approximately 1 transmission per month following the initial settling period during which the subject's routine is characterised (dropping below 1 tranmission per day within 3 days). Future work is aimed at a deeper understanding of current biomechanical and physiological behaviour change metrics as well as deployment of the proposed algorithm on a wearable platform and associated evaluation.

References

[1] Shimmer Research, "Fequently Asked Questions (What is the power consumption of the Shimmer?)." http://www.shimmer-research.com/links/faqsH1.

[2] E. I. Gaura, J. Brusey, and R. Wilkins, "Bare Necessities—Knowledge-driven WSN design," in *In Proceedings of IEEE Sensors*, 2011.

[3] L. Bao and S. S. Intille, "Activity recognition from user-annotated acceleration data," in *Proceedings of the Second International Conference on Pervasive Computing (PERVASIVE 2004)*, (Linz/Vienna, Austria), pp. 1–17, Berlin, Heidelberg: Springer-Verlag, April 2004.

[4] T. Huýnh, U. Blanke, and B. Schiele, "Scalable recognition of daily activities with wearable sensors," in *Proceedings of the 3rd international conference on Location-and context-awareness (LoCA'07)*, (Oberpfaffenhofen, Germany), pp. 50–67, Berlin, Heidelberg: Springer-Verlag, 20–21 September 2007.

[5] J. Pansiot, D. Stoyanov, D. McIlwraith, B. Lo, and G. Yang, "Ambient and wearable sensor fusion for activity recognition in healthcare monitoring systems," in *Proceedings of the 4th International Workshop on Wearable and Implantable Body Sensor Networks (BSN 2007)*, vol. 13, (RWTH Aachen University, Germany), pp. 208–212, Berlin, Heidelberg: Springer-Verlag, 26–28 March 2007. ISSN: 1680-0737.

[6] M. Ermes, J. Pärkkä, J. Mäntyjärv, and I. Korhonen, "Detection of daily activities and sports with wearable sensors in controlled and uncontrolled conditions," *IEEE Transactions on Information Technology in Biomedicine*, vol. 12, pp. 20–26, 7 January 2008. ISSN: 1089-7771.

[7] K. V. Laerhoven, D. Kilian, and B. Schiele, "Using rhythm awareness in long-term activity recognition," in *Proceedings of the 12th International Symposium on Wearable Computers (ISWC2008)*, (Washington, DC, USA), pp. 63–66, Los Alamitos, CA: IEEE Computer Society Press, 28 September–1 October 2008.

[8] I. C. Gyllensten and A. G. Bonomi, "Identifying types of physical activity with a single accelerometer: evaluating laboratory-trained algorithms in daily life," *IEEE Transactions on Biomedical Engineering*, vol. 58, pp. 2656–2663, September 2011.

[9] I. H. Campbell, D. Austin, T. L. Hayes, M. Pavel, T. Riley, N. Mattek, and J. Kaye, "Measuring changes in activity patterns during a norovirus epidemic at a retirement community," in *33rd Annual International Conference of the IEEE EMBS*, (Boston, Massachusetts, USA), pp. 6793–6796, August 30–September 3 2011.

[10] J. H. Shin, B. Lee, and K. S. Park, "Detection of abnormal living patterns for elderly living alone using support vector data description," *IEEE Transactions on Information Technology in Biomedicine*, vol. 15, pp. 438–448, May 2011.

[11] A. Madan, M. Cebrian, D. Lazer, and A. Pentland, "Social sensing for epidemiological behavior change," in *12th ACM International Conference on Ubiquitous Computing (Ubicomp 2010)*, (Copenhagen, Denmark), 26–29 September 2010.

[12] D. Goldsmith and J. Brusey, "The Spanish Inquisition Protocol—model based transmission reduction for wireless sensor networks," in *Proceedings of IEEE Sensors 2010* (T. Kenny and G. Fedder, eds.), (Waikoloa, HI, USA), pp. 2043–2048, 1–4 November 2010.

[13] A. M. Khan, Y.-K. Lee, S. Y. Lee, and T.-S. Kim, "A triaxial accelerometer-based physical-activity recognition via augmented-signal features and a hierarchical recognizer," *IEEE Transactions on Information Technology in Biomedicine*, vol. 14, pp. 1166–1172, September 2010.