

Long-Term Activity Monitoring with a Wearable Sensor Node

Kristof Van Laerhoven, Hans-W. Gellersen
*Computing Department,
Lancaster University*
{kristof,hwg} @ comp.lancs.ac.uk

Yanni G. Malliaris
*Institute of Psychiatry,
King's College London*
y.malliaris @ iop.kcl.ac.uk

Abstract

This paper introduces an encapsulated sensor node that is devised to monitor and record motion patterns over long, quotidian periods of time with potential application in psychological studies. Its design fuses different sensing modalities to allow efficient capturing of tilt and acceleration stimuli, as well as embedded algorithms that abstract from the raw sensory data to indicative features. By combining tilt switches and accelerometers with customized processing techniques, it is argued that a power-efficient yet information-rich approach is reached for the observation and logging of human motion-based activity.

1. Introduction

Many believe that posture and motion patterns of the human body are key indicators for a person's emotional state and wellbeing [1]. Capturing the gestures and pose of dancers, actors, and more recently also patients, is receiving an increasingly amount of attention in the research community.

Common approaches that are currently applied in motion capture involve tags and external sensors (such as video cameras [2] or ultrasound receivers), or highly accurate body suits [3] measuring rotation and acceleration. It is, with today's technology, costly and intrusive to capture all aspects of body motion and pose throughout the day, in unpredictable environments and situations, with these traditional methods.

The sensor node that is described in this paper provides data that is accurate enough to tell what type of activities the wearer has been doing, but not as fine-grained and detailed as commercially available motion capture suits. The presented approach revolves around the abstraction of sensory data in the sensor module itself, and a dynamic switching of the most appropriate input modality.

2. Overview

The research that has led to this paper includes both hardware and algorithm decisions which are closely tied together: The algorithms do not just output pre-processed sensory data, but also control which sensors to activate, and at which speed the micro-controller should operate.

A diagram of the presented approach is given below in Figure 1 as an overview of the paper – justification and a more detailed explanation will follow from the next section onwards. The output of the sensor node varies with the mode it is in: the node remains silent in sleep mode, a three-dimensional orientation (pose) is given in tilt mode, and characteristic pose and peak descriptions are produced in acceleration mode.




Figure 1. Diagram of the three possible states of the sensor node (boxes), conditions for transition between them (edges), and output types (clouds).

The sensor node is foremost designed to provide a low-energy method for sensing motion and posture of the body it is attached to. As in many low-power designs, this need for preserving as much energy as possible has affected the choice of algorithms and individual hardware components (in particular microcontroller and sensors). In this case, it has also propagated to the data analysis, exploiting the fact that human motion patterns do not always contain gestures.

3. Application

Despite many advances in the genetics of psychiatric diseases [4], there is still no single laboratory test for the diagnosis of any psychiatric “disease”, unlike other areas of medicine. Diagnosis and therapy monitoring rely on self-report measures, semi-structured interviews, and careful observation of a patient’s behaviour. Given the “behavioural” nature of these disorders, the “status-quo” of psychiatric assessment is unlikely to change any time soon. This is where the promise of activity sensors presented in this paper lies.

Mood disorders make a fine example. For instance, a depressive episode is characterised by major changes in mood, cognition and behaviour (see DSM-IV-TR¹ for major diagnostic criteria). Although activity sensing may never be able to tap in a patient’s feelings or suicidal thoughts (but see Picard’s work for emerging solutions to this, [5]), it can easily monitor changes in psychomotor activity. Anergic and agitated states as well as disturbances in sleep, are not only highly distressing symptoms of mood disorders in the patients themselves, but they are also core symptoms that have high predictive value for further relapses in remitted patients [6].

Indeed, there is some emerging work in the application of activity sensors (Actigraphy) in psychiatry ([7], [8]). However, research so far has relied on commercial off-the-shelf single sensor accelerometers with limited recognition capabilities (such as the Actiwatch and the Motionlogger, [9]). The commercial cost of such devices (£400-£800) has thwarted further research development, and the application of such work remains restricted in the realms of well-funded research projects rather than real life clinical practice.

Although the low-cost and expandability of the presented activity sensors are major advantages for many applications in psychiatry, there are also important theoretical benefits. From a theoretical point of view, psychiatric disorders provide a good pool of extreme behavioural and emotional states (such as vegetative depressed and highly activated manic states, rigid rituals of OCD and Autistic patients and so forth). Such extreme states probably represent easier targets for recognition by quantitative models used in the current sensor network than “normal” behaviours and emotions found in healthy populations.

The challenge lies in the successful collaboration between the two fields (computer scientists / engineers

and clinical psychologists / psychiatrists). And also in the development of quantitative models, informed by activity sensors, which will successfully characterise and distinguish between different illness states and disorders.

4. Hardware Description

The biggest challenges in engineering such a body-worn activity sensor node are (1) keeping its power consumption down and (2) keeping its form-factor unobtrusive to wear comfortably. In order to be able to log activity data over long periods of time, the sensor node is built to switch between two sensing modes: a rough pose-measuring mode and a more accurate acceleration sensing mode. A third mode powers the node down for periods of no activity.

The first type of sensor present in the sensor node is the tilt switch: the combination of 9 tilt switches positioned at 45 degree angles from each other cover a very crude sense of tilt for the board in three planes (as illustrated in Figure 2 left). The easiest way to visualise how this works, is in one plane (Figure 2 right); the other two orthogonal planes work the same, only 9 instead of 12 switches are required since we have common tilt switches between the three planes.

The accelerometer (Analog Devices ADXL202JE) covers both acceleration and pose in two axes, but at a more fine-grained resolution than the cluster of tilt switches.

There are two important reasons for incorporating these sensors this way. Both are related to resources:

Energy. The tilt switches require less power: with heavy pull-up resistors, they draw only a few microAmperes, compared to a few hundred for the accelerometer.

Processing. The tilt switches’ output is binary, and thus easier and faster to process and store, especially when combining multiple switches. The microcontroller also does not need to run at a fast speed for reading the switches’ state, which results in an additional way to energy preservation.




Figure 2. Left: The sensor node (25x25mm) with accelerometer (a) and tilt switches (t). Right: the arrangement of the tilt switches allows for a coarse orientation or pose measurement.

¹<http://www.behavenet.com/capsules/disorders/dsm4TRclassification.htm>

The microcontroller (PIC16F628A) has specifically been chosen because of its capability to change its processing speed. For the tilt switches, it runs at a low 48kHz, while it needs to switch to 4 MHz for processing the accelerometer data. The microcontroller can also go into a low-power sleep mode if no change is detected for a while in the datastreams of both switches and accelerometers.

The speed at which any microcontroller runs affects its power consumption heavily (measured with a 3V coin cell battery attached): In sleep mode, it consumes about 200 nA. It can be woken up by the switches, in which case it changes to the 48kHz mode using approximately 20 uA. If motion patterns (peaks) are detected, the accelerometer is activated and the microcontroller switches to 4Mhz for reading and processing, which requires about 1 mA.

This sensor board is designed to piggyback on a wireless body sensor network node (BSN node, developed at Imperial College, with similar specs as the Telos Mote), to which it is interfaced via RS232. It contains capabilities for both logging the data locally (4Mbit, enough for logging 3 hours in worst case) and wirelessly distributing the data over a network.

5. Algorithm Description

The switches' states are read at a speed of around 200 Hertz; this information is not stored, but is instead processed into two values: one which describes the most often occurring set of states during the last 200 measurements, and another value that reflects how many switches have toggled during that period. The first is denoted the medoid, the second the cumulative hamming distance; the medoid occupies 9 bits, while the cumulative hamming distance is assigned to 3 bits.

The accelerometer's two signals are read at approximately 50 Hertz; out of the raw sensory data, the basic statistics (minimum, maximum, mean, and variance), and peak features (length, height, sign, and area) are calculated over a sliding window. These features for the two acceleration sensors are stored in 96 bits (6 bytes for basic statistics, 6 bytes for peak features of the last peak). Figure 3 illustrates how new values cause the features to be recomputed.

For the calculation of average and variance, it suffices to keep the sum of the past n elements (sum) and the sum of squares ($sumsq$) as the variance can be calculated by $(sumsq - (sum^2/n))$. Note that the complete array needs to be transversed only when minima and maxima need to be re-calculated (when the outgoing element used to be minimum or maximum). Otherwise only the first and last values of the first-in-first-out array need to be examined.

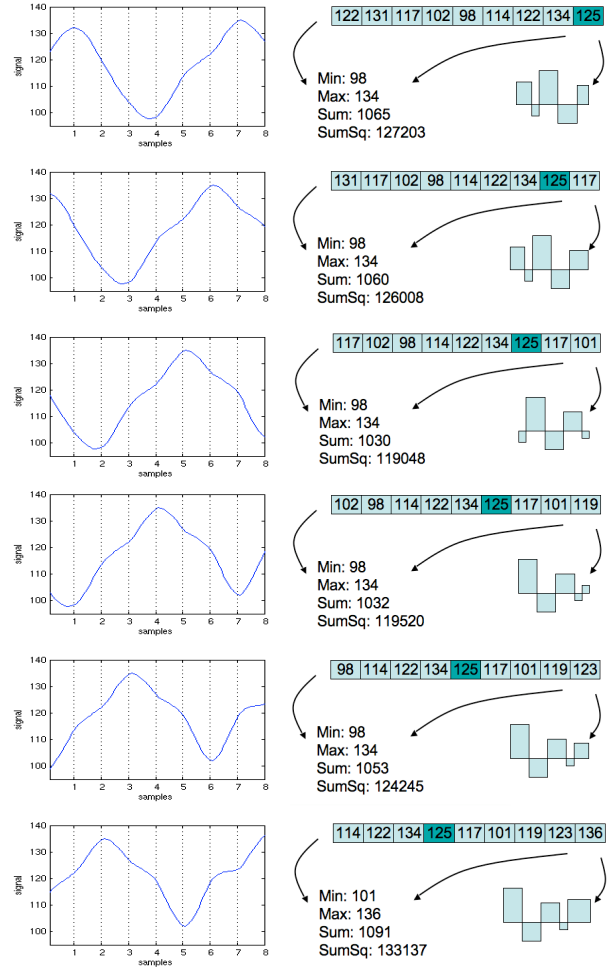


Figure 3. Step-by-step example in which a sliding window (first-in-first-out array) of past acceleration values is used to compute the necessary variables for basic statistics and peak features.

The output of the activity sensor node is by default a packet that contains both the tilt and acceleration information (15 bytes in total, including a 1-byte pre-ample). After a long period of inactivity (defined by the lack of significant peaks), the device will switch to output from just the tilt switches (3 bytes in total, including pre-ample). If this pose state remains the same for a long period, the device will finally go into a power-saving “sleep” mode.

The sleep mode can only be interrupted by activity on the tilt switches, in which case it will operate with tilt switch output only. The cumulative hamming distance measure is then used as a way to spot accelerometer-based activity: if a threshold is reached, the device goes into a mode that outputs the full tilt- and acceleration information. These modes thus allow a reduction in power consumption *and* data storage.

6. Long-Term Activity Monitoring

The motivation for this work lies in the recording of long-term activity patterns. In this aim, the challenge is to find an optimal position in the trade-off between conserving as much of the activity information as possible (but requiring much storage and energy) and doing the implementation with limited resources (but throwing away potentially useful data).

To illustrate how much information is retained, Figure 4a plots the data produced by an activity sensor node, attached to the ankle, while the wearer was walking and climbing stairs. The current implementation of the sensor board sends an output package (if in sensing mode) up to five times per second to the BSN node, as opposed to the common 30 to 50 Hertz found in work such as [10]. This bandwidth has been experimentally established to avoid missing significant peaks in the accelerometer signals.

Figure 4b shows reconstructed peak patterns from the features for both activities, showing promising differences that could be exploited by classification algorithms to distinguish between the two activities. Note that both activities would be difficult to separate if just the basic statistics were used as features.

When the user remains motionless for a while, no data is supposed to be recorded. The mechanical nature of the tilt switch, however, sometimes causes single tilt switches to oscillate, thereby preventing the device to go into sleep mode. In our preliminary experiments so far, this has occurred only on a few occasions.




Figure 4. a) Example data of the wearer while walking and climbing stairs, and b) graphical representations of typical peak patterns using the peak height and length features.

7. Conclusions and Future Work

This paper has focused on the design of the sensing component for the recognition of valuable features in acceleration and tilt measurements. Special care has been given on integration of both sensing and feature processing, so that as much activity-related data is retained, using a minimal amount of storage capacity.

With the sensing hardware completed, the next steps would be to implement storage and communication facilities on the BSN wireless node to start integration toward long term testing.

This work was partially sponsored by the UK's DTI-funded UbiCare centre (<http://www.ubicare.org>) and the MRC London/Newcastle eMonitoring Trial (<http://www.bipolarlab.com>). Details on hardware and software are available at: <http://www.comp.lancs.ac.uk/~kristof/research/notes/porcupine>

8. References

- [1] Schouwstra, S J. Hoogstraten J. (1995) Head position and spinal position as determinants of perceived emotional state. *Perceptual & Motor Skills*. 81(2):673-4
- [2] Song Y., Feng X. and Perona. P. (2000) Towards Detection of Human Motion. In *IEEE Proc. of Conf. Computer Vision and Pattern Recognition*, Vol. I, pp. 810-817.
- [3] Animazoo Gipsy, <http://www.animazoo.com>. 11/2005
- [4] Kelsoe, J. R. (2004) Genomics and the Human Genome Project: implications for psychiatry. *International Review of Psychiatry*. 16 (4): 294-300
- [5] Picard, R. W. (1997) *Affective Computing*, MIT Press, Cambridge.
- [6] Jackson, A. , Cavanagh, J., Scott, J. (2003) A systematic review of manic and depressive prodromes, *Journal of Affective Disorders*. 74 (3), 209-17
- [7] Teicher, M.H. (1995) Actigraphy and motion analysis: New tools for psychiatry. *Harvard Review of Psychiatry* 3, 18-35
- [8] Miller, A., Espie, C.A., Scott, J. (2004) The Sleep of Remitted Bipolar Outpatients: A Controlled Naturalistic Study Using Actigraphy. *Journal of Affective Disorders*. 80, 45-153
- [9] Benson, K., Friedman, L., Noda, A., Wicks, D., Wakabayashi, E., Yesavage, J. (2004) The Measurement of Sleep by Actigraphy: Direct Comparison of 2 Commercially Available Actigraphs in a Nonclinical. *Journal of Sleep and Sleep Disorders Research*. 27 (5), 986-989
- [10] Tapia, E. M., Marmasse, N., Intille, S. S., Larson, K. (2004) MITes: Wireless portable sensors for studying behavior. In *Proceedings of Extended Abstracts Ubicomp*.