D3.2
Early Fusion & Mining v1

Dementia Ambient Care: Multi-Sensing Monitoring for Intelligent Remote Management and Decision Support

Dem@Care - FP7-288199

Deliverable Information

Project Ref. No.		FP7-288199
Project Acronym		Dem@Care
Project Full Title		Dementia Ambient Care: Multi-Sensing Monitoring for Intelligence Remote Management and Decision Support
Dissemination level:		Public
Contractual date of delivery:		Month 24, October 31, 2013
Actual date of delivery:		Month <24>, < October 31, 2013>
Deliverable No.		D3.2
Deliverable Title		Early Fusion & Mining v1
Type:		Report
Approval Status:		< Approved >
Version:		9
Number of pages:		74
WP:		WP3
Task:		T3.3 Early Fusion and Mining of Health & Lifestyle Data
WP/Task responsible:		CS
Other contributors:		DCU, PENB
Authors (Partner)		Dafni Stampouli, Jennifer Renoux (CS), Eamonn Newman(DCU), Erik Schuijers (PENB)
Responsible Author	**Name**	Dafni Stampouli(CS)
	Email	Dafni.stampouli@cassidian.com
Internal Reviewer(s)		Francois Bremond (INRIA)
EC Project Officer		Gerard Cultot
Abstract (for dissemination)		This deliverable contains the first version of algorithms for early fusion and mining of multi-modal sensor data, in order to extract and recognise specific patterns and detect unusual and critical events. Initially, work is described to combine the Philips DTI-2 data and the Gear4 Renew Sleepclock in order to measure parameters of the PwD that relate to their daily activity patterns. Work on WIMU sensors is presented, describing algorithms created for processing accelerometer, gyroscope and magnetometer data. Finally the CEP engine is introduced. CEP is used to combine information and infer events. Detailed implementation and results are shown and future work is discussed.

Version Log

Version	Date	Change	Author
1	16-09-2013	Added first input on early fusion	Erik Schuijers
2	17-09-2013	Added CS input and additional information about early fusion	Dafni Stampouli
3	14-10-2013	Additional Editing	Dafni Stampouli
4	16-10-2013	Added (sub-)conclusions	Erik Schuijers
5	18-10-2013	Added CEP Implementation	Jennifer Renoux
6	18-10-2013	Final Editing	Dafni Stampouli
7	28-10-2013	Including internal reviewer comments	Erik Schuijers
8	30-10-2013	Including internal reviewer comments	Eamonn Newman
9	30-10-2013	Final Editing to release to PMB	Dafni Stampouli

Executive Summary

This deliverable contains the first version of algorithms for early fusion and mining of multi-modal sensor data, in order to extract and recognise specific patterns and detect unusual and critical events. Fusion of different wearable sensor data is presented in the different sections. Initially, work is described to combine the Philips DTI-2 data and the Gear4 Renew Sleepclock in order to measure parameters of the PwD that relate to their daily activity patterns. Detailed experiments are described along with work to overcome specific technical difficulties. Results are also presented and described in detail. Furthermore work on the WIMU sensors is presented, describing algorithms created for processing accelerometer, gyroscope and magnetometer data. The algorithms exploit low-level fusion of sensor data to detect different body postures and different modes of activity and inactivity. Finally the CEP engine is introduced. CEP is used to combine events detected by the two earlier sensors. For the first attempt of early fusion DTI-2 and Gear4 data collected over a large period of time is to be combined to assess the correlation between the daily exercise and quality of sleep. These are two of the main daily activities identified as critical by the clinical partners and the aim of this work is to examine whether early fusion can provide early indications to dependencies or problems between these two activities. Initial implementation is presented along with some early results. The report closes with some concluding remarks and description of future work for the WP3.

Abbreviations and Acronyms

CEP	Complex Event Processing
DTI	Discreet Tension Indicator
WIMU	Wireless Inertial Measurement Unit
PwD	People/Person with Dementia
PDF	Probability Density Function
IS	Inter-daily Stability
IV	Intra-daily Variability
AMP	Amplitude

Table of Contents

List of Figures

Figure 46 Rule Edition Interface

List of Tables

1 Introduction

The aim of this report is to describe the first version of algorithms and the developed technologies for early fusion of multi-modal data obtained by the wearable sensors in order to extract and recognise specific daily activity patterns and detect unusual events and correlations between events.

This document is separated in different sections; each describing a set of technologies. Initially in this report, in Chapter 2, the work to combine and fuse the Philips DTI-2 data and the Gear4 Renew Sleepclock data is described, in an attempt to augment the DTI-2 data with the Gear4 data. Detailed experiments are described along with work to overcome specific technical difficulties. Results are also presented and described in detail.

Furthermore work on the WIMU sensors is presented. Chapter 4 presents algorithms created for processing accelerometer, gyroscope and magnetometer data. The algorithms exploit low-level fusion of sensor data to detect different body postures and different modes of activity and inactivity. The results obtained and presented are well in accordance to the proposed needs within Dem@Care. Further work on this topic is also discussed in the chapter.

At Chapter 5, CEP processing is introduced. CEP is used to combine events detected by the two earlier sensors. For the first attempt of early fusion DTI-2 and Gear4 data collected over a large period of time is to be combined to assess the correlation between the daily exercise and quality of sleep. These are two of the main daily activities identified as critical by the clinical partners and the aim of this work is to examine whether early fusion can provide early indications to dependencies or problems between these two activities. Initial implementation is presented along with some early results.

The report closes with some concluding remarks and description of future work.

2 Early fusion

2.1 Modelling daily activity patterns

2.1.1 Introduction

In Deliverable D3.1 [2] the use of the wrist-worn Philips DTI-2 is proposed as a means to monitor physical activity (see left panel Figure 1). The Philips DTI-2 is a multimodal sensing device, containing 3-axis accelerometers, skin conductance, temperature and light sensors.

Figure 1 – Philips DTI-2 (left) and Gear4 Renew Sleepclock (right)

Bouten et al [5] describe how the physical activity can be approximated based on the accelerometer data (see also [2]). Obtaining a 24 hour view of the daily activity pattern of People with Dementia (PwD), i.e. day and night [1], would require a continuous wearing of the Philips DTI-2 device. Despite the fact that the wearing comfort of the Philips DTI-2 is relatively high, it is still undesirable to wear the device during night-time. Additionally, the limited battery life of the device would prohibit a full 24 hours of recording. As a means to monitor the night-time activity without making use of the Philips DTI-2, the Gear4 Renew Sleepclock was proposed (see right panel Figure 1). Together, using the Philips DTI-2 during daytime and using the Gear4 Renew Sleepclock at night, a complete 24 hour view of the physical activity of the PwD can be obtained. It is however not trivial how to interpret the data in a uniform manner. The Philips DTI-2 accelerometer data can be converted to so called activity counts at a fixed epoch length (e.g. 30 or 60 seconds) whereas the Gear4 Renew Sleepclock generates sleep/wake classifications on a 30 second epoch basis.

2.1.2 Fusing Philips DTI-2 and Gear4 Renew Sleepclock data

In order to fuse the Philips DTI-2 data and the Gear4 Renew Sleepclock data, a first set of lab data has been generated as follows. A Gear4 Renew Sleepclock device was put at the bedside of a lab employee. When this person went to bed, it would first start to wear the Philips DTI-2 and turn it on. As soon as this person was in bed, the Gear4 Renew Sleepclock was turned on as well. After the person woke up, the Gear4 Renew Sleepclock was turned off. For a number of days the Philips DTI-2 was being worn until the moment that the battery died, for other days the Philips DTI-2 was turned off at the moment that the Gear4 Renew Sleepclock was also turned off. In total 14 nights of simultaneous DTI-2 and Gear4 recordings were realised in this way (see Figure 2).

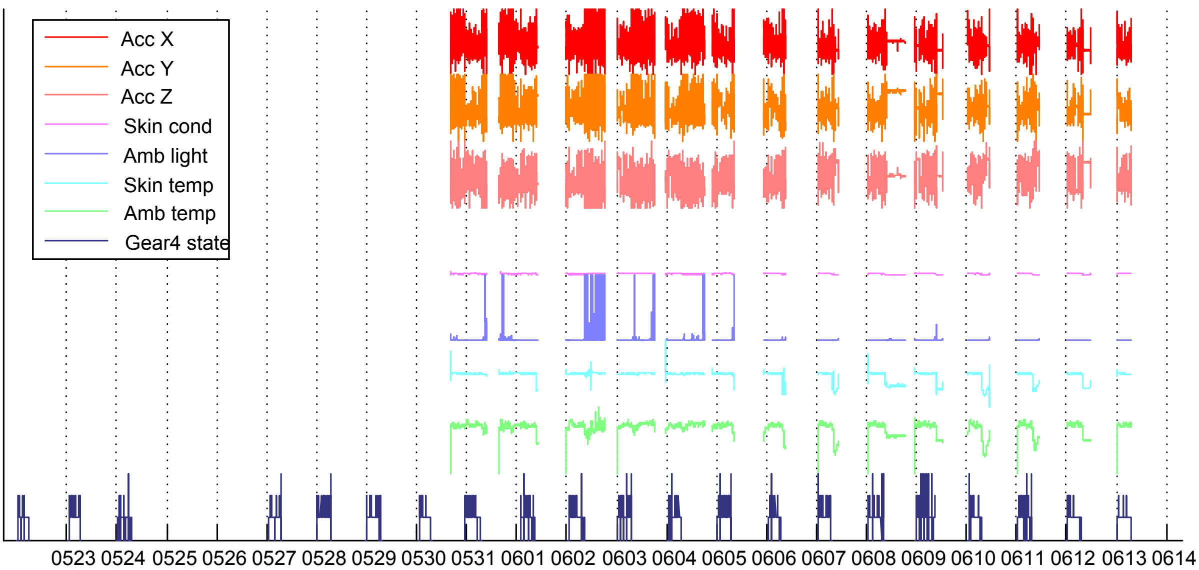

Figure 2 – Visualization of DTI-2 and Gear-4 daa, the x-axis represents time (grid lines represent midnight MMDD format), the y-axis represents the signals.

The above data allows establishing a relationship between the Gear4 sleep states and the Philips DTI-2 parameters. Gear4 is the sensor used to detect different stages of sleep. It is a commercial sensor that is using an iPhone on a dock station. It measures the breathing rate and movement without the need for any extra equipment. The patient is supposed to place the sensor in the bedside table in order that the bed area to be monitored during the night. The sensor differentiates three states:

- Awake

- Sleep (light)
- Deep sleep

The system uses a patented contactless sleep monitoring sensor so there's no need to sync, charge or wear armbands/headbands. The Renew SleepClock provides you with detailed information on how the patient has been sleeping, stores the statistics and gives helpful hints and tips on how to sleep better.

In order to match the Gear4 Renew Sleepclock states to the Philips DTI-2 accelerometer data, the DTI-2 accelerometer data is converted into activity counts per epoch of 30 seconds. The activity counts are separated into three different groups, corresponding to the classifications of the Gear4 states. The resulting probability density functions (PDFs) are shown in Figure 3.

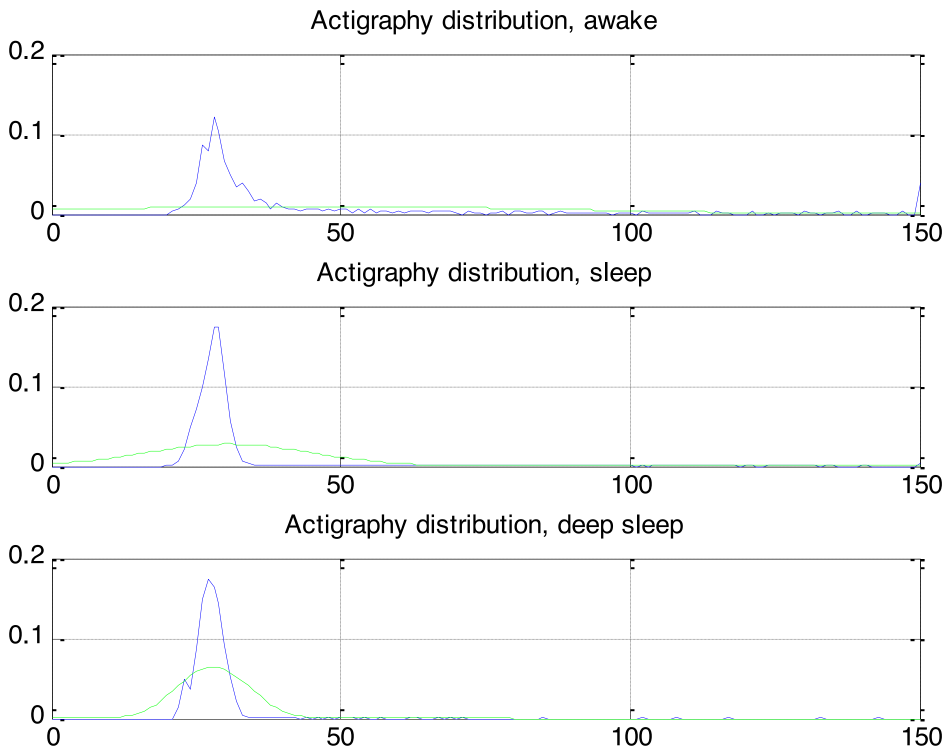

Figure 3 – Distribution of activity counts as a function of Gear-4 states (awake/sleep/deep sleep), the x-axis represents the activity count value, the y-axis represents the relative probability

As expected, the activity counts PDF for deep sleep shows slightly lower values than that of the light sleep which in turn is generally lower than that of the awake state. The PDFs of the activity counts have been modelled using a Gaussian fit (green lines Figure 3). It can be observed that the fit is very poor, especially for the awake state due to a poor match to a normal (or Gaussian) distribution.

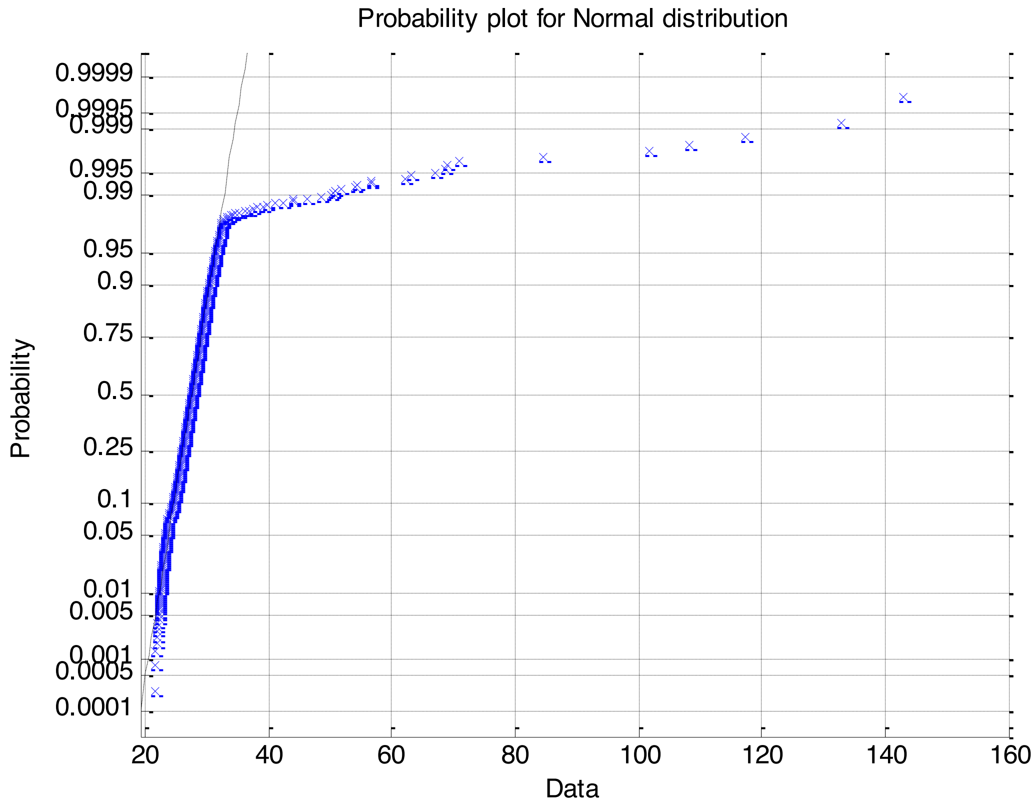

Figure 4 – Probability plot for normal distribution, deep sleep, the x-axis represents the activity count value, the y-axis represents the cumulative probability

Figure 4 depicts the histogram information during deep sleep in a different way, the x axis shows the activity count data (for calculations, see [2]), the y axis shows the probability if the data is to be modelled according to a normal/Gaussian distribution. For an ideal normal distribution this visualization would show a straight line, since the y-axis is mapped according the a normal distribution. It can be observed that for activity count data values roughly smaller than 30 the normal model corresponds quite well, whereas for larger values it doesn't. Figure 5 and Figure 6 show similar trends for the light sleep and awake states respectively, where the deviation is largest for the awake state.

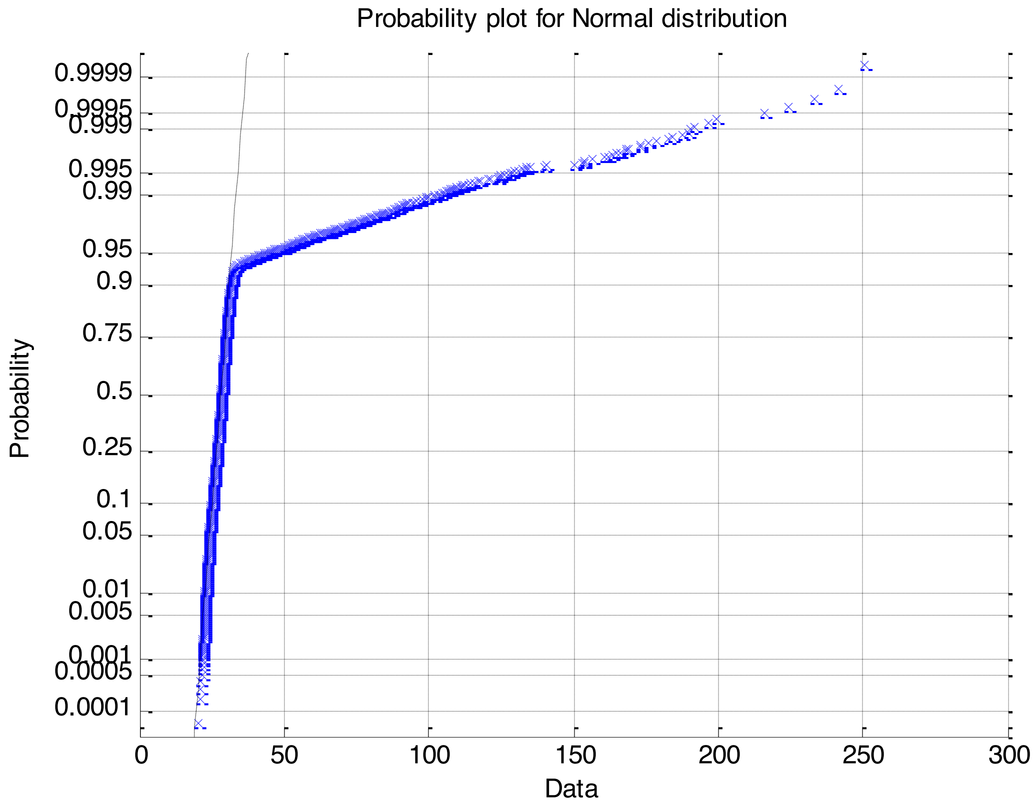

Figure 5 – Probability plot for normal distribution, (light) sleep, the x-axis represents the activity count value, the y-axis represents the cumulative probability

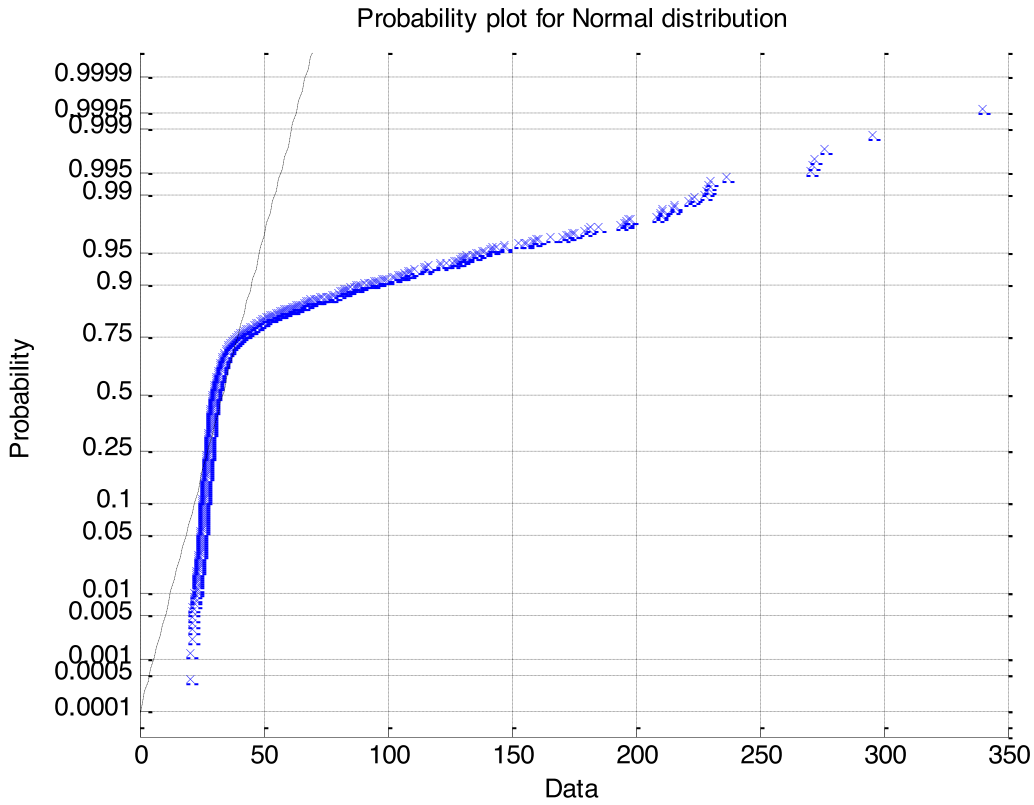

Figure 6 – Probability plot for normal distribution, awake, the x-axis represents the activity count value, the y-axis represents the cumulative probability

By remapping (or warping) the data points onto the ideal line, a much better normal approximation can be obtained. This remapping can be approximated by linearly mapping the points above the (manually determined) cut-off point onto the ideal line. The following steps are followed:

1) Manually determine cut-off point
2) 1st order polynomial fit of two portions (below and above cut-off point)
3) Map data of second portion (above cut-off) to line of first portion

The resulting probability plots are provided in Figure 7, Figure 8 and Figure 9 for deep sleep, light sleep and the awake state respectively. As can be observed they are much closer to the ideal straight line.

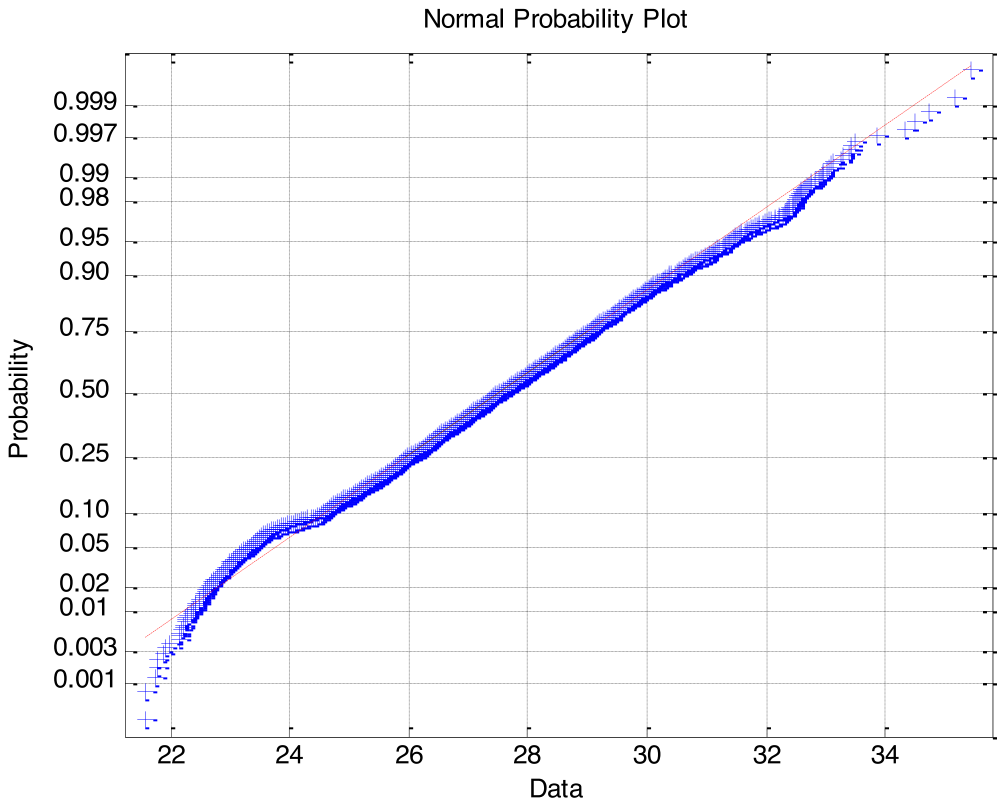

Figure 7 - Probability plot for normal distribution after remapping, deep sleep, the x-axis represents the activity count value, the y-axis represents the cumulative probability

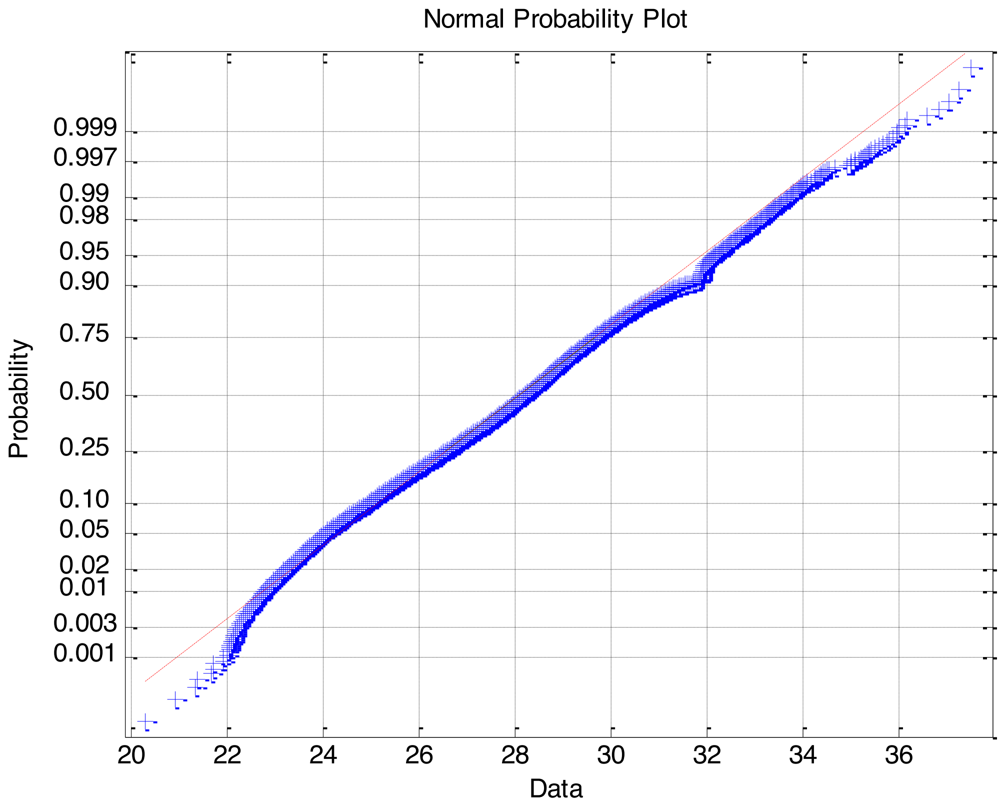

Figure 8 – Probability plot for normal distribution after remapping, (light) sleep, the x-axis represents the activity count value, the y-axis represents the cumulative probability

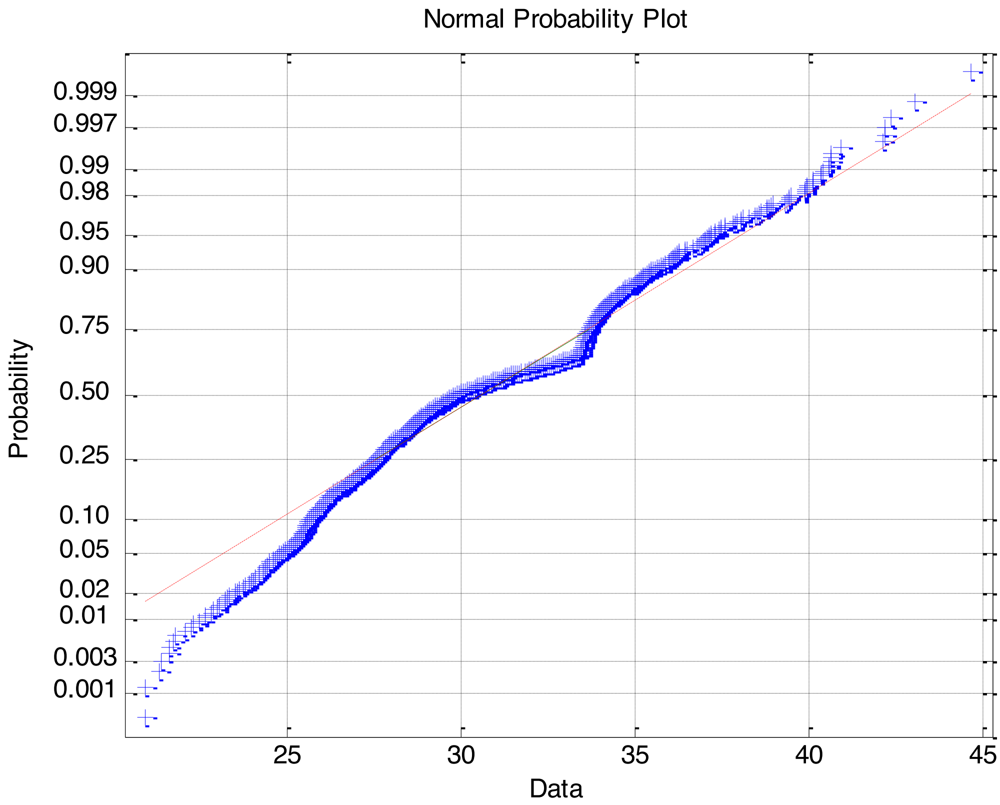

Figure 9 – Probability plot for normal distribution after remapping, awake, the x-axis represents the activity count value, the y-axis represents the cumulative probability

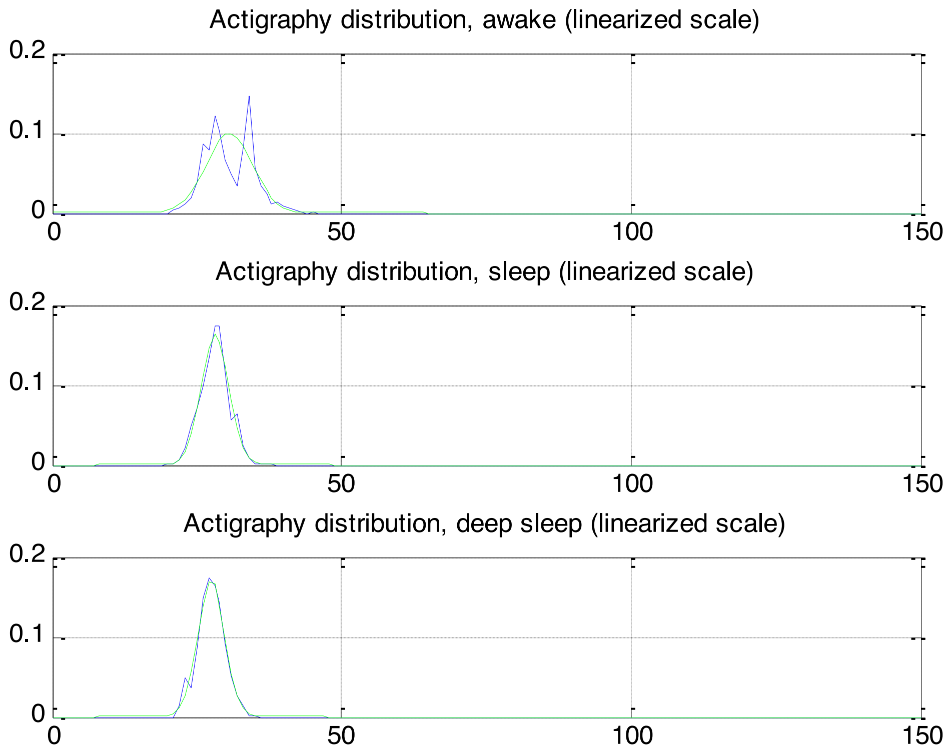

Figure 10 – Distributions after remapping, , the x-axis represents the value, the y-axis represents the relative amount of occurence

Figure 10 provides the remapped values and their Gaussian fit. As can be seen the Gaussian fit now closely matches the PDFs. Based on this fit, it is now possible to estimate the activity counts from the Gear4 states using a stochastic representation. This is visualized in Figure 11. White noise with a normal distribution and unity norm is scaled (sigma) and shifted (mu). Then, values above the predetermined threshold are remapped using the linear equations as described above. Each state has its own parameters (sigma, mu, mapping equations). Figure 12 provides PDFs of this realization. It is noted that the synthetic PDFs roughly follow the original PDFs.

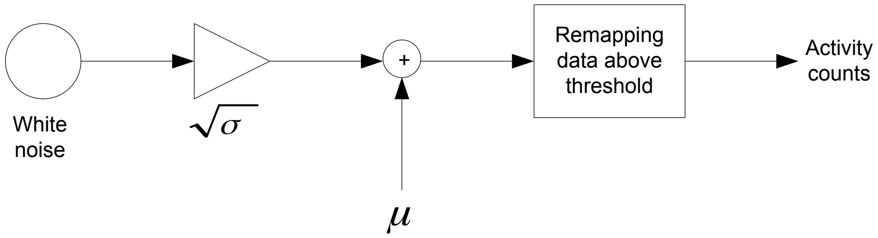

Figure 11 – Block diagram of synthetic realization

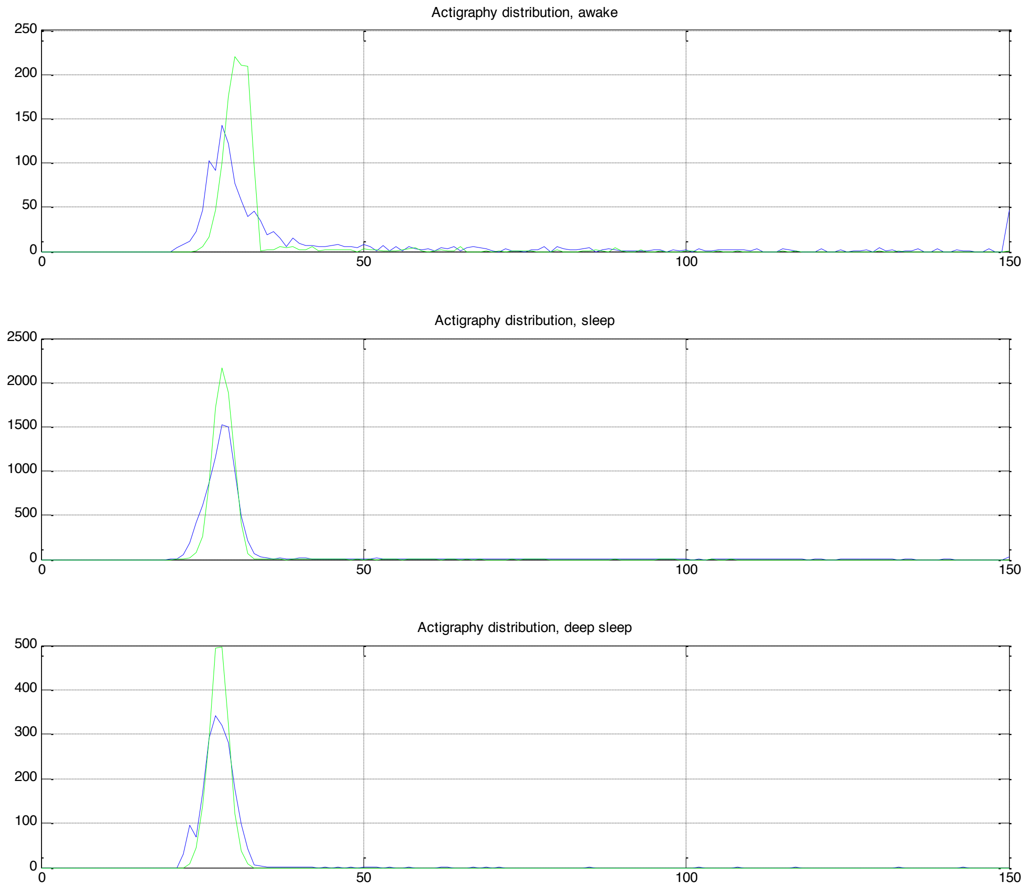

Figure 12 – PDFs of original (blue) and a remapped realization (green), the x-axis represents the activity count value, the y-axis the amount of occurence

A second lab data set was generated using the following equipment (see Figure 13):

- Philips DTI-2 (daytime)
- Gear4 Renew Sleepclock (night-time)
- MIO Alpha heart rate watch (night-time)[1]

[1] The MIO Alpha heart rate data is not included in the data fusion activity. It may be used later on during the project.

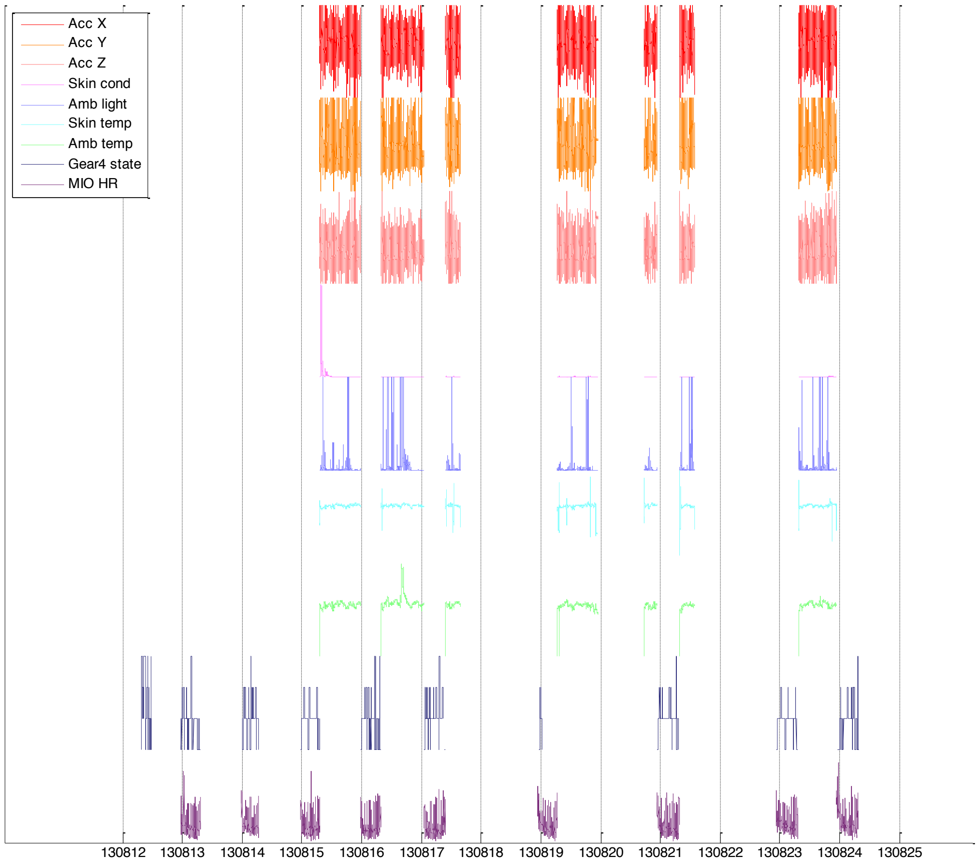

Figure 13 – Lab data, DTI-2 daytime, Gear4 & MIO night-time, the x-axis represents time (grid lines represent midnight MMDD format), the y-axis represents the signals

Using the procedure described above, the Gear4 sleep states (Figure 14, top pane) are converted to 30 second activity counts (Figure 14, lower pane, green). In addition, the DTI-2 accelerometer data are also converted to 30 second activity counts (Figure 14, lower pane, blue).

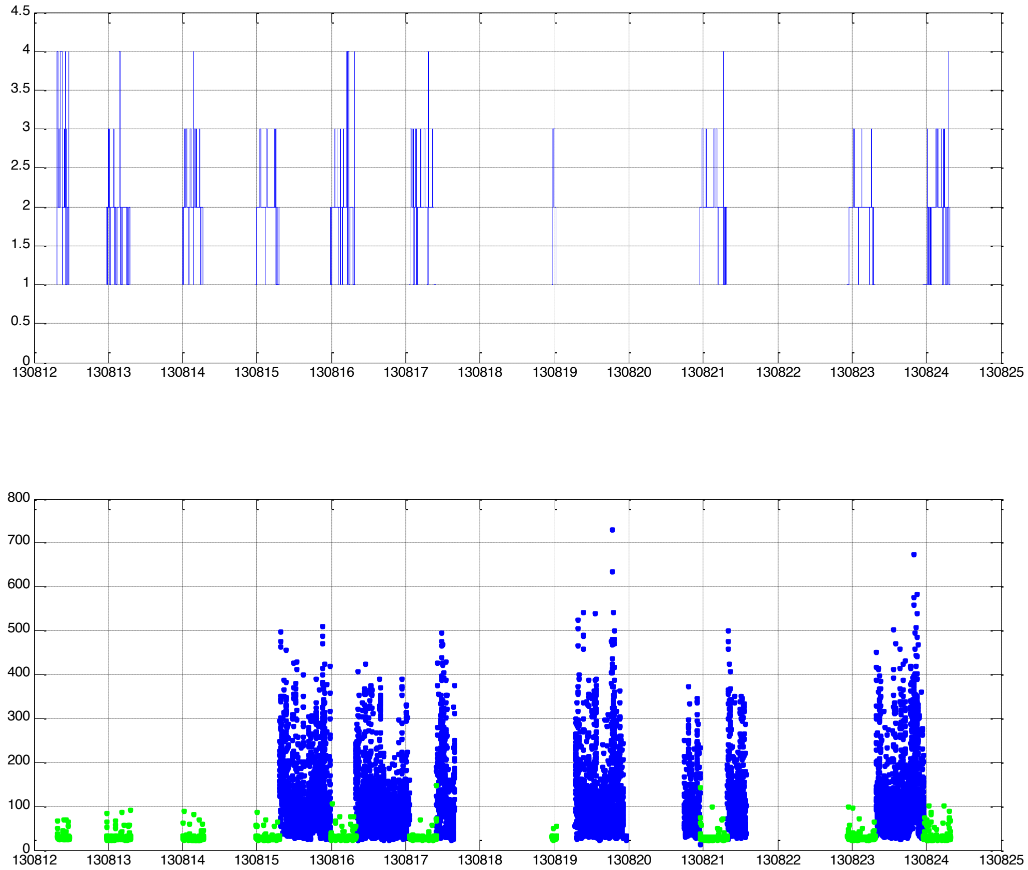

Figure 14 – Top pane: Gear4 states, lower pane: original activity counts (blue) and generated activity counts (green) from Gear4 states, the x-axis represents time (grid lines represent midnight MMDD format)

Based on the fused activity counts (Figure 14) features that are relevant for PwD can now be derived. Van Someren et al [3][4] showed that a number of parameters that can be derived from the activity counts values are relevant for PwD: inter-daily stability, intra-daily variability and amplitude of rest-activity rhythm.

In order to derive these features, first the 30 second activity count epochs are converted to an hourly grid using averaging. All hours for which more than half an hour of epochs are available are included. Figure 15 shows the resulting average hourly activity count profile (green). In order to show the amount of variation in addition the average activity count plus a single standard deviation is shown (blue).

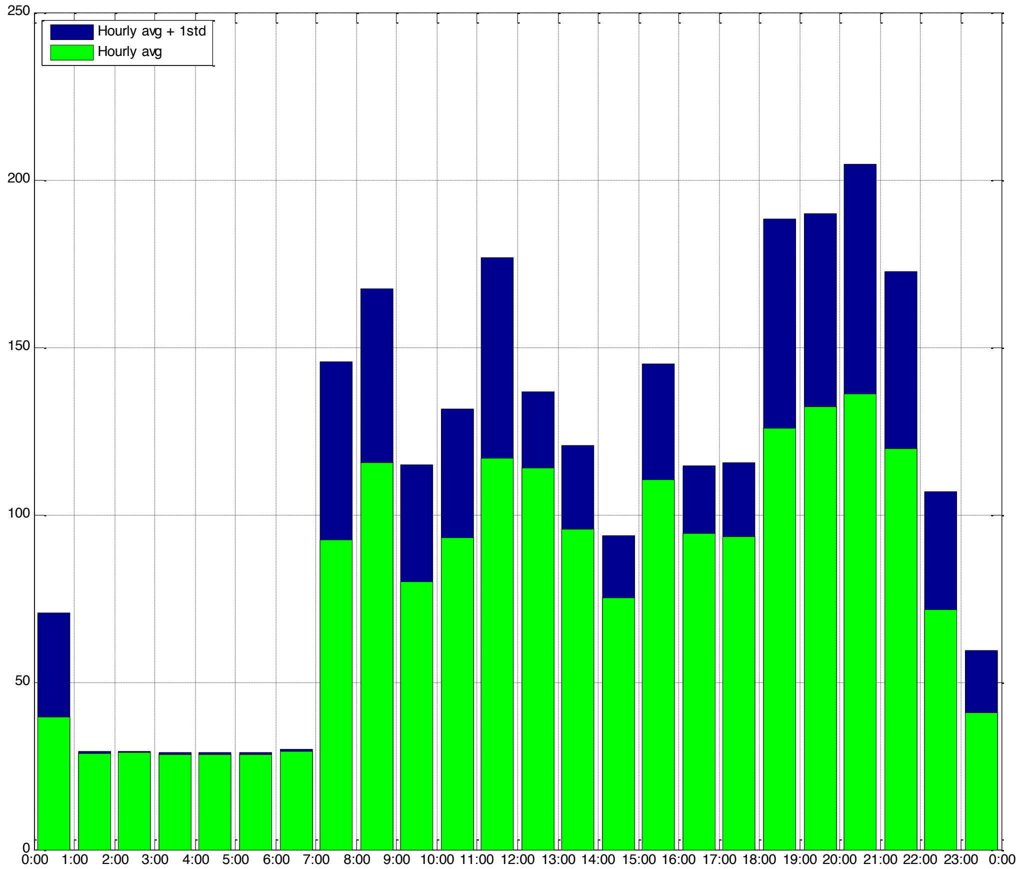

Figure 15 – Hourly average activity count for lab data

According to Van Someren et al, the inter-daily stability is calculated as:

$$IS = \frac{n\sum_{h=1}^{p}\left(\bar{x}_h - \bar{x}\right)^2}{p\sum_{i=1}^{n}\left(x_i - \bar{x}\right)^2},$$

where $\bar{x}$ represents the overall average hourly value, $\bar{x}_h$ represents the average value for hour h, n is the number of hourly values, p is the number of hours and x_i represents the hourly value at index i. Van Someren et al do not take into account any missing data. Therefore, in order to prevent imbalance, to calculate $\bar{x}$, the overall average value, the average of the hourly values $\bar{x}_h$ is calculated, instead of calculating the average of x_i. For this data, the above procedure leads to a relatively low IS of 0.565, primarily caused by a lot of variation in the sleep onsets.

The above procedure assumes that the IS does not vary significantly over time. However, if the IS is to be tracked in a longitudinal fashion, it is recommended to calculate the IS over a fixed interval length, e.g. a single week.

The intra-daily variability is calculated as:

$$IV = \frac{n\sum_{i=2}^{n}\left(x_i - x_{i-1}\right)^2}{\left(n-1\right)\sum_{i=1}^{n}\left(x_i - \bar{x}\right)^2}$$

This assumes that consecutive values exist for all i. Since our data may be incomplete, only available data pairs can be employed. Hence a slightly different definition is employed, wherein the n in the numerator becomes the length of all values and the (n-1) in the denominator becomes the length of amount of consecutive data pairs. For this data, the IV becomes a relatively low 0.746, due to irregularity of the daily patterns.

Finally, the parameter amplitude (AMP) is calculated as the difference between the means of the most active 10-hour period (M10) and the least active 5-hour period (L5). It is noted that this parameter becomes dependent on the normalization of the accelerometer data since it conveys an absolute and not a relevant difference between the active and least active periods. For the lab data, AMP is equal to 81.07. The AMP parameter, in contrast to the IS and IV parameters, is not normalized. Hence, no judgement can be made based on its value. Visually however, it is clear that there is a clear contrast between the resting (sleeping) hours and the active hours.

2.2 Multi-sensory assessment of Philips DTI-2 device usage

2.2.1 Introduction

As outlined in the previous section and in Deliverable D3.1 [2] the Philips DTI-2 is a multimodal sensing device. Despite the richness of the produced output data, the DTI-2 does not intrinsically indicate whether the device was worn or not. In principle, the data from the accelerometer could indicate whether or not the device was moving, implying that the device was being worn. However, if the accelerometer showed low or virtually no movement levels, it is unclear whether this means that the device has not been worn, or whether the person wearing the device was inactive (e.g. resting/sleeping).

2.2.2 Fusing DTI-2 sensor modalities for wearing of device

The DTI-2 allows for multiple cues to whether or not the device has been worn:

- A non-zero skin conductance, indicating that the electrodes of the device are touched by skin.
- Non-zero motion, indicating movement of the device.
- A skin temperature measured in the order of 35-38 degrees Celsius.

Based on these observations an algorithm has been devised to check whether the device was worn. The goal of the algorithm is to invalidate any data corresponding to an assumed non-wearing of the device. The algorithm currently operates on a 30 second epoch grid. A block diagram of the algorithm is shown in Figure 16. In a first step the (3-axis) accelerometer, skin conductance and skin temperature signals are segmented into 30 second epochs. Starting with the skin conductance measurement, if the minimum value of the skin conductance is larger than 0 nS, the wearing flag is initially set to 'worn'. If the minimum value of the skin conductance was 0 nS, the power of the 3D accelerometer vector after subtracting the mean of the vector is calculated. If the power is larger than a predefined threshold, set to 1.0, the initial flag is set to 'worn'. If the power is smaller than the predefined threshold, the average skin temperature is calculated. If this average value is larger than 30 degrees Celsius, the initial flag is set to 'worn'. If not, the initial flag is set to 'not worn'. At this point of the algorithm each epoch has been initially classified as 'worn' or 'not worn'. To prevent the situation that spurious 'worn' or 'not worn' decisions occur, these are filtered.

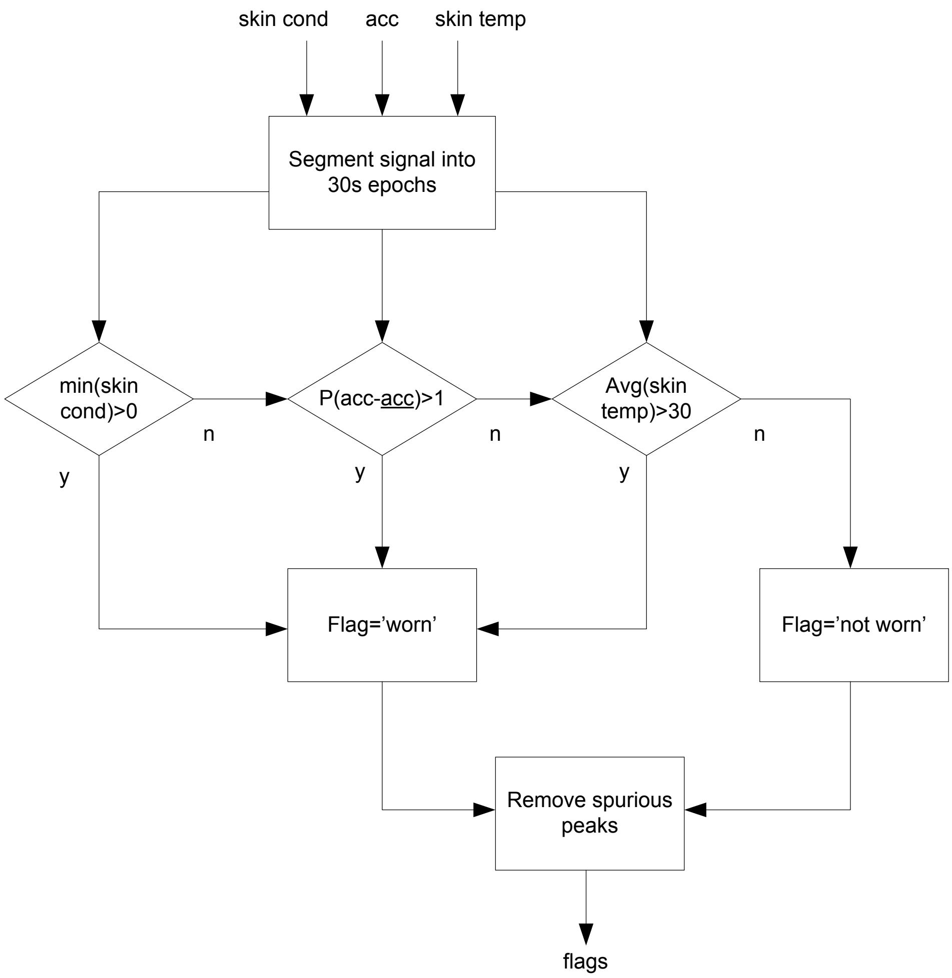

Figure 16 – Block diagram of algorithm to determine wearing of Philips DTI-2

An example of the algorithm results are shown in Figure 17. The top three panes show the minimum skin conductance, the power of the accelerometer signal and the skin temperature, all measured in epochs of 30 seconds respectively. The blue line in the lower pane shows the initial flags (1 is 'worn', 0 is 'not worn') before removal of spurious peaks. The green line shows the final flags after removal.

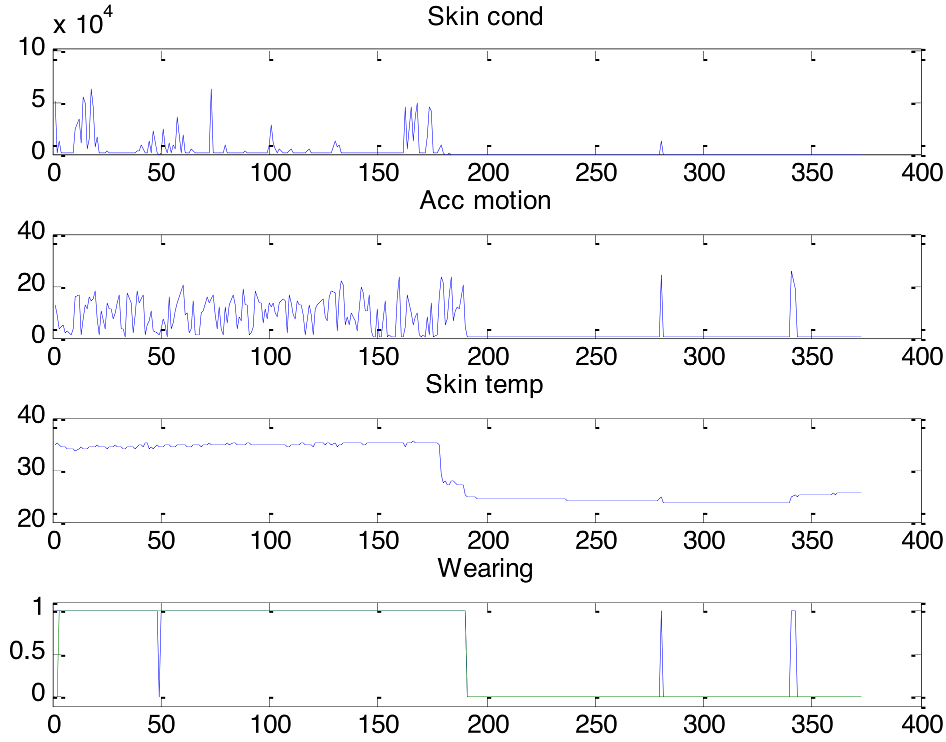

Figure 17 – Detection of wearing of DTI-2 based on skin conductance (1st pane), accelerometer motion (2nd pane) and skin temperature(3rd pane) , the x-axis represents epoch index

The algorithm has been verified on 48 recordings of approximately one hour from the Nice trial.

2.3 Signal synchronization

2.3.1 Introduction

One of the challenges of the Dem@Care system is to synchronize the clocks of all the different sensors/sensor modalities. This is especially true for devices that do not (yet) allow for real-time readout of the sensor data. The current firmware version of the Philips DTI-2 does not allow real-time readout. Instead, it saves the recorded data onto an integrated micro SD card which can only be read after connecting the device via USB to a PC. The recorded data includes timestamps from the internal clock of the DTI-2 device. The initial value of the clock of the device can be set by means of a synchronization procedure before a recording. However, due to clock drifts or (complete) battery drainage, differences with respect to the system clock (typically a PC or network time) may occur. Ideally the Dem@Care system should automatically compensate for any clock differences that may have occurred in the individual components of the system.

2.3.2 Philips DTI-2 and static camera synchronization

In order to automatically align the Philips DTI-2 clock to the Dem@Care system clock, it is beneficial to synchronize to one of the devices that is directly connected to the Dem@Care system, such as the fixed camera. Therefore, similar features need to be extracted from the DTI-2 and the fixed camera. For this purpose the video data, at 8fps, is converted into a motion estimate (see Figure 18). Each video frame is converted to gray scale. Subsequent gray scale video frames are subtracted from each other and thresholded to find regions of motion. An erosion filter is applied to remove spurious pixels.

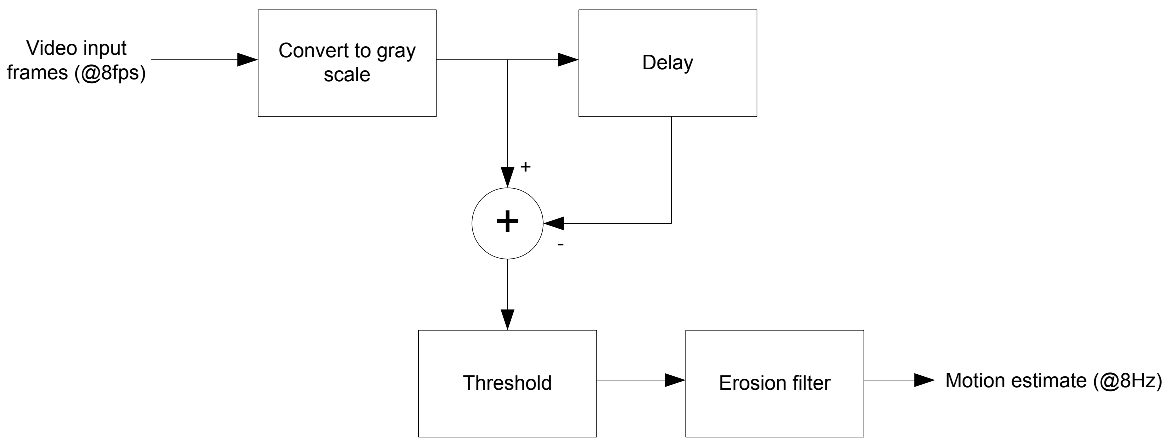

Figure 18 – Block diagram motion video motion estimation

An example of the processing is shown in Figure 19. The original colour input frame is shown in the top left pane. The top right pane shows the gray scale image. The bottom left pane shows the thresholded difference image frame. The bottom right pane shows the thresholded difference image frame after the erosion filter.

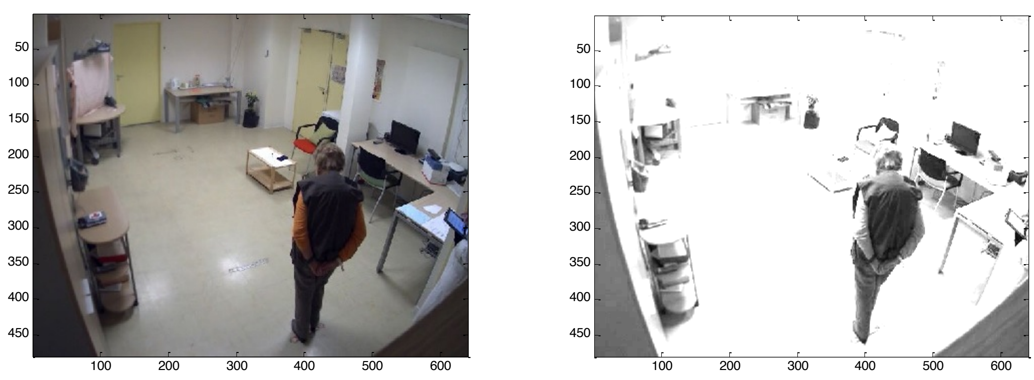

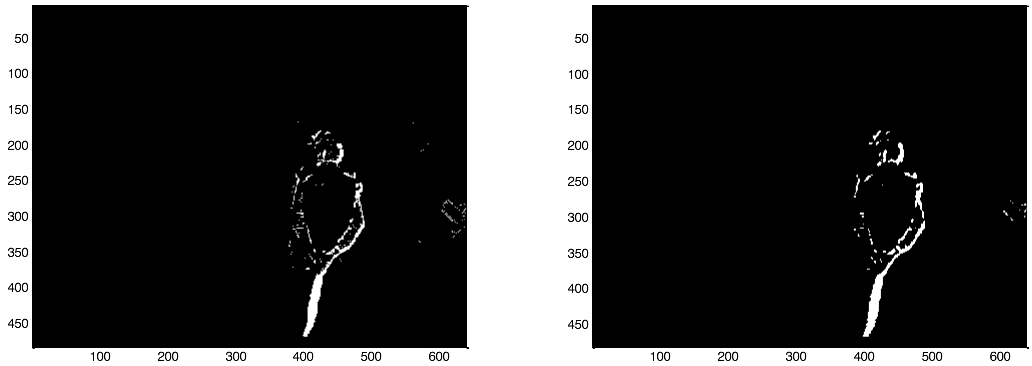

Figure 19 – Example video motion estimation

In a similar manner as described in Deliverable D3.1 [2][5], the DTI-2 accelerometer data is converted to activity counts at a fixed rate.

The camera motion feature and the DTI-2 activity count values ideally need to be synchronized on the same sampling rate. Experiments have shown that best results are obtained when the camera motion data is converted to 2Hz and the activity counts are directly calculated at 2Hz.

Figure 20 provides the camera motion feature and the Philips DTI-2 activity count, both at 2Hz for one of the recordings of the Nice lab trial.

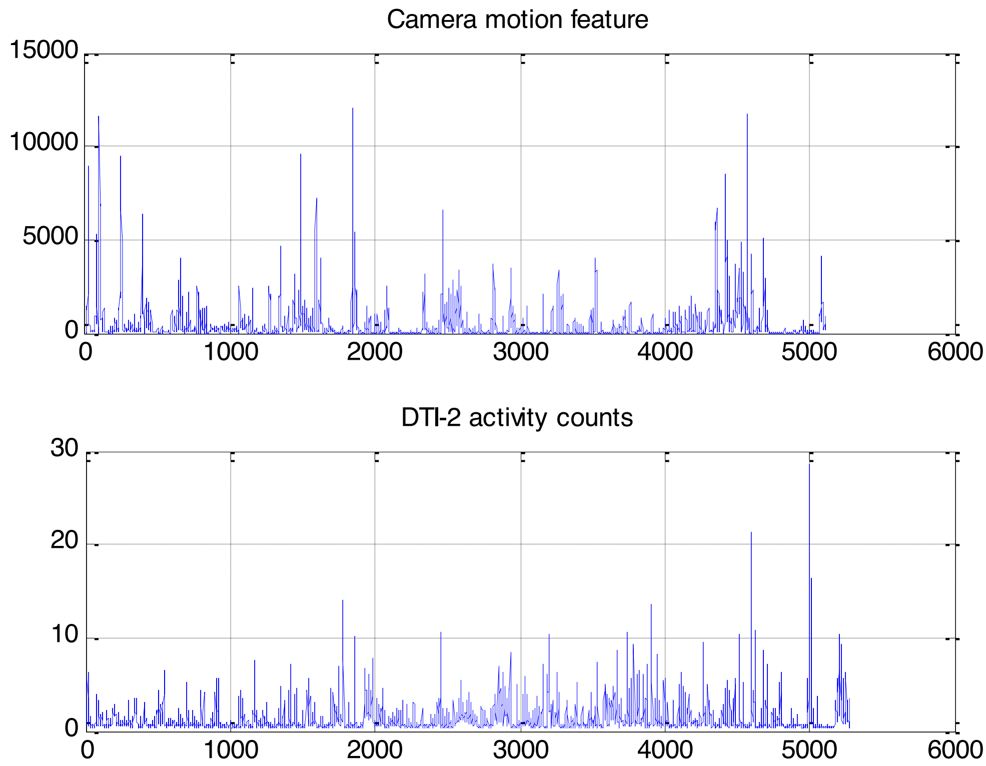

Figure 20 – Camera motion feature and DTI-2 activity counts on same sampling rate (fs=2Hz) for one subject lab data, the x-axis represents the sample index

The delay between the camera motion feature and the DTI-2 activity counts can be calculated using a standard non-normalized cross-correlation function with lag parameter *k*:

$$C_{xy}[k] = \sum_{\forall n} x[n] \cdot y[n-k].$$

Figure 21 shows the cross-correlation of the sequences in Figure 20. The peak of the cross-correlation can be found at a lag of -98, corresponding to a delay of -49s.

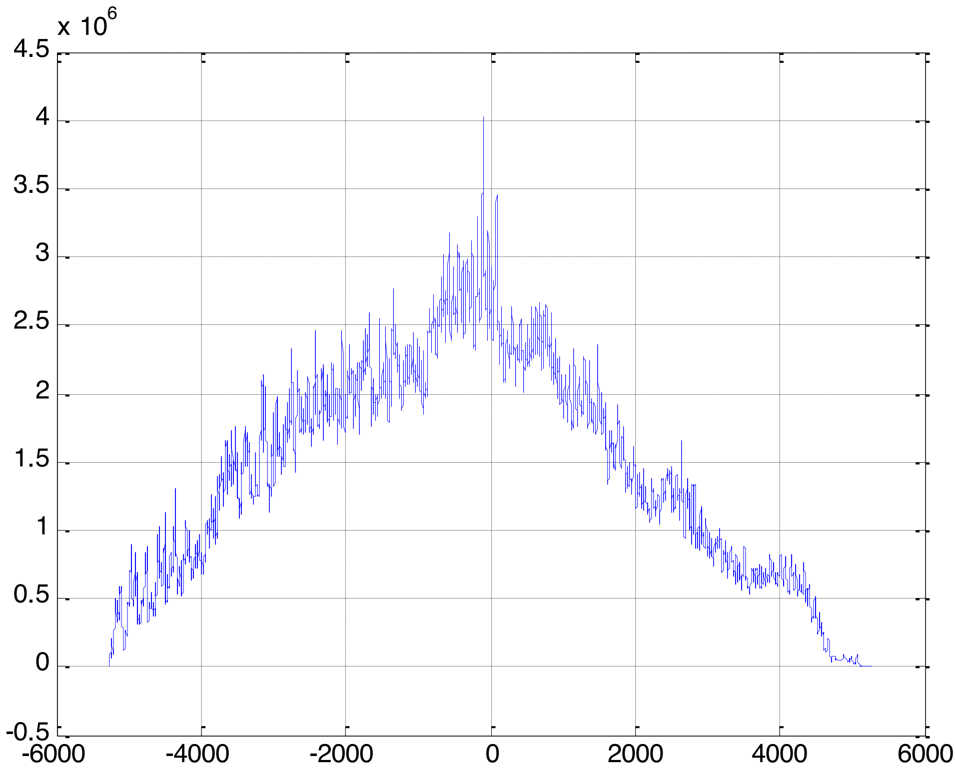

Figure 21 – Cross-correlation of video motion and DTI-2 activity count as a function of lag *k*, peak at -49s (offset -98; f_s=2Hz)

Table 1 provides a comparison between the manually aligned video/DTI-2 data (column "Manual annotation") and the cross-correlation described above (column "Xcorr"). For most signal pairs, the cross-correlation delay nicely corresponds to the delay of the manually aligned data. For three signal pairs however, there is a significant mismatch.

Table 1 – Manual annotation versus automatic delay estimation

Patient ID	Manual annotation	Xcorr	Adjusted Xcorr	Binary input
20121129a	-53.64	-2885	-53.5	-53.5
20121203a	-153.56	-152.5	-152.5	-152.5
20121203b	-57.96	-57.5	-57.5	-57.5
20121204a	-78.32	-78	-78	-78
20121211a	-105.52	-104.5	-104.5	-104.5
20121214a	-48.84	-49	-49	-49
20121217a	-35.44	-35	-35	-35
20121217b	-55.92	-56	-56	-56
20130109a	-73.08	-72.5	-72.5	-73
20130114a	-47.44	-47.5	-47.5	-47.5

20130121a	-78.72	-78.5	-78.5	-78.5
20130122a	-122.44	-122	-122	-122
20130128b	-118.96	-118	-118	-118.5
20130129a	-155.36	-154	-154	-154
20130204a	-67.44	-65	-65	-64.5
20130205a	-78.92	-78	-78	-78.5
20130206a	-119.36	-118.5	-118.5	-114
20130211b	-74.8	-75	-75	-74
20130212a	-68.52	-68	-68	-68
20130218a	-39.32	-38.5	-38.5	-38.5
20130219a	-48.16	-48	-48	-48
20130225a	-38.16	-38	-38	-37.5
20130304b	-47.84	-48	-48.5	-48.5
20130306a	-62.04	-62	-62	-62
20130311a	-65.6	-65	-65	-65
20130319a	-211.76	-211.5	-211.5	-211.5
20130326a	-50.56	-50	-50.5	-50.5
20130327a	-23.4	-23.5	-23.5	-23
20130402a	-118.28	-118.5	-118.5	-118.5
20130403a	-118.2	-334	-118.5	-118.5
20130409a	-825.48	-1019	-1017	-825
20130416a	-24.44	-24	-24	-23.5
20130416b	-65.68	-65	-65	-65
20130417a	-90.2	-89.5	-89.5	-89.5
20130422a	-48.44	-47	-47	-46.5
20130507a	-25.88	-25	-25	-25
20130514a	-63.44	-61.5	-61.5	-61.5
20130603a	-84.6	-85	-85	-84
20130604a	-23	-21	-21	-22
20130611a	-180.4	-179.5	-179.5	-179.5

To improve the results, an alternative cross-correlation function has been calculated. It uses the standard cross-correlation $C_{xy}[k]$ as input and filters it in such a way that peaks with a relatively low neighbourhood are emphasized. An example adjusted cross correlation output is shown in Figure 22 for the sequences of Figure 20. The results are also shown in Table 1 (column “Adjusted Xcorr”). Comparing to the original cross-correlation method, now only a single sequence pair provides an incorrect delay estimate.

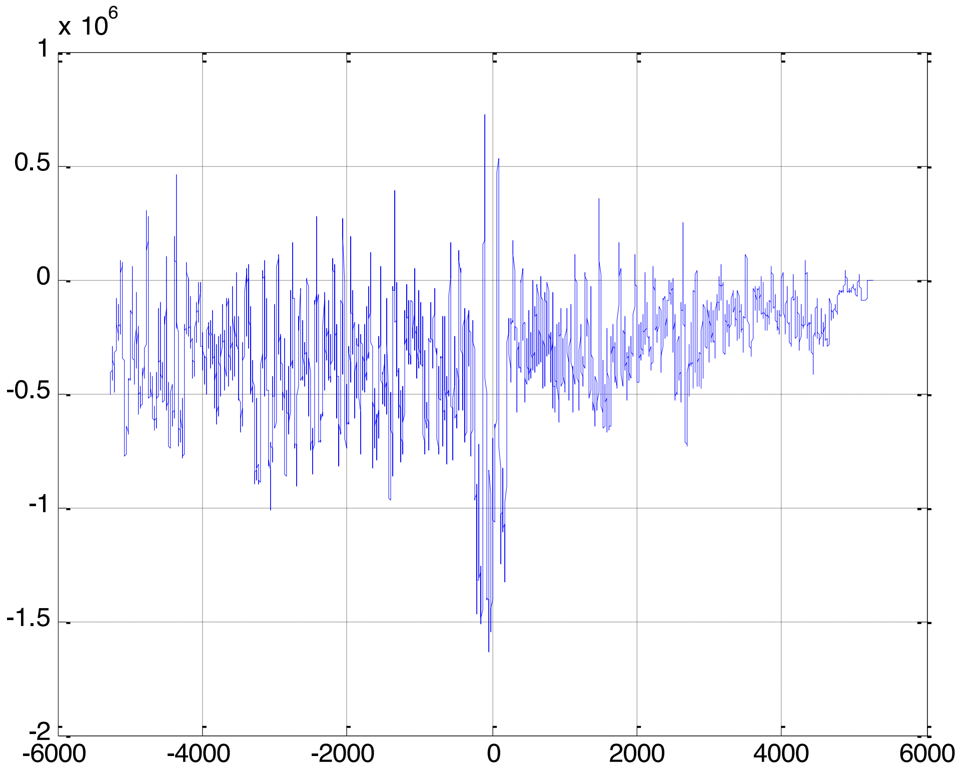

Figure 22 – Adjusted cross-correlation of video motion and DTI-2 activity count as a function of lag parameter *k*, peak at -49s (offset -98; f_s=2Hz)

To further improve the delay estimate, the motion profiles are converted to binary motion profiles by simple thresholding (see Figure 23) followed by the standard cross-correlation (see Figure 24). As can be observed from Figure 24, the peak associated to the delay between the camera and DTI-2 motion data is much more distinct than the standard cross-correlation method (Figure 21) or the adjusted cross-correlation method (Figure 22). Table 1 (column “Binary input”) indicates that the cross-correlation of the binarized motion data provides correct delay estimates for all of the Nice data.

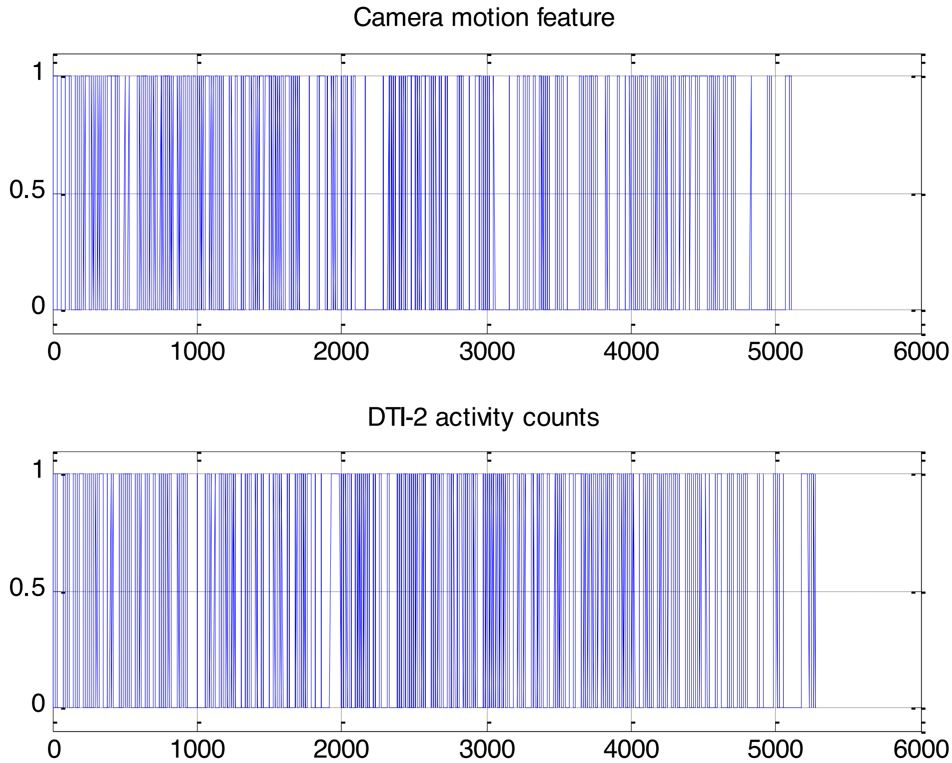

Figure 23 – Binarized camera motion feature and DTI-2 activity counts on same sampling rate (fs=2Hz) for one subject lab data

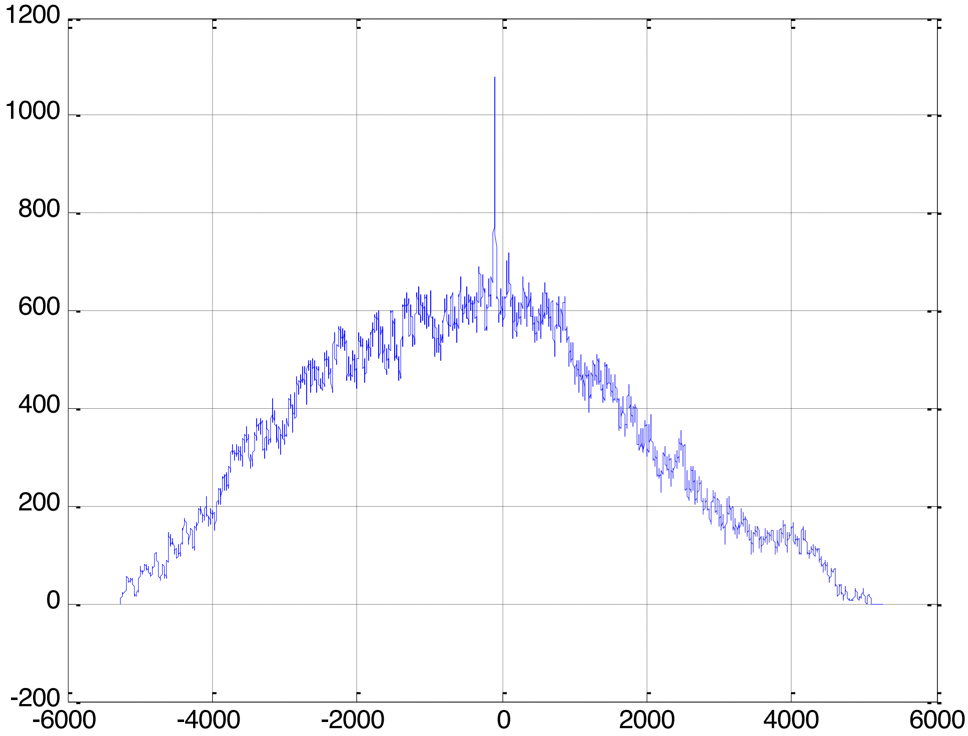

Figure 24 – Cross-correlation of binarized video motion and DTI-2 activity count as a function of lag parameter *k*, peak at -49s (offset -98; f_s=2Hz)

2.4 Conclusion

It has been shown that the Philips DTI-2 and Gear4 Renew Sleepclock data can be efficiently combined to distil meaningful parameters about the daily pattern of PwD. By statistically predicting activity counts from the sleep state values it is not necessary to where the DTI-2 24/7, but only during the time out of bed. Apart from leading to increased user comfort, this has the advantage that recharging of the Philips DTI-2 can be done during nighttime, allowing it to become part of daily routine. Future work will consist of relating the parameters derived from daily life from the fused data to the cognitive health assessment scores.

The DTI-2 device has a disadvantage that it does not autonomously detect whether the device was worn or not. This could result in incorrect classification of sleep/rest segments. By fusing data from different sensor modalities: skin conductance, accelerometers and skin temperature, a reliable way of detecting the wearing of the device is described.

Finally, a method to describe synchronization of the DTI-2 with static camera imagery is described. Since the DTI-2 has its own independent clock, deviations may occur to the system clock. This may results in incorrect measurement values especially in case of tasks of short duration, like the single/dual task of the laboratory trial. By automatically synchronizing the DTI-2 data to the camera imagery such problems can be prevented.

Based on the requirements and evaluations from the system future work may comprise the implementation of above technology into the system.

3 Fusion of accelerometer, magnetometer and gyroscope data - WIMU data

3.1 WIMU update

Tyndall Wireless Inertial Measurement Unit (WIMU) devices are used to gather data on the motion people with dementia. The WIMUs provide a platform to wirelessly collect data from inertial sensors (i.e., accelerometer, gyroscope, and magnetometer) for post-processing on the host computer.

WIMUs (Figure 25) are used to capture information such as motion and non-motion events of the wearer, changes in orientation and in posture (e.g. sitting, standing, walking, lying), time taken to complete a change of posture, pedometer as well as gait analysis, all of them obtained by advanced algorithms fully implemented by Tyndall National Institute, UCC, Ireland (a research associate of DCU).

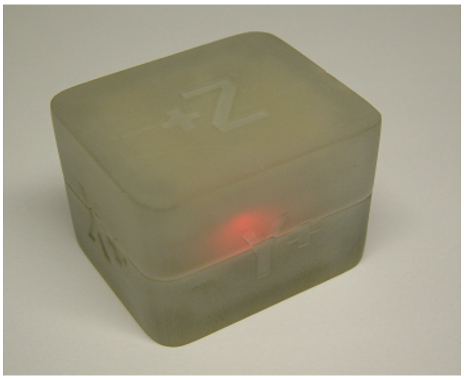

Figure 25: Tyndall Rev3 WIMU

A simplified user interface (Figure 26) has been prepared for clinical end-users. The purpose is to allow the user to interact with WIMUs in the simplest way without programming any code. Detailed documentation provides instructions on the operation of the WIMUs. This document also contains Troubleshooting and Recharging Procedure sections.

Furthermore, a short guide was written to explain how and where to attach the WIMUs on the body (Figure 27). It also gives steps to be taken before capturing data in order to simplify the features extraction (e.g. body posture and so on) in post-processing.

Several laboratory tests concerning simple physical routines (including standing, sitting, lying, running, jumping, walking and rotations) have highlighted that most of desired information for Dem@Care can be extrapolated by wearing the WIMUs in 3 positions simultaneously: on the chest, the thigh, and the lower leg.

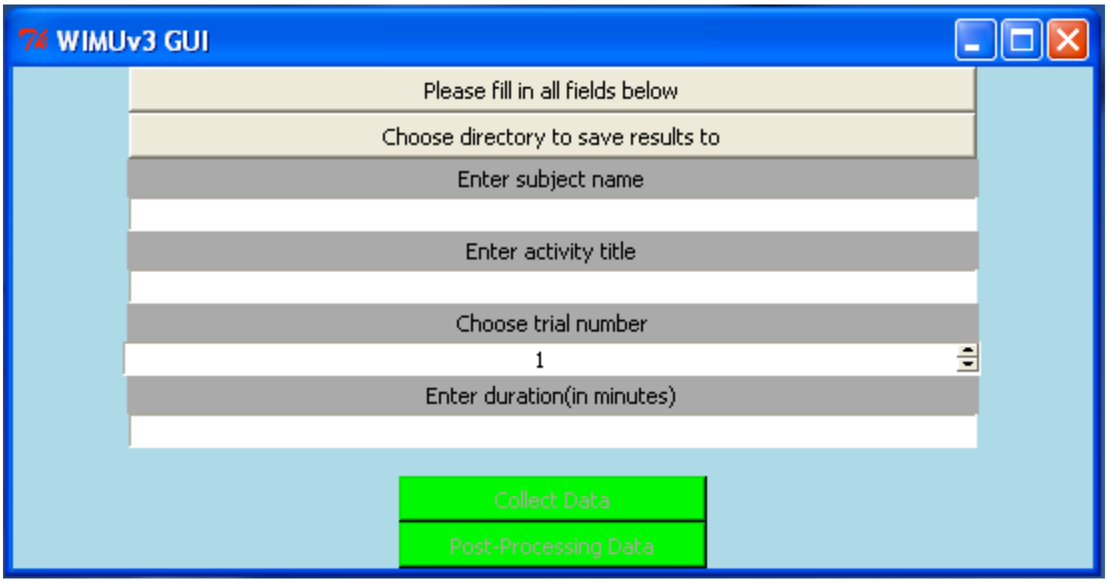

Figure 26: WIMU Graphical User Interface

Figure 27: Platforms on the body

Two complete hardware sets (including 4 WIMUs with relative enclosures, receptacles, Base Stations, and Velcro straps – one of these WIMUs is provided as a replacement) have been fully tested and assembled and sent to two project partners in Thessaloniki, Greece (to the Informatics and Telematics Institute of the Centre for Research and Technology Hellas-CERTH) and in Nice, France (to the Memory Resource and Research Centre of the University Hospital in Nice) (CHUN), in order to record activities of daily living with healthy and Alzheimer's patients.

Based on data gathered at CERTH and on data from additional in-house laboratory tests, Activity/Inactivity, Orientation, Pedometer, and Body Posture algorithms have been fully implemented and tested. Such algorithms will be commented exhaustively in the following sections.

3.2 Activity / Inactivity

Detection of human body movement and inactivity periods is a critical step for human body monitoring applications. When body movement is being monitored using inertial sensors, their output signals can be used to discriminate periods where the monitored subject is static from those where he is moving.

Such distinction is extremely important in Dem@Care, since it provides an important parameter that may be useful to quantify the progress of dementia. As an example, a person with advanced dementia may generally exhibit less activity than a subject with early-stage dementia.

Detection algorithms can be classified according to the sensor they use as an input:

- Acceleration Moving Variance Detector (AMVD) and Acceleration Magnitude Detector (AMD) use the accelerometer data. Thus, there may be erroneous detection of instants where there is no acceleration but the gyros are measuring angular velocity change.
- Angular Rate Energy Detector (ARED) uses the angular velocity signals (gyroscopes), so there may be erroneous classification of instants where there is little or no angular velocity change but accelerometers are sensing acceleration.
- Stance Hypothesis Optimal Detector (SHOD) uses both the acceleration and angular velocity signals to increase the precision of the detector.

The mathematical core of each of the detectors is as follows:

- AMVD

  A sliding window is applied over the signal in which the variance of the acceleration is computed. The figure of merit of the detection algorithm is computed as follows,

  $$V(n) = \frac{1}{N}\sum_{k=1}^{N}[|a_k - \mu|]^2 \leq \gamma$$

  where **n** is the time instant, **ak** is the acceleration vector at instant **k**, **µ** is the mean of the acceleration of the frame at instant **n**, N is the length of the frame and **γ** is the predefined threshold that characterizes the decision based on the resultant value of the figure of merit.

- AMD

  The magnitude of the gravity acceleration vector is subtracted from the magnitude of the acceleration vector which is computed at every instant.

  The figure of merit used as the input of the classifier can be computed as

$$V(n) = \frac{1}{N\delta}\sum_{k=1}^{N}(||a_k|| - g)^2 \leq \gamma$$

  Where g is the magnitude of the gravity acceleration (1 g) and δ is the variance of the acceleration signal noise that is used as a scaling factor to make the threshold less sensitive to noise.

- ARED

  The squared magnitude of the angular rate vector at each instant is compared with a predefined threshold. This can be expressed in the following way

$$V(n) = \frac{1}{N\delta}\sum_{k=1}^{N}(||w_k||)^2 \leq \gamma$$

  Where *wk* is the angular rate vector at instant *k* and δ is the variance of the angular rate noise signal, which is also used as a scaling factor.

- SHOD

  The figure of merit used as the input of the classifier is

$$V(n) = \frac{1}{N\delta_a}\sum_{k=1}^{N}\left\|a_k - g\frac{\mu}{||\mu||}\right\|^2 + \frac{1}{N\delta_w}\sum_{k=1}^{N}(||w_k||)^2 \leq \gamma$$

All the detectors have been implemented in MATLAB. Each code consists of the same structure explained below:

1. Gather Data (sample shown in Fig. 3.4)
2. By using a sliding window of N samples, the figure of merit is calculated according to the equations indicated above.

3. Compare the calculated figure of merit with a predefined threshold. For every instant k, if the value is lower than the threshold, the instant is marked as "static" (or 0); otherwise, it is marked as "active" (or 1).

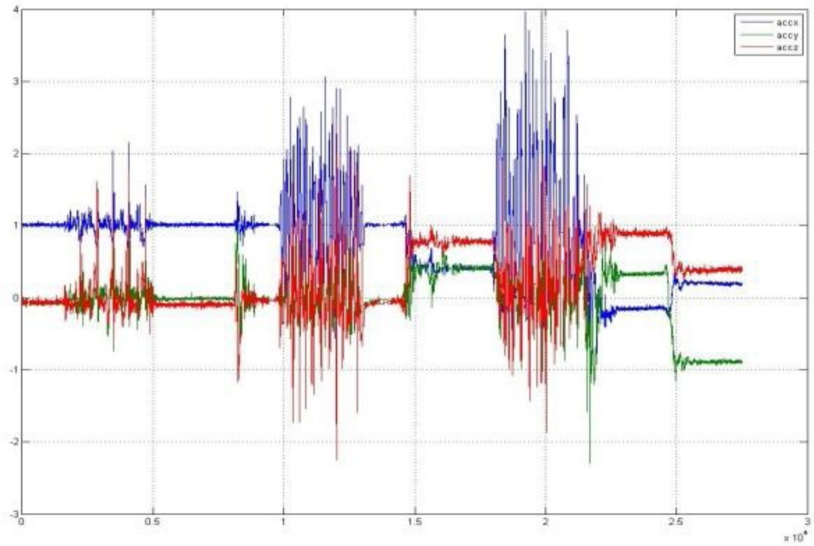

Figure 28: Example of accelerometer data collected during the lab tests (thigh case)

A comparison of the algorithms was carried out. Accuracy scores of between 92% and 95% have been achieved (Figure 29), with the SHOD detector having the best. This result is expected as the SHOD detector uses the information from all six sensors, whereas all other detectors are using only subsets of this.

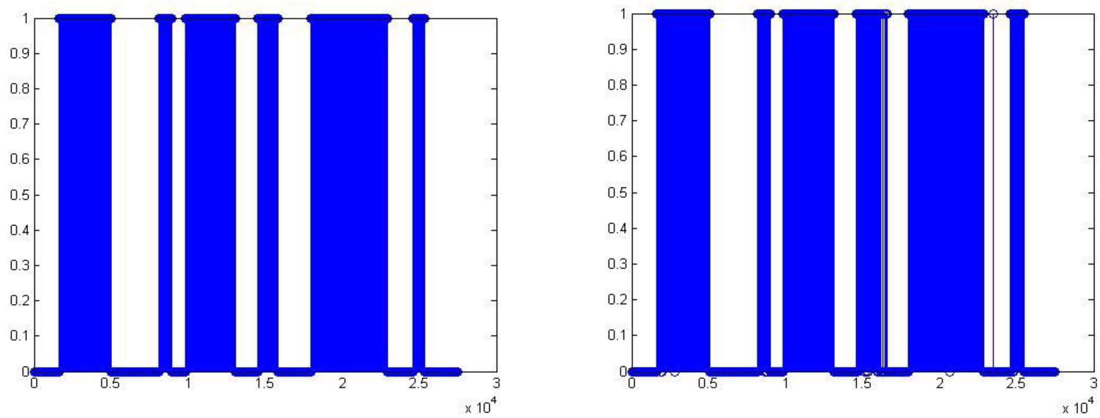

Figure 29: Activity/Inactivity - SHOD: Target (left), Estimated Results (right)

All of the detectors perform real-time classification since they base their calculations on the set of samples immediately proceeding the current time instant. Other detectors could be built that where the calculation may also use samples that follow the current time instant. This would allow for an increase in accuracy, but real-time classification would be sacrificed (i.e. data collection must be completed before any activity calculations could be performed. A small (statistically insignificant) laboratory test showed that overall accuracy could be increased from 95% to 97% with such a detector.

3.3 Orientation

A person's orientation (also indicated as heading or yaw) is the direction that the subject is pointing relative to the local magnetic North direction.

Two coordinate frames (Figure 30) need to be defined to compute the orientation of the subject:

- *Body Frame*. This frame has its origin at the inertial sensor; and each axis points along the sensitive axis of the sensor, respectively. It is a body fixed frame moving together with the body segment where the sensor unit is mounted. It is regarded as a local coordinate in all calculations.

- *Global Frame (or Horizontal Frame)*. This frame is attached to a plane normal to the earth's gravity vector. It is a fixed global coordinate. An accelerometer in its steady state can be directly used to measure the gravity vector *g,* which is always vertical to this plane. Tilt angles are then calculated from three orthogonal acceleration components as:

$$a = \sqrt{a_x^2 + a_y^2 + a_z^2}$$

$$\theta = \sin^{-1}\left({a_x}/{a}\right)$$

$$\varphi = \sin^{-1}\left({a_y}/{a}\right)$$

where a_x, a_y and a_z are the component of the gravity vector measured in the body frame, and θ and φ are pitch and roll in the global frame.

For the magnetometer sitting in the global plane, the yaw angle is always computed as:

$$\tau = \tan^{-1}\left({b_y}/{b_x} \right)$$

where b_x and b_y are the component of the geomagnetic vector measured in the global frame.

In order to calculate the yaw angle at any position, the magnetometer orientation needs to be mathematically rotated to the global plane through a transformation matrix, based on the tilt angle measured by the accelerometer, which rotates the geomagnetic vector measured in the local frame.

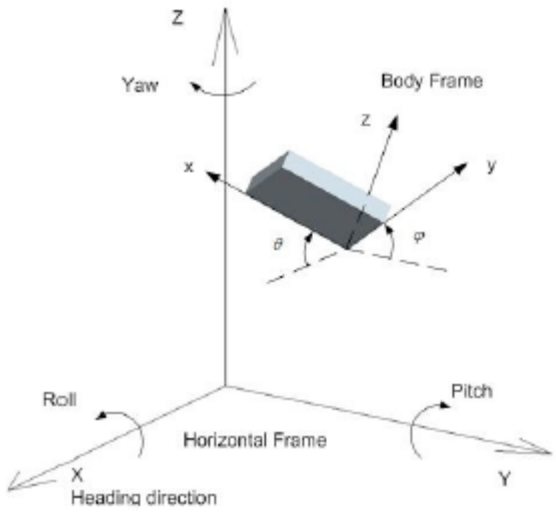

Figure 30: Body and Global Frame distinction

It is worth clarifying that the above algorithm is only valid under the steady-state assumption. Any external acceleration (i.e. any acceleration not caused by the earth gravity field), will affect the accuracy of tilt angles and consequently result in heading errors. Therefore, only the data gathered from the WIMU on the chest are taken into account for the orientation estimation, since it is the body part which presents less dynamism (e.g. low-*g* motion) compared to the other two body segments whose movements are recorded by the WIMUs (thigh and shin).

Heading accuracy may be affected by magnetic sensor errors, platform tilt errors, geomagnetic variation, and nearby ferrous distortions. Calibration of the sensors in situ reduces errors caused by tilt and nearby hard-iron distortions, whilst removing any soft-iron materials near the sensor reduce soft-iron distortions.

Figure 31a shows the estimated yaw during a test consisting of a series of 4 rotations, each one of 90 degrees. The input data are collected from the inertial sensors in the platform attached to the chest, owing to the low-g motion of the mentioned body segment (Figure 31b).

Appropriate sensors calibration and the removal of soft materials have mitigated the typical sources of error that influence heading accuracy. This has resulted in an average error rate of less than 7 degrees.

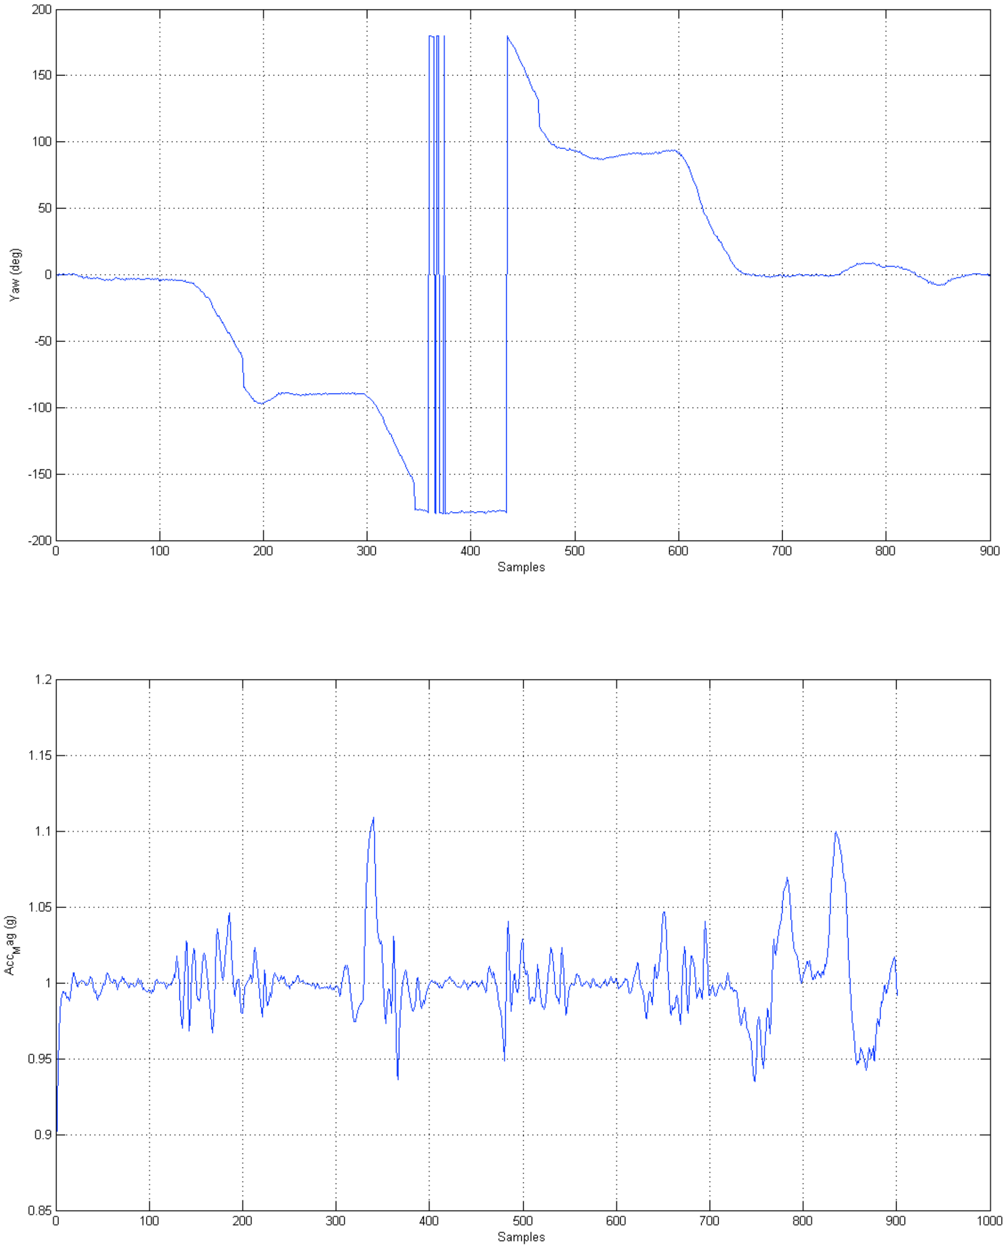

Figure 31: Estimated Yaw (a) and correspondent Acceleration Magnitude (b)

3.4 Pedometer

Dem@Care requires a long-time monitoring system which enables the analysis of the patients' locomotor coordination. This system should facilitate the quantification of rehabilitation precedures effect and measure the patient's health status. The automatic step detection would be a crucial component for the mentioned project.

Pedometers generally detect steps from vertical acceleration at the human trunk, hip, thigh or shin. The detection mechanism can be classified as either mechanical (spring, lever arm or contacts) or electrical (accelerometer). Accelerometer-based pedometers have been deemed to me more accurate and reliable as step counters. However, the accelerometer signals show considerable difference in morphology and amplitude among the individuals. Therefore, the design of a fast and robust algorithm, suitable for individual patients, and without any user-specified parameters is a must. In the present project, the acceleration signal of the shin segment has been taken into consideration for the following evaluation.

Currently, in literature, three techniques are widely adopted:

1. Pan-Tompkins Method.
    - This algorithm includes a series of filters and methods tha perform low-pass, derivative, squaring, integration for preprocess and adaptive threshold for peak-searching.
2. Template-Matching Method.
    - The main concept of the method is to generate a template, which represents a typical step cycle. An event is declared to be detected when there is a match between the accelerometer signal and the template to certain degree.
3. Peak-Detection Method based on combined dual-axis signals.
    - The third method is based on the observation that negative wave in the x- and z-axial acceleration signal are coincident within a step cycle. The preprocess procedure consists of bandpass filtering, summing up negative elements, moving-window integration and squaring operation. The preprocess procedures transform the complex signal of each step to a sloping wave of great amplitude. The peaks in the preprocessed signal are detected using a threshold, set at one quarter of the maximum in array, which defines the peak-searching interval (e.g. since there is only one peak for each peak-searching interval, the step count is equal to the number of defined peak-searching intervals).

The first method is simple to implement but fluctuations in the signal, yielding the positive and negative slopes as useful features, can result in false peak-searching intervals.

Under the condition that the initial template is correct, the second approach has the great advantage of being capable of detecting the steps and updating the template self-adaptively. However if the initial template is only estimated (as the step cycle is unknown), it may not be close enough to the real step signal, and thus will fail to perform detection accurately.

The third method is the fastest and simplest one, and provides good performance. Nevertheless, the adoption of a threshold rigidly connected to the maximum value might involve some miscounting. More specifically, subjects do not have constant step cycles, and this consideration is particularly relevant in the case of persons with dementia. Because of this, the maximum value in the signal could have a large amplitude (caused by a single very high-*g* motion; for instance, a fall) compared to other peaks in the signal, and so the correspondent threshold would be greater than the other peaks, and thus cause a miscount of the number of steps.

A comparison of four variants of the peak-detection method (method 3) was undertaken. The study uses seven distinct scenarios:

A. Walking along a line backwards and forwards
B. Walking in a circle
C. Rotating the leg taking steps
D. Rotating the leg without taking any steps
E. F. G. Random walks (mixture of previous scenarios and shuffling)

The algorithms all have the same general method:

stage 1) Input data collected from the x- and z-axial acceleration signal on the shin segment (sampling frequency 30 Hz).

stage 2) Pre-processing, equal for all the algorithms, and composed by low-pass filtering, positive values removal (set to zero), two arrays summed up entry-by-entry, and squaring operation.

stage 3) Threshold-based Peak-searching intervals. As already mentioned, the threshold is needed to extrapolate the peak-searching intervals where the foot hits the ground, and its definition is different for each algorithm.

The distinctions between stage 3 of each of the variants are as follows:

Variant 1

1. Maximum Value calculation
2. Threshold = Max / 4

Variant 2)

1. Division of pre-processed data in non-overlapping windows
2. Maximum Value calculation in each window. Max_i, where is window number.
3. Threshold_i = Max_i / 4

Variant 3)

1. Normalization
2. Create Histogram (n=100 bins) to give a model of the data distribution.
3. Threshold = level which represents x% of total energy of the histogram. (x = 85, 86, 87, 88)

Variant 4)

1. Normalization
2. Sliding windows for variance calculation (window size = 5)
3. Normalization
4. Create Histogram (n=100 bins)
5. Threshold = level which represents 85% of total energy of histogram

The low-pass filtering in the pre-process step is the same for each variant. Since the sampling frequency is 30 Hz and the bandwidth filter is 2 Hz, the *smoothing factor* in the simple infinite impulse response filter implemented is correspondent to 0.295. It can be easily proved that a cut-off frequency of 2 Hz is sufficient for the studied application. Assuming that the average walking speed for a person is 6 km/h (= 167 cm/s) and that his/her average step length is equal to 80 cm/step, it is easily deducted that the step frequency is 2 steps/sec (considering both legs). Since the utilised WIMU is attached only on one shin (so, only the step run by one leg are counted), the result needs to be divided by two (assuming there is no difference in the legs movement), which means that the final step frequency is 1 Hz.

Table 2 summarizes the results for all the algorithms in every scenario. For each scenario the actual number of steps made by the subject is indicated in brackets.

Table 2 Results of each scenario

	A (100)	B (50)	C (50)	D (50)	E (137)	F (166)	G (200)	Mean Relative Error [%]	Worst Relative Error [%]
1	106	55	35	60	131	142	64	21.83	68
2	87	50	33	62	121	152	173	14.94	34
3(85%)	101	53	44	48	138	197	223	7.7	18.67
3(86%)	103	53	45	45	138	197	205	7.27	18.67
3(87%)	94	46	43	48	126	177	205	7.02	14
3(88%)	91	46	39	48	118	156	190	9.69	22
4	132	58	57	75	149	205	199	20.68	50
5(86%)	95	50	47	48	126	174	209	4.58	8.02

The possibility to adopt a Savitzky-Golay filter, a well-known method of data smoothing based on local least-squares polynomial approximation, instead of a simple low-pass filter, was also investigated. Variant 3 has shown the best results with the simple low-pass filter so it was used as the basis for Variant 5, where the only difference is that a Savitzky-Golay filter (with order zero and frame size equal to five) has been implemented in the second part of the algorithm while the last part remains unvaried (x = 86%).

An average error greatly reduced to 4.58% (with 8% in the worst case) has been achieved. Although the method proved to be reliable and with good performance, capable of mitigating the drawbacks present in the original approach, the amount of subjects resorted for testing is rather low, making the study non-optimal in statistical terms.

3.5 Body Posture

Body posture analysis using a single accelerometer has poor accuracy and requires extremely complex operations, making real-time detection of body posture impossibile. Using a combination of signals from three accelerometers fixed on the trunk, the thigh and the shin, it is possible to reliably discriminate body posture (lying, sitting, standing) and physical activities (walking, posture transitions), providing a quantitative assessment of the daily actions. The fusion of data from the three devices provide information that would not be otherwise attainable (i.e., three devices is the minimum required).

The implemented approach includes three phases: data collection, signal processing, and final body posture detection.
Initial raw data are collected from three tri-axial accelerometers (making a total of 9 signals) attached to distinct body segments.

The signal processing is based on the adoption of:

- Savitzky-Golay (S-G) smoothing filters, used to remove noise related to high-frequency movement artifacts while preserving abrupt level changes
- Numerical gradient, used to detect the dc level changes of the acceleration signals (where the dc signal indicates the acceleration only due to the earth's gravity), shown as relatively large peaks.

More specifically, all 9 input signals are filtered by applying a S-G filter (with polynomial order zero and frame size equal to 141). The negative norm value of the numerical gradient (conveniently post-filtered with a second S-G averager) applied to the three filtered versions of each accelerometer signal, allowed the detection of postural transitions, as the time of the occurance of large amplitude peaks.

Figure 32 shows an example of the signal processing phase adopting as input the data gathered from the inertial sensor on the chest.

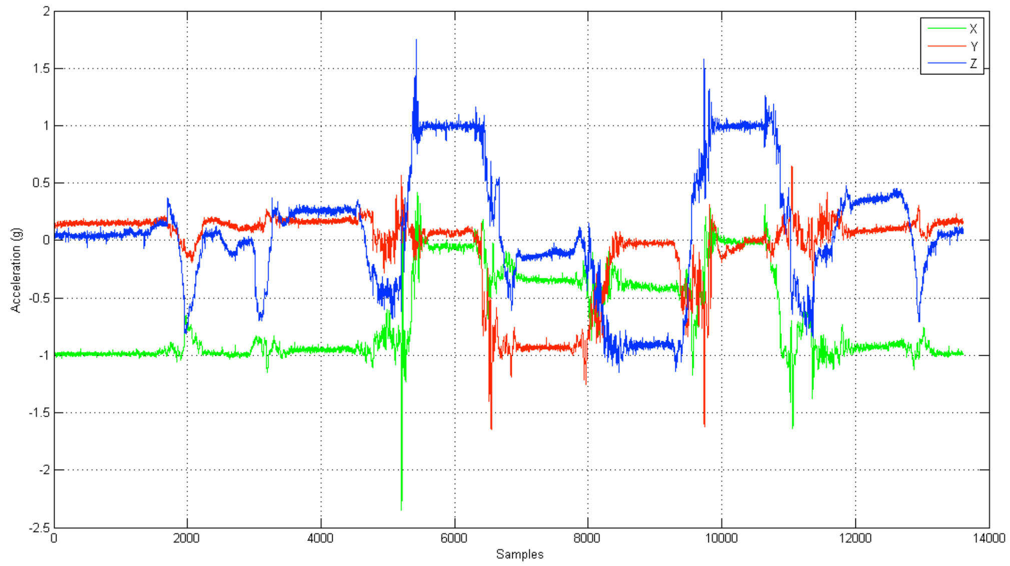

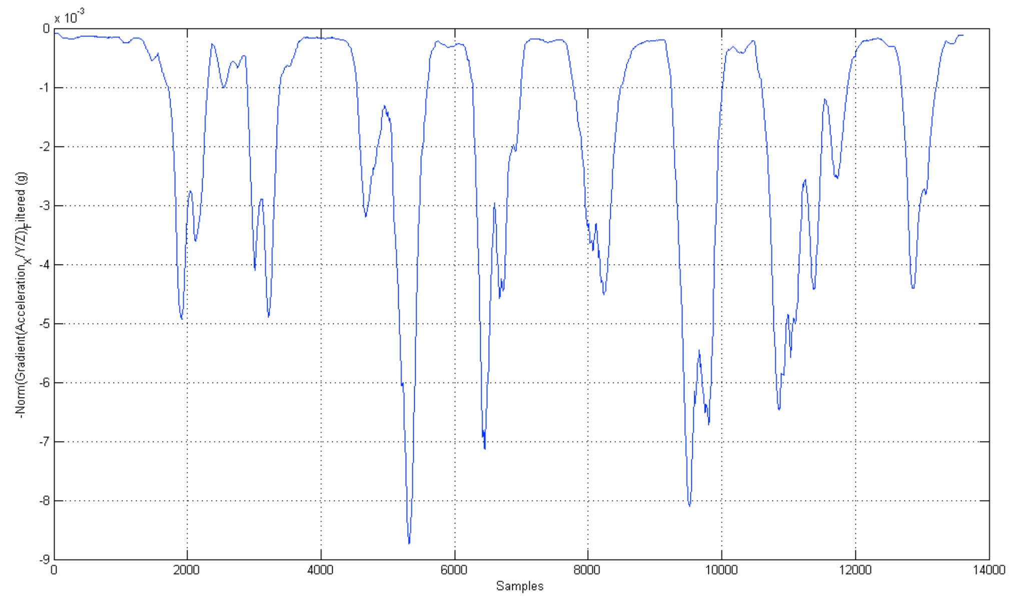

Figure 32: Accelerometer signals collected on the chest (top) during a lab test. Signal processing phase output, called Zq_{chest} (bottom). It is worth noting that the same approach has been applied also to the thigh and shin acceleration signals (producing the correspondent output Zq_{thigh} and Zq_{shin})

In order to fuse the information present in the three outputs into only one signal, the following calculation has been performed:

$$Z_q = \sqrt{\frac{Zq_{chest}\, Zq_{thigh} + Zq_{chest}\, Zq_{shank} + Zq_{shankt}\, Zq_{thigh}}{3}}$$

again post-filtered with another S-G filter. Figure 33 shows the final signal processing result.

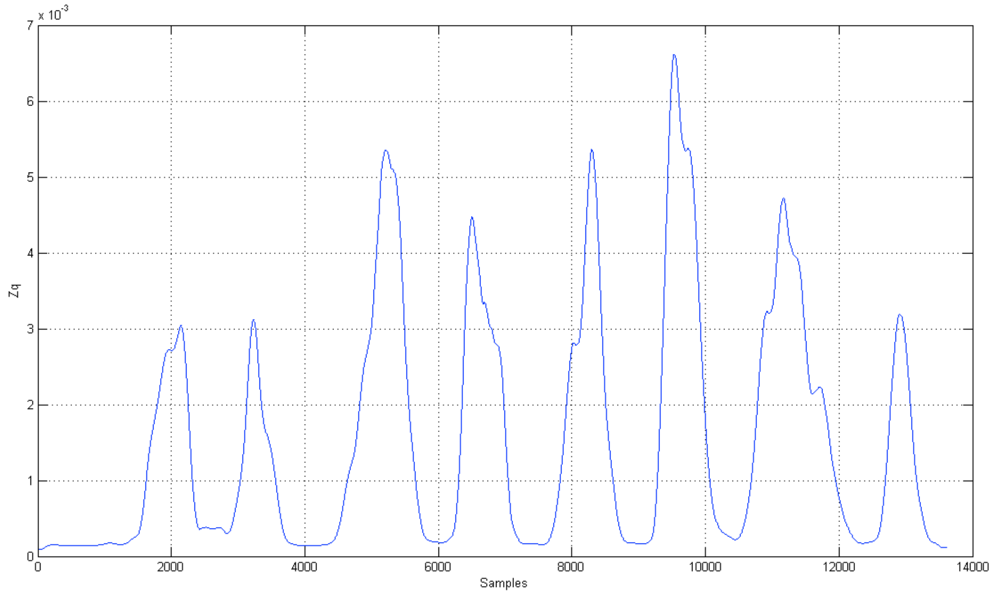

Figure 33: Final signal processing output (Zq) showing clear divisions between intervals of activity (peaks) and non-activity (troughs)

Finally, by means of a pre-defined threshold δ, the output signal is divided, in the time domain, into several sections, each one associated to different activities (static or dynamic). Figure 34shows an example of such a discrimination.

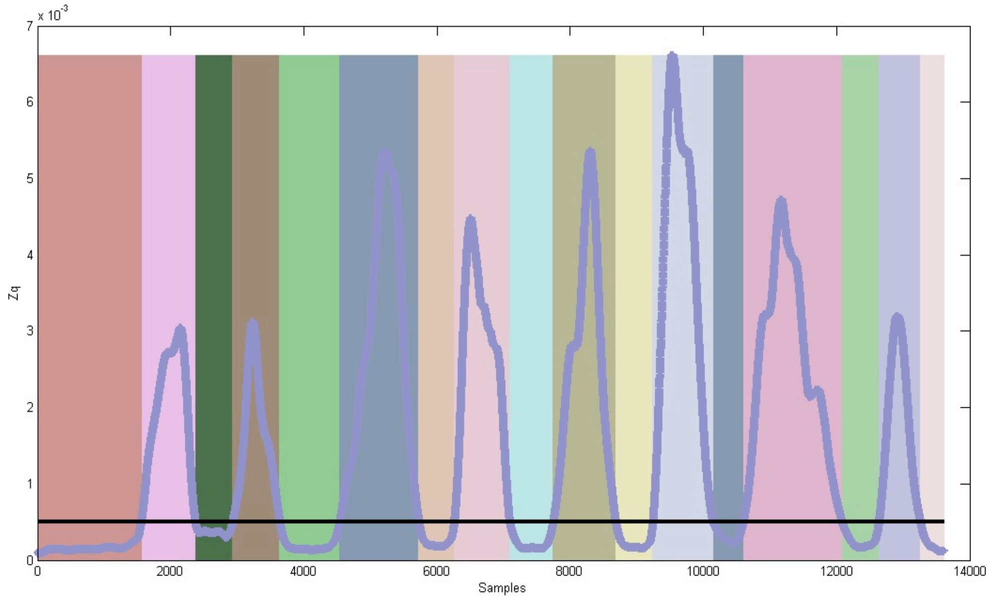

Figure 34: Activities division in sections (each one highlighted with a different color). Threshold δ = 0.0005 (black line)

The connection between each estimated section and the physical activity typology is performed in the last part of the algorithm, which consists of a body posture detector. A flowchart of the detector is shown in Figure 35. Such a phase of the approach involves two steps: dynamic motion detection and pose detection.

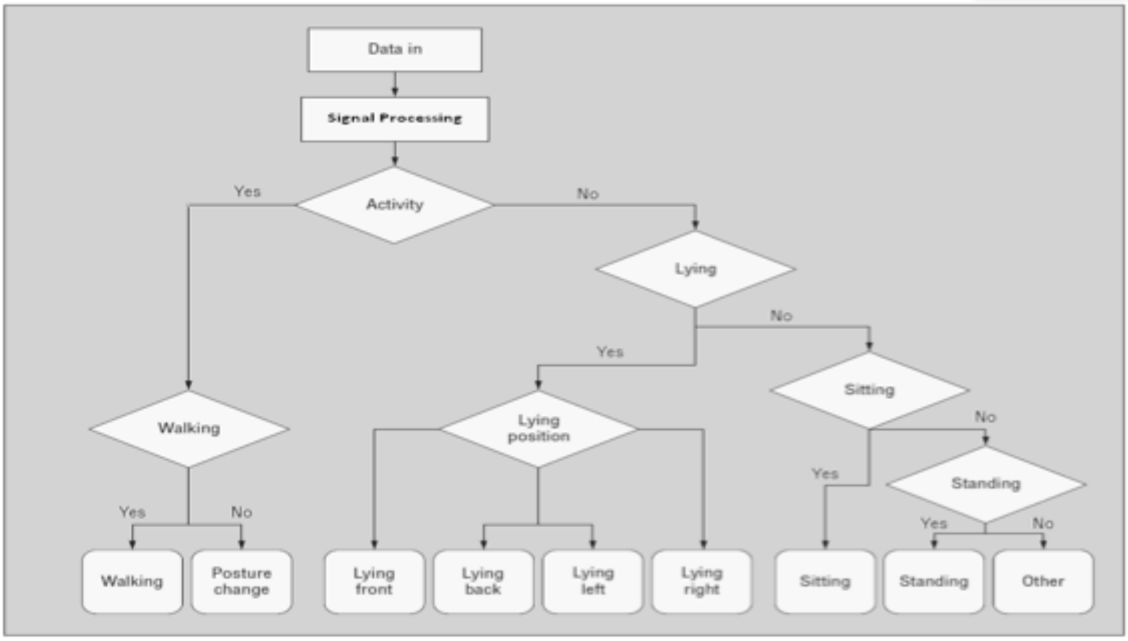

Figure 35: Body Posture detector flowchart

The dynamic motion detection is based on the assumption that when a body is static, the accelerometers respond only to gravity, producing a constant 1*g* total acceleration; on the other hand, when in motion, external acceleration is evident in the collected data, and the greater the motion, the greater the change.

Thus, the activity detector calculates for each estimated section, the relative standard deviation associated to the three initial raw magnitude acceleration data (from chest, thigh and shin). If at least in two cases out of three the relative standard deviations (RSD) are larger than a threshold ϑ, the corresponding section is marked as “active”; otherwise it is marked as “static”.

An example of the activity detector output is shown in Figure 36. The horizontal line at y=0.025 shows the predefined threshold ϑ. The sections where the RSDs do not exceed the threshold ϑ correspond to the same sections in Figure 34 where the blue curve is constantly below the threshold δ (static model).

The final stage of pose detection discerns among several poses, such as standing, sitting, lying and walking.

The walking pattern can be detected by taking advantage of the pedometer algorithm already implemented. It is sufficient to apply this algorithm to the time segments representing activity, and consider the final step counter. If this counter exceeds a specific value (4 in the current implementation) the section is marked as “Walking”; otherwise, it is considered as a “Posture or Orientation Change”.

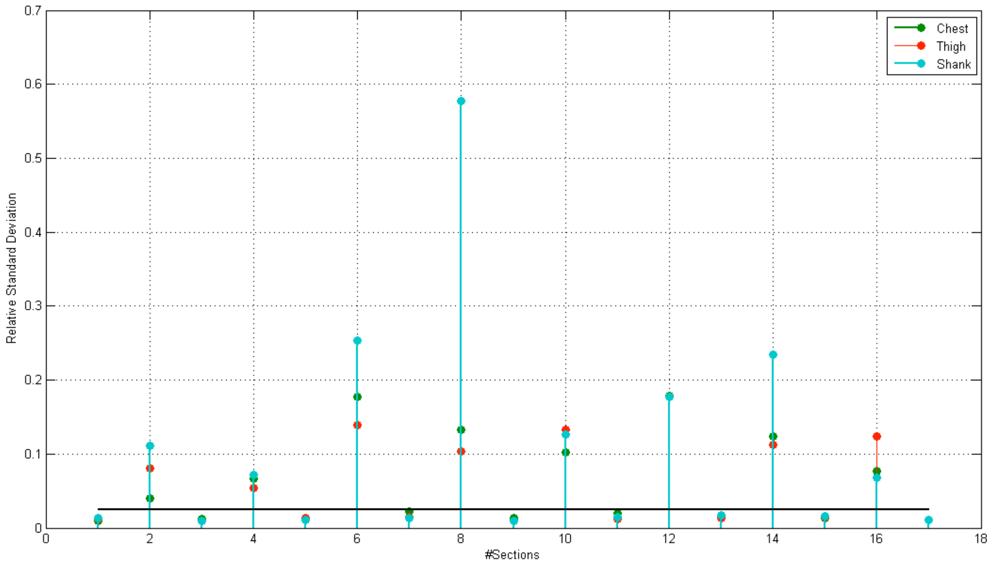

Figure 36: Example of the Activity Detector Output

For static poses, the discrimination among basic body position could be performed by means of a comparison of the collected accelerometer data with the expected values.

When a person is standing, the body should be perpendicular to the ground, and the acceleration measured on the z-axis (see the Global Frame in Figure 30) is ideally 1*g,* and is 0*g* for the other axes, for all WIMUs.

When a person is sitting, the trunk and the shin segments should be perpendicular to the ground such that each thigh is parallel to the ground; therefore, the ideal 1*g* gravity is measured on the chest and shin WIMUs z-axis and on the thigh WIMU x-axis.

When a person lies down, the body is typically parallel to the ground, and the acceleration measured on the x-axis is ideally 1*g* for all the WIMUs.

These conditions, summarized in Table 3, can help the system to recognize whether the person is currently standing, sitting or lying.

Table 3. Summary of Conditions

Axis		***x***	***y***	***z***
Standing	Acc_chest	0	0	$\sim 1g$
	Acc_thigh	0	0	$\sim 1g$
	Acc_shin	0	0	$\sim 1g$
Sitting	Acc_chest	0	0	$\sim 1g$
	Acc_thigh	$\sim 1g$	0	0
	Acc_shin	0	0	$\sim 1g$
Lying Front/Back	Acc_chest	$\sim 1g$1	0	0
	Acc_thigh	$\sim 1g$	0	0
	Acc_shin	$\sim 1g$	0	0
Lying Left/Right	Acc_chest	0	$\sim 1g$	0
	Acc_thigh	0	$\sim 1g$	0
	Acc_shin	0	$\sim 1g$	0

For each section, determinated by the signal processing phase, the mean acceleration value is calculated for each axis in any WIMUs. Following the Figure 35, the algorithm, firstly, checks the x- and y-axis acceleration signals to observe a possible lying condition. If so, it discriminates between the four possible positions (front, back, left and right) with a similar procedure and mark the section with the appropriate posture. Likewise, the feasible sitting and standing positions are investigated. In case no basic posture respects the mentioned conditions, the section is marked as “Other”, which indicates an unforeseen activity.

The detection rules consider the presence of inclination thresholds, empirically defined, which simplify the body posture analysis, mitigating the presence of sources of error. Such errors make the detection more problematic and derive, mainly, from the degree of the normal forward/backward body segments tilt and the relative movement of the WIMUs respect to the body (owing to the clothes, not skin-tight Velcro straps, or the movement of surface tissues relative to the skeleton). A graphical representation of the inclination threshold arrangement for the lying position is shown in Figure 37.

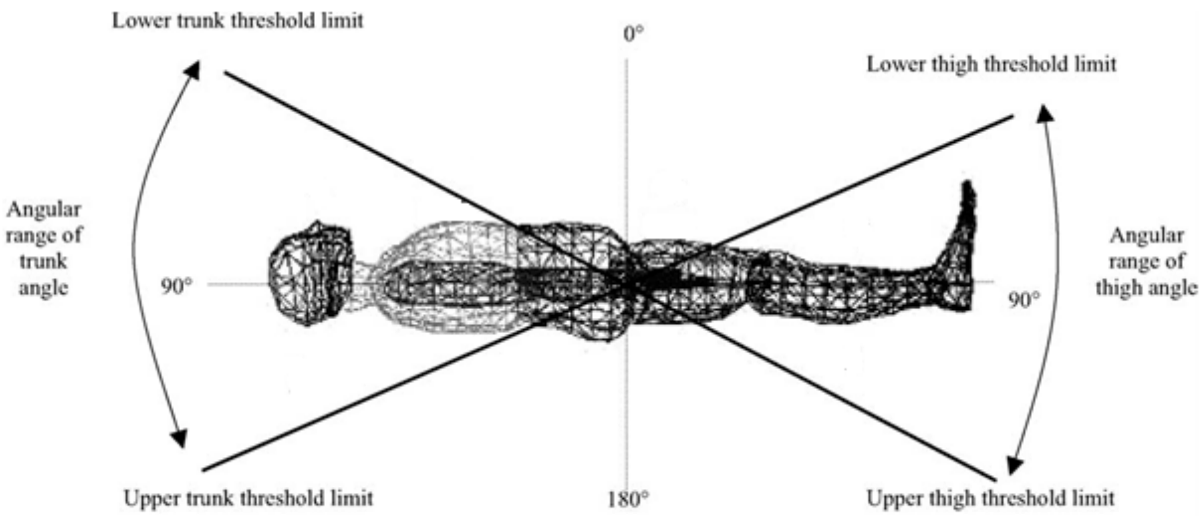

Figure 37: The inclination threshold arrangement for lying position. If the trunk, thigh and shank average x- or y-axis acceleration are included in the range [0.7-1.3] *g*, the subject is deemed to be lying.

The algorithm has been implemented in MATLAB. To determine its accuracy, seven different experiments have been performed, each one consisting of a distinct movement sequence. The real activities and the estimated ones are indicated in the following Tables 3.III-3.IV.

Table 4 Real and Estimated Activities

Experiment 1		Experiment 2		Experiment 3		Experiment 4		Experiment 5		Experiment 6		Experiment 7	
Actual	***Estimate***	***Actual***	***Estimate***	***Actual***	***Estimate***	***Actual***	***Estimate***	***Actual***	***Estimate***	***Actual***	***Estimate***	***Actual***	***Estimate***
Standing	Standing	Sitting	Sitting	Standing	Standing	Standing	Standing	Standing	Standing	Standing	Standing	Sitting	Other
Transition	Transition	Transition	Transition	Transition	Transition	Transition	Transition	Transition	Transition	Walking	Walking	Transition	Transition
Sitting	Sitting	Lying Front	Transition	Standing	Standing	Sitting	Sitting	Sitting	Sitting	Standing	Standing	Standing	Standing
Transition	Transition	Transition	Transition	Transition	Transition	Transition	Transition	Transition	Transition	Transition	Transition	Transition	Transition
Standing	Standing	Lying Left	Lying Left	Standing	Standing	Standing	Standing	Lying Front	Lying Front	Standing	Standing	Sitting	Sitting
Transition	Transition	Transition	Transition	Walking	Walking	Walking	Walking	Transition	Transition	Walking	Transition	Transition	Transition
Standing	Standing	Lying Front	Lying Front	Standing	Standing	Standing	Standing	Sitting	Sitting	Standing	Standing	Lying Front	Lying Front
Transition	Other	Transition	Transition			Transition	Transition	Transition	Transition			Transition	Transition
Sitting	Sitting	Sitting	Sitting			Standing	Standing	Standing	Standing			Lying Right	Lying Right
Transition	Transition					Walking	Transition					Transition	Transition
Standing	Standing					Standing	Standing					Lying Back	Lying Back
												Transition	Transition
												Lying Front	Lying Front
												Transition	Transition
												Sitting	Sitting
												Transition	Transition
												Standing	Standing

Table 4 Real and Estimated Activities

On the whole, the correct activities detections have been 66 out of 71 total actions (accuracy=93%).

Although the algorithm has demonstrated to provide good performance, its robustness should be still proved in statistical terms. Moreover, some parameters need to be specified, such as the threshold value for the division of the post-processed signal in temporal sections, the threshold value for detecting the kinematic/static model, or the inclination threshold arrangement for the static body posture distinction, which, if incorrectly defined, will involve a series of activities misdetection. A procedure to automatically set these parameters is planned as a future development.

3.6 Data Analysis

In February 2013, a multi-sensor Activities of Daily Living recording took place at the Greek Alzheimer's Association in Thessaloniki, with approximately 15 healthy participants and 15 Alzheimer's patients (Figure 38).

The participants were asked to perform the listed activities, while wearing a WIMU system (as well as other sensors):

Enter room, Open a snack, Eat, Discard to bin, Fix a drink, Drink, Pick up the phone, Dial number, Hang up phone, Open the cupboard, Take a book from the cupboard, Read while seated, Handshake.

An example of a collected raw data set for one subject is shown in three figures: Figure 39-Figure 41.

Figure 38: A participant to the multi-sensor Activities of Daily Living recording at the Greek Alzheimer's Association in Thessaloniki (Feb. 2013)

The data gathering procedure showed several inconsistencies. Firstly, the association between each body segment and its WIMU has not been univocal and constant for all the subjects. Although, this ambiguity is not particularly grave, it further complicates the final data analysis. Moreover, the experiments were run on a Windows 8 operating system, which, for some unforeseen reason, meant that the time-stamp on each value provided by the inertial sensors in the platforms was void. Therefore, the time synchronization between collected data can be only estimated in post-processing, so a quantification of the estimation accuracy is not achievable.

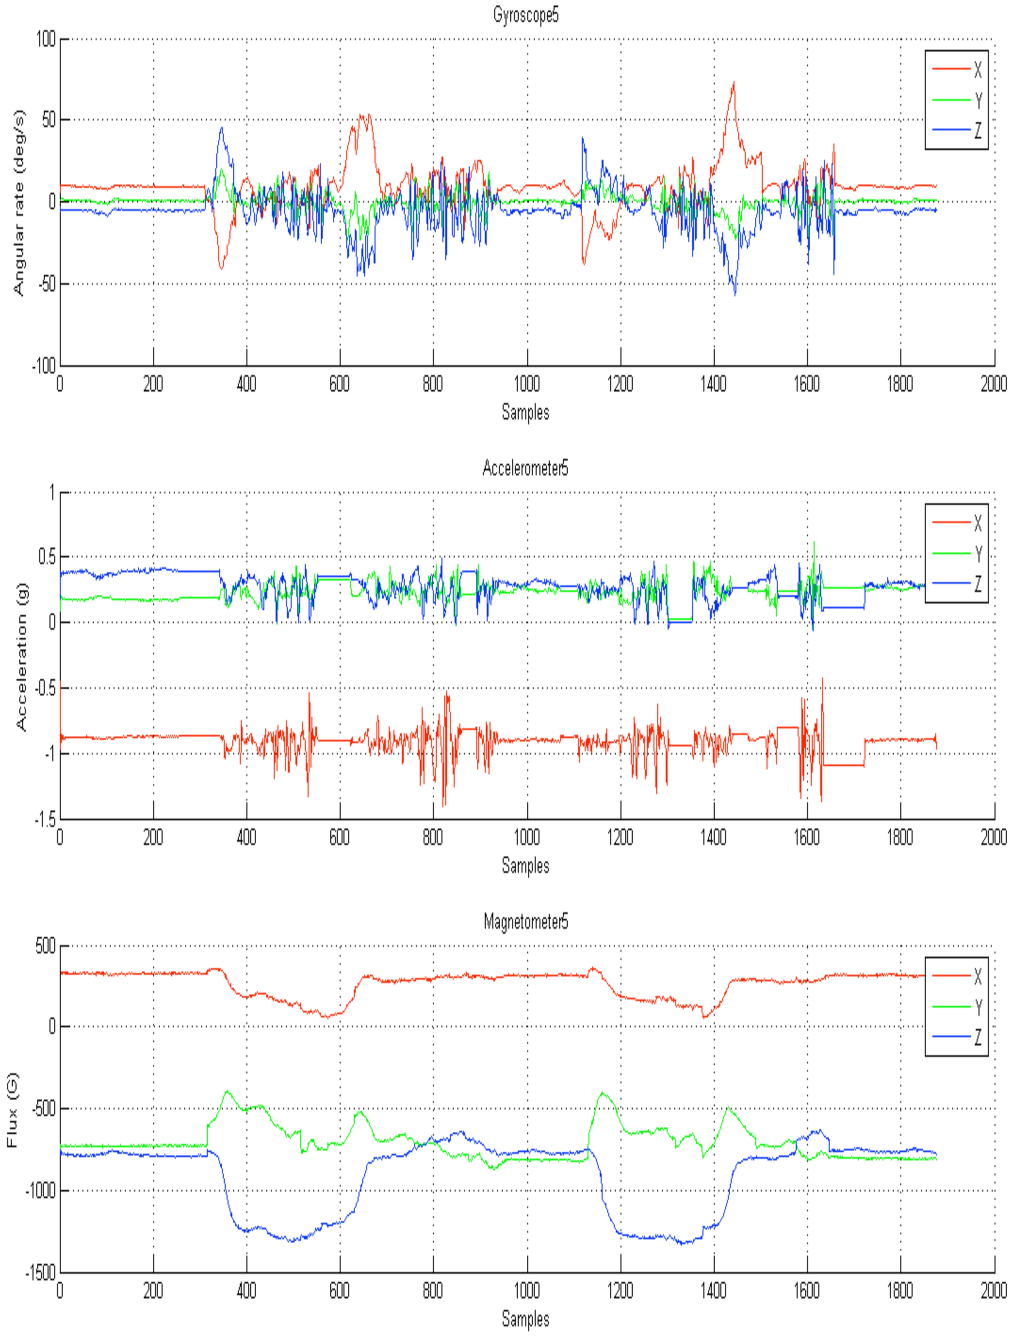

Figure 39: An example of a sample set of initial raw data collected in Thessaloniki (trunk segment – subject 6) showing the correspondence between the sensors for different movements.

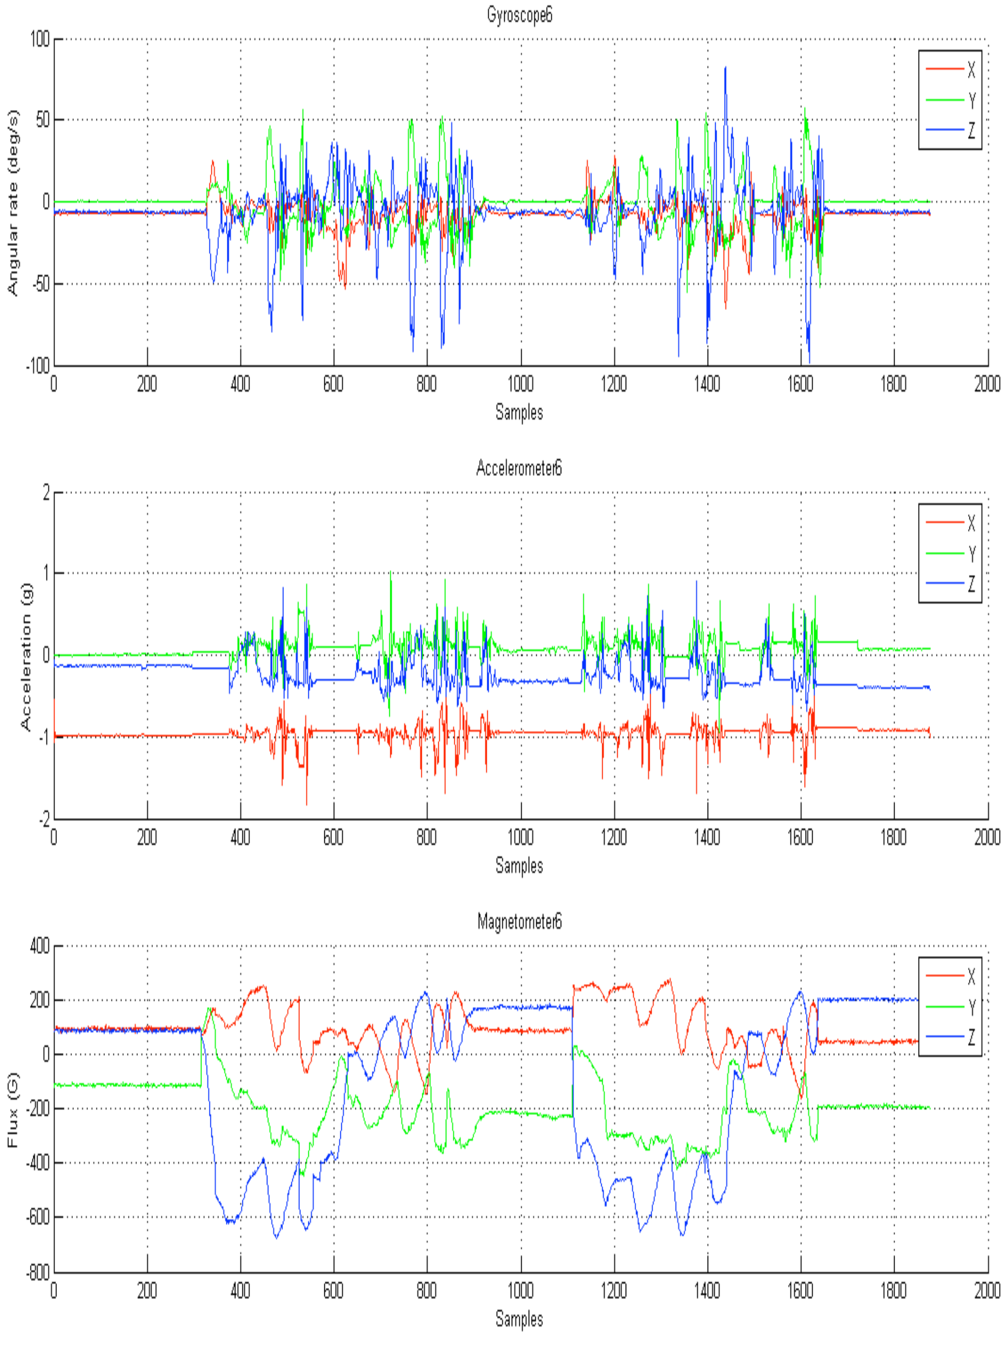

Figure 40: An example of a sample set of initial raw data collected in Thessaloniki (thigh segment – subject 6) showing the correspondence between the sensors for different movements

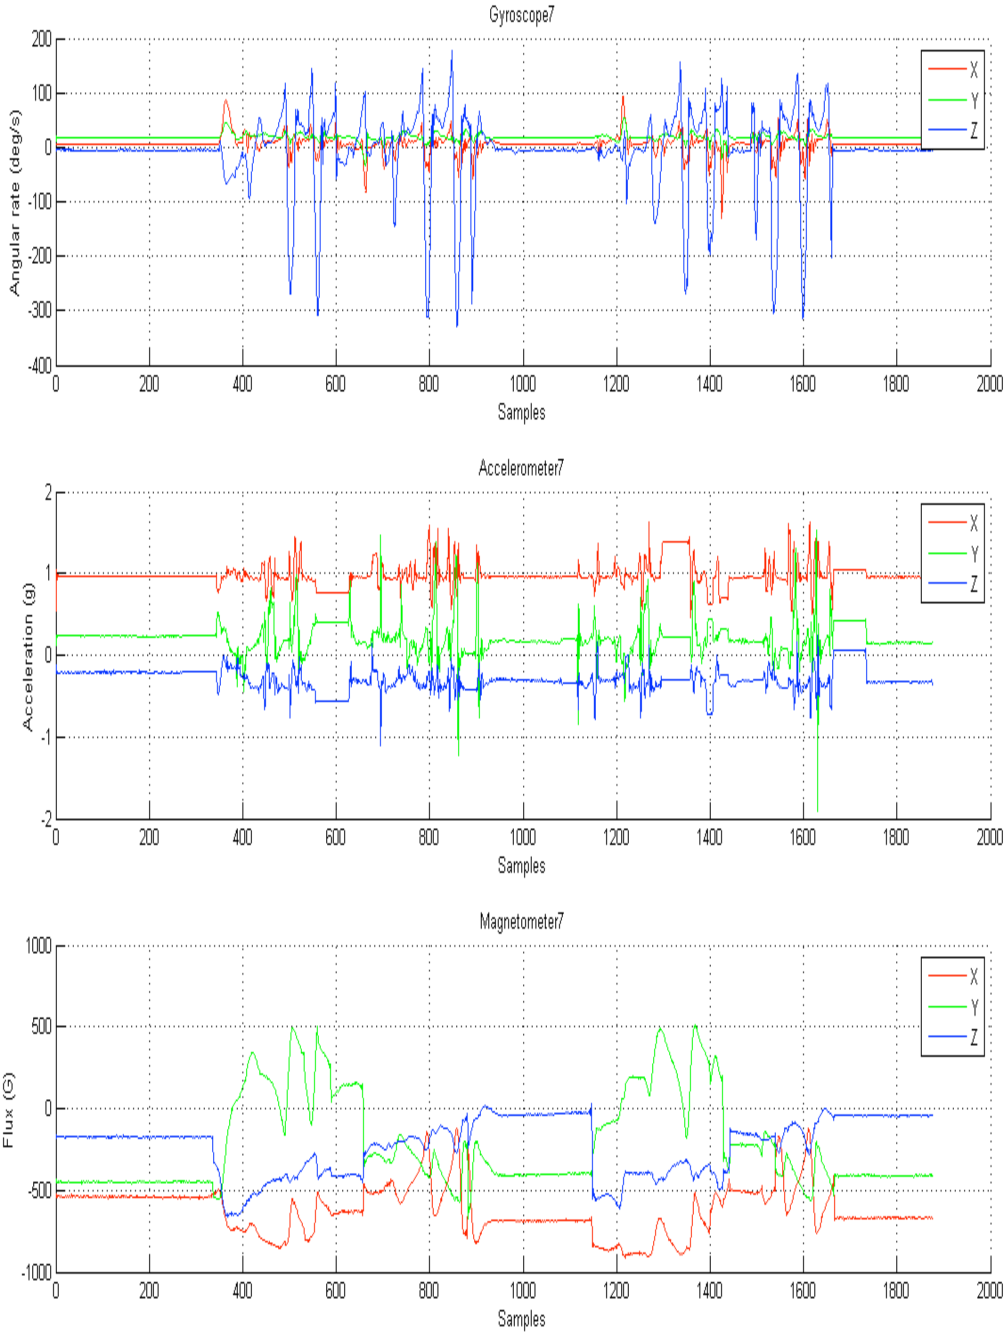

Figure 41: An example of a sample set of initial raw data collected in Thessaloniki (shin segment – subject 6) showing the correspondence between the sensors for different movements.

The deployed system had been fully tested on different OSs, such as Windows 7, Vista, and XP, without that the indicated problem has never been detected. The use of Windows 8 had not been planned.

The technologies involved in the experiment have been tested separately, which involves the absence of a benchmark, for instance, produced by a static camera video, for the following WIMUs data analysis.

Furthermore, some incongruities among the performed activities that have been declared and the gathered information are present. As an example, the transition "Taking a book in the cupboard – Reading in a sitting position" is not visible in the data-gathering, since the 0→1 and 1→0 abrupt changes in the thigh and shank acceleration signals are not shown, which indicates that either the listed activities are different from the ones really performed or the experiment duration has been set incorrectly, so that the data-gathering finished before the end of the test. The lack of any benchmarks cannot provide helpful information as far as that is concerned.

Finally, the set experiment duration (averagely 1 minute) has been too short to extrapolate useful knowledge from the data. Indeed, it is expected that an ambulatory monitoring of physical activity should store data for hours (or days) to obtain some reliable outputs about the patient's health status. It is expected in the pilots that WIMU data will be collected at set points during the day while particular tasks are being undertaken. This will allow for some consistent and continuous data for analysis.

For these reasons, the adoption of the collected data in the developed algorithms testing phase has been particularly limited. Hence, they could not be used to improve the methods' robustness. Since information about the real orientation, body posture and step counting were not evident, only the activity/inactivity algorithms have been tested using the data collected in Greece. Table 5 and Table 6 show the comparison between causal and acausal SHOD method. Results are shown for two subjects and are discriminated between low- and high-dynamic cases (under the supposition that the chest presents a reduced motion compared to the typically high-kinematic model recorded on the shin), presenting, as expected, a mean accuracy increment equal to 2.45%, even though such a conclusion is statistically non-optimal.

Table 5 Results of SHOD method

	Low-Dynamic		High-Dynamic	
SHOD (causal)	Window Size	90	Window Size	15
	Threshold	70.84	Threshold	1027.469
	Accuracy	88.67 %	Accuracy	93.92 %
SHOD (acausal)	Window Size	256	Window Size	38
	Threshold	72.6023	Threshold	947.0984
	Accuracy	93.78 %	Accuracy	95.02 %

Table 6 Results of SHOD method

	Low-Dynamic		High-Dynamic	
SHOD (causal)	Window Size	275	Window Size	185
	Threshold	198.2109	Threshold	1372.7669
	Accuracy	89.06 %	Accuracy	92.46 %
SHOD (acausal)	Window Size	275	Window Size	191
	Threshold	198.2109	Threshold	1408.62
	Accuracy	91.02 %	Accuracy	94.12 %

3.7 Conclusion

The algorithms conceived for processing accelerometer, gyroscope and magnetometer data from the WIMU devices have developed well in accordance to the proposed needs within Dem@Care. The algorithms exploit low-level fusion of sensor data to detect different body postures and different modes of activity and inactivity.

Further work will develop the accuracy of the algorithms, as they will be tested with greater amounts of data collected from more people over longer periods of deployment.

4 Introduction to Complex Event Processing

Early fusion in Dem@care will allow the extraction of conclusions concerning the person's state by mining the physiological and lifestyle data that can be used for the early detection of unusual events or patterns. Complex Event Processing (CEP) technology is proposed as first candidate for merging and extracting information from sensors and raw data.

In this chapter we introduce Complex Event Processing, its concepts, its proposed implementation and its usage. Also we present the proposed implementation of this technology in the project.

4.1 Fusion and CEP

Information fusion is the process of combining information and knowledge from multiple sources representing the same real-world object in order to obtain a consistent, accurate, and useful representation of the situation. Fusion aims to achieve a more complete result than simply the sums of all the inputs. Recent information fusion research is also concerned with the representation of uncertainly and imprecision in the situation representation.

Complex Event Processing [7] is an architectural concept that is part of data management. It groups together the methods to realize event-driven information systems.

4.2 Complex Event Processing

Event processing is a method of tracking and analyzing streams of input information about things that occur (events), in order to derive a better understanding of the situation from them. In Dem@Care, CEP will be used to combine data from multiple sources (sensors, detected activities) in order to infer events or patterns that represent the condition of the PwD, and allow the doctor to make better informed decisions. The CEP is based on the definition of a set of rules which take a set of events as the input and produce a set of complex events at a higher level as output. The clinical partners will need to contribute in the definition of such rules.

A complex event is as an event which can occur only if other observed events occur. It is the abstraction or the aggregation of events i.e. it is generated by the occurrence of several events.

For the implementation, the processing engine for CEP requires three items to work: a set of events, a set of rules and a rule engine [6].

The set of events are stored in a database of facts (as events are not yet processed by the engine, they can be considered as facts). To the arrival of an event in the system, the latter is added to the base of facts. This base must thus be able to contain a huge amount of events and to be updated regularly. The database of facts is also called Working Memory.

The set of rules, just like the events, are stored in a database of rules, also called Knowledge Base because rules represent the intelligence of the system. It is generally created with the initialisation of the system but the majority of implementations are able to modify the rules dynamically. The database of rules can also be represented in the form of a network which allows to identify the links between rules (the output for one is the input of other). This representation also allows can demonstrate the links between events and their causality. Thus the arrival of an event in a rule will trigger the generation of a new event which will then activate other rules.

The inference engine is the core of the system. Indeed, it is about the algorithm making the link between rules and facts. Various algorithms are available, that makes the specific characteristic of the various implementations of CEP. The inference engine is connected with a diary which manages the order according to which are tested and carried out the rules in order to optimise the search of the validated rules.

The implementation of CEP consists of three steps: identification of significant events, analysis of their impact and real-time decision-making. These three parts lead to the definition of an event model, an action model and the linking of the two models. The event model corresponds to the specification of the event type, which is expected by the system. The action model corresponds to the rule definition that explicitly defines the action to perform when the associated pattern is matched.

The pattern is used to detect the fact that an event or a combination of events has occurred.

But patterns can also detect the lack of expected events.

An event can be defined as “something that happens”. Actually there is distinction between the event in the real life and its representation in information system, the event object. The events can usually come from three sources: hardware, software or human. In hardware, an event is the result of a change of state of a process. The state is defined by the value of some parameters or some sensor, when a value exceeds a threshold there is a change of state. An example of such event is a motion detection sensor detecting a person entering the room that the PwD sits. In software, events correspond to messages send by components; for example, the detection of a face to face interaction from the relevant component. An event from a human can be the information resulting from a human observation, such as the doctor or care taker noticing a change of state in PwD and informing the Dem@Care system.

Furthermore, each event has metadata which generally gives the time when the event occurred, what produced it and its relation to other events. For example, when an event generates another event, the second one has the reference of the first one in its metadata, this is a causality relation. But events can also have time relation, hierarchical relation or affiliation relation if they are caused by the same event. To sum up, there are three aspects to an event: its relations with other events, a meaning (that is the description of the activity it represents) and its list of attributes.

One of the advantages of CEP is that it can be represented graphically to allow understating by the decision maker. A graph can be created to represent the notion of the complex events. Another advantage of CEP is that it is easily adaptable to distributed architecture. A CEP application can be easily added to an existing system.

Finally, CEP is proposed as a solution to assist the diagnosis or monitoring of the progress of AD within Dem@Care system. The feasibility of applying CEP will be investigated in task T3.3 as rule-based detection like CEP needs expertise to describe and write the rules that fit the desired results. Therefore, it may be difficult to set rules that are human readable.

4.3 CEP Implementation and Results

The Dem@care system will use the Drools CEP engine. This engine has been encapsulated in a new web-service and integrated to the overall system.. Observations are sent to the CEP Webservice. Drools CEP engine processes these observations and creates new higher level observations. These new observations are added to the CEP knowledge base and may be reused to infer even higher level observations, by other Dem@Care components.

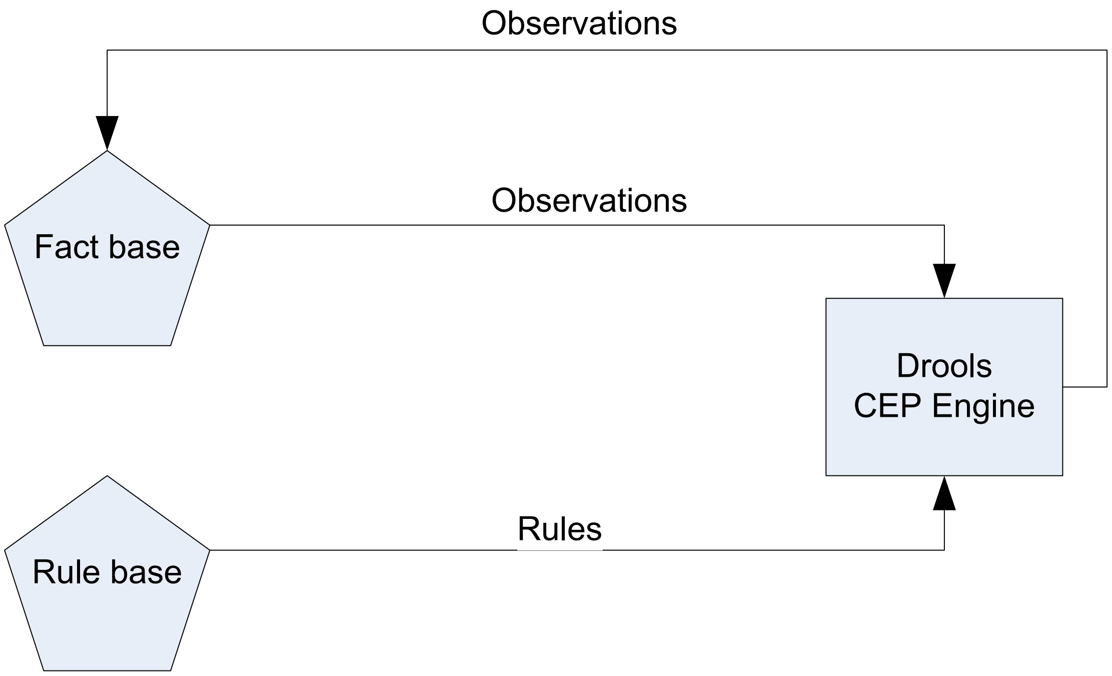

Figure 42: Drools global functioning

The CEP engine depends on a set of rules, and creates new observations based on known facts, as demonstrated in Figure 42. Examples of rules created could include:

- During a time window t_i to t_j, *heart_rate* > x_1 + *moving_speed* > y_1→ *Walking_Exercise*
- During a time window t_k to t_l, *heart_rate* > x_2 + *lying* → *Sleeping*

To visualise the observations created a log system has been developed and integrated in ServiceMix. Figure 43 and Figure 44 show an example of some rules being applied.

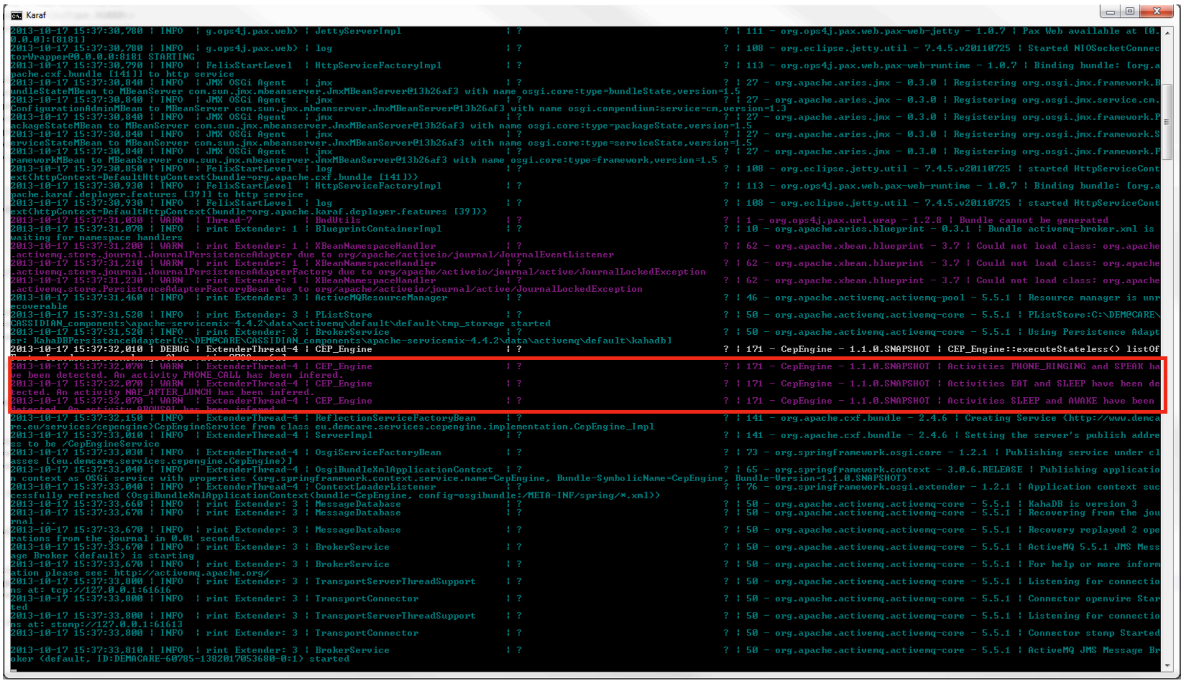

Figure 43 Global Log of Service Mix

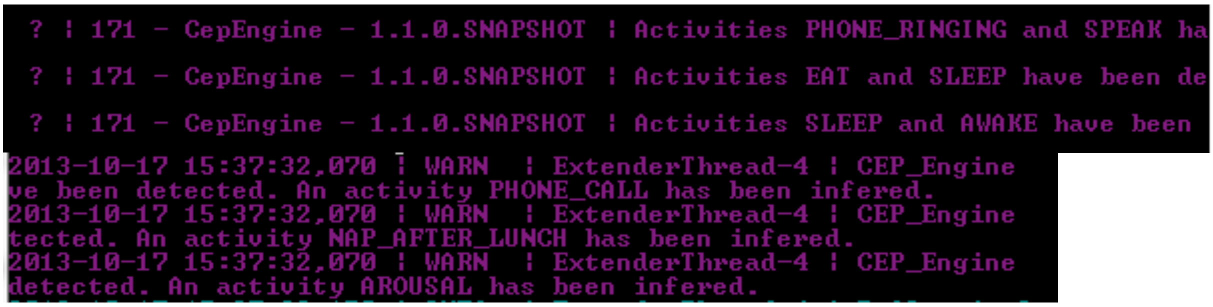

Figure 44 Extension of the ServiceMix Logs showing the applied rules.

Some rules have been created to demonstrate the concept, but to complete the rule base, clinical and technical partners need to work together to list all rules that are needed by the system. Afterwards these rules will be translated in the rule language and added to the rule base.

A graphical interface has also been integrated to make easier rule management. This graphical interface makes it possible to see all rules existing in the system (Figure 45), to edit them (Figure 46) and to manage them (addition, modification). This interface is also helpful to allow not technical people to visualise and understand the rules implemented in the system and provides the flexibility to dynamically modify these rules.

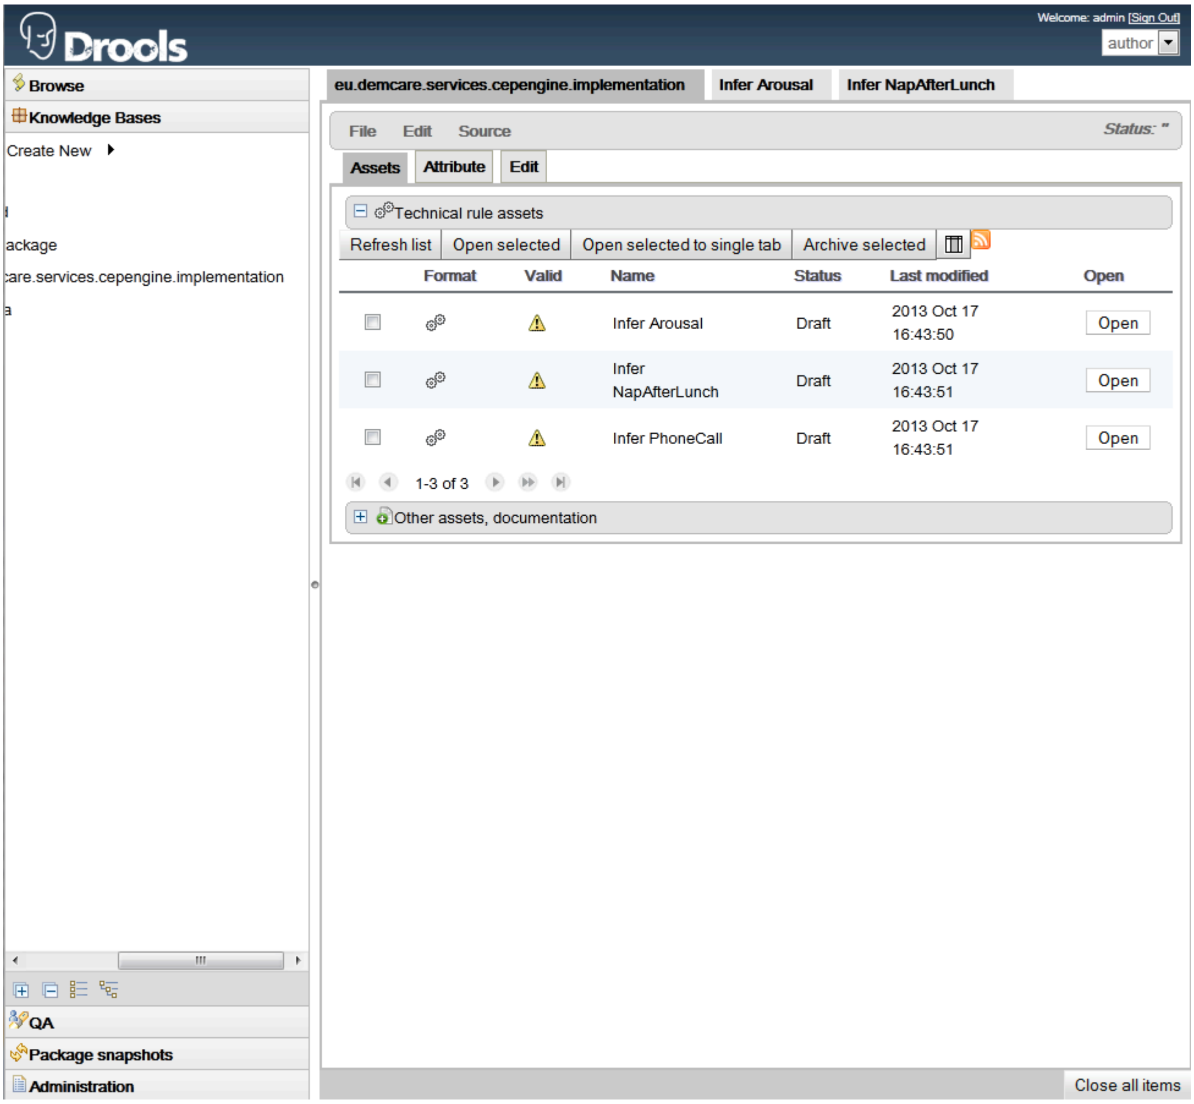

Figure 45 Rules Overview Interface

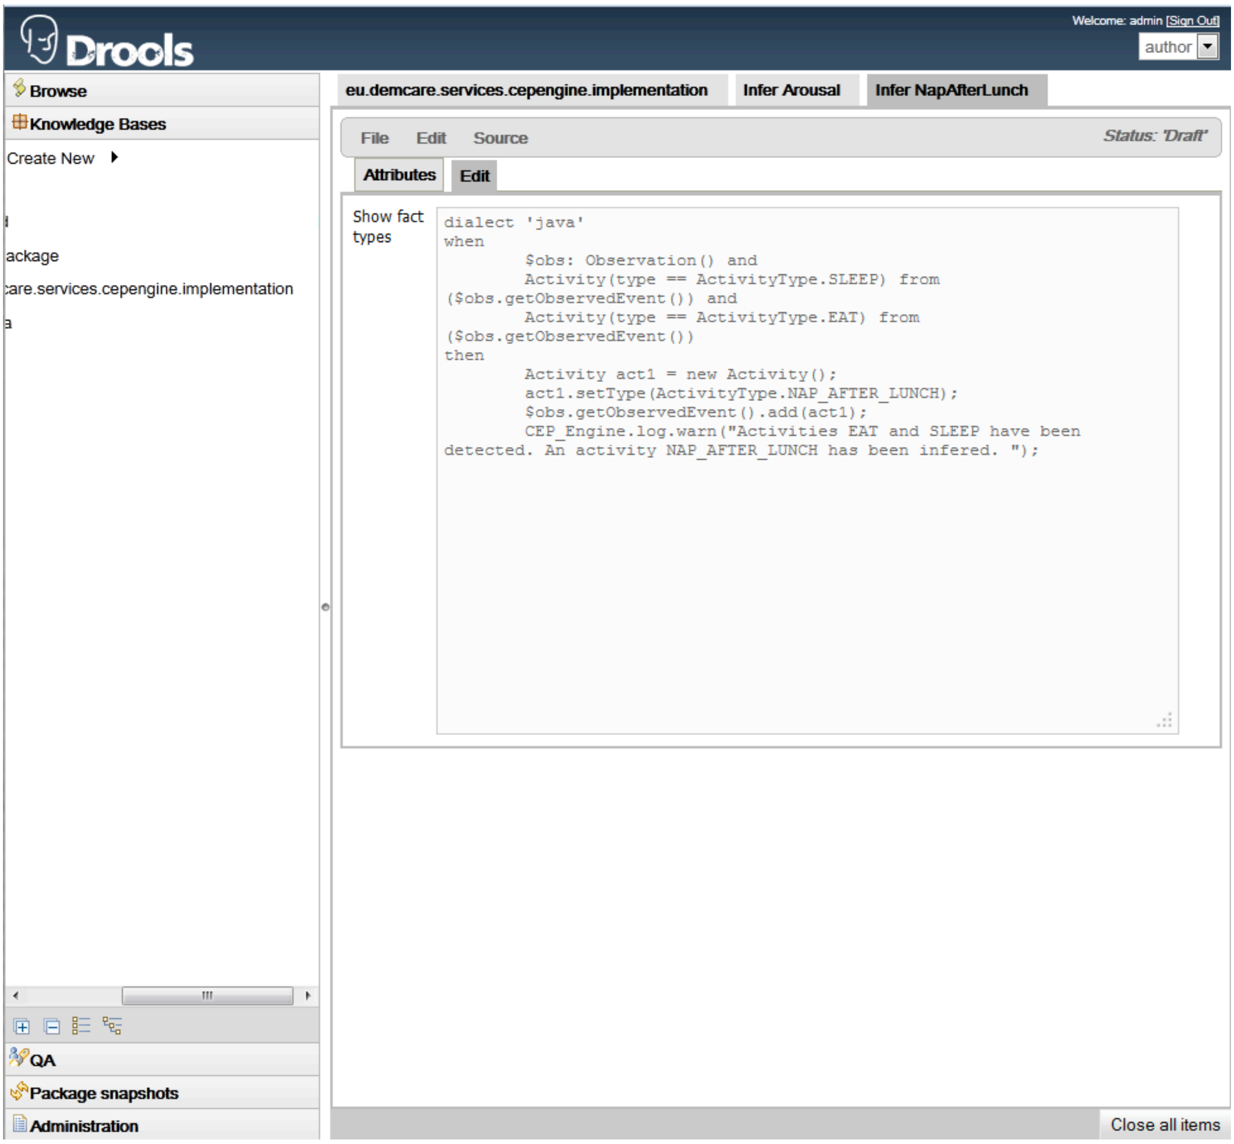

Figure 46 Rule Edition Interface

4.4 Conclusion

The complex event processing technology is an efficient way to make real-time information fusion. Rules are written according to the system's needs and make it possible to infer more high-level information. CEP engine has already been integrated to the system, and it is planned that it is used in the second prototype. The major future work concerns the definition of the rules, with the help of the clinical experts. On top of that, future work may focus on a way to improve the CEP engine to take into account semantic data and uncertainty.

5 Conclusion

This report presented work in progress and the algorithms developed require further refinement and testing.

Fusion of the Philips DTI-2 and the Gear4 Renew Sleepclock allows to measure parameters of PwD that relate to their daily patterns. Apart from thereby adding functionality, this also increases user comfort, as the DTI-2 does not need to be worn at night.

Incorrect detection of sleep/rest segments due to the fact that the device has been turned on, but is not being worn is prevented by fusion of skin conductance, accelerometer and temperature sensors of the DTI-2.

In order for the DTI-2 data to be interpreted correctly it is crucial that the data recorded on the device is completely aligned with the rest of the system. A mechanism to automatically synchronize the DTI-2 data to the camera data has been described.

The WIMU based algorithms created for processing accelerometer, gyroscope and magnetometer data are well in accordance to the proposed needs within Dem@Care. The algorithms exploit low-level fusion of sensor data to detect different body postures and different modes of activity and inactivity. Further work will be needed to improve the accuracy of the algorithms, as they will be tested with greater amounts of data collected from more people over longer periods of deployment.

The complex event processing technology has been developed as an efficient way to make real-time information fusion. Rules are written according to the system's needs and make it possible to infer more high-level information. Future work may focus on a way to improve the CEP engine to take into account semantic data and uncertainty.

Based on the feedback of the first pilot, as well as new requirements set by the clinical partners, parts of the technology described above will be further matured and integrated into the overall systems, by means of new, updated or fusion components.

References

[1] Irving, K., O'Mahoney D., Journier V., Sävenstedt S., "*D2.2 Functional Requirements and Clinical Scenarios v1*", Dementia Ambient Care: Multi-Sensing Monitoring for Intelligent Remote Management and Decision Support, Dem@Care – FP7-288199.

[2] Schuijers, E., Newman E., Contat M., "*D3.1 Physiological & Lifestyle Monitoring v1*", Dementia Ambient Care: Multi-Sensing Monitoring for Intelligent Remote Management and Decision Support, Dem@Care – FP7-288199.

[3] Eus J.W. Van Someren, Annemarieke Kessler, Majid Mirmiran, and Dick F. Swaab, "*Indirect Bright Light Improves Circadian Rest-Activity Rhythm Disturbances in Demented Patients*", BIOL PSYCHIATRY 1997;41:955—963.

[4] Van Someren E J, "*Actigraphic monitoring of movement and rest-activity rhythms in aging, Alzheimer's disease, and Parkinson's disease*", IEEE transactions on rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society, Vol. 5, No. 4, pp. 394-8, December 1997.

[5] Bouten, C. V.; Westerterp K.R.; Verduin, M.; Janssen J.D., "*Assessment of energy expenditure for physical activity using a triaxial accelerometer*", Medicine and Science in Sports and Exercise, Vol. 26, No. 12, pp. 1516-1523, 1994.

[6] David Luckham, The Power of Events, Addison-Wesley Professional; 1st edition ,May 18, 2002.

[7] http://www.thetibcoblog.com/2007/04/23/what-is-complex-event-processing-part-1/. (Accessed 10/09/2012)