

An Integrated Multi-Sensing Framework for Pervasive Healthcare Monitoring

M. ElHelw*, J. Pansiot, D. McIlwraith, R. Ali, B. Lo, L. Atallah
Department of Computing, Imperial College,
London, UK
{me, jpansiot, dm05, smrali, benlo, latallah}@doc.ic.ac.uk

Abstract — Pervasive healthcare provides an effective solution for monitoring the wellbeing of elderly, quantifying post-operative patient recovery and monitoring the progression of neurodegenerative diseases such as Parkinson's. However, developing functional pervasive systems is a complex task that entails the creation of appropriate sensing platforms, integration of versatile technologies for data stream management and development of elaborate data analysis techniques. This paper describes a complete and an integrated multi-sensing framework, with which the sensing platforms, data fusion and analysis algorithms, and software architecture suitable for pervasive healthcare applications are presented. The potential value of the proposed framework for pervasive patient monitoring is demonstrated and initial results obtained from our current research experiences are described.

Keywords-component: *pervasive sensing, body sensor networks, vision sensor networks, activity recognition, data fusion, behaviour profiling, data visualisation*

I. INTRODUCTION

Recent advances in sensing technologies offer significant potential for pervasive healthcare delivery in terms of versatility and cost-effectiveness. For instance, wearable Body Sensor Networks [1] are used for measuring postoperative recovery for patients after elective minimal invasive surgery where early signs of post-operative complication can be captured in a home environment because patients are usually discharged much more quickly compared to conventional surgery [2]. Pervasive sensing also provides an appealing approach for monitoring the wellbeing of the elderly, where the general trend of an increasingly ageing population has placed significant burdens on the current healthcare systems. To this end, smart homes employing different sensors can monitor patient interaction with the surrounding environment [3]. Such pervasive healthcare environments support elderly independent living, encourage the maintenance of physical fitness and enable observing social activity while alleviating workload of healthcare professionals. They also aid identifying transient behaviour abnormalities that may indicate adverse events.

In practice, pervasive systems are only realizable by integrating multiple sensing modalities and existing research has shown that there is a complementary relationship between wearable and ambient, *i.e.* background, sensing paradigms.

Wearable sensing enables continuous monitoring of patient motion and physiological parameters through a network of body-worn sensors wirelessly linked to each other. They are usually based on accelerometers [4, 5], pulse oximeters (SpO₂) [5], ECG [6] and temperature sensors [4]. However, wearable sensors only provide limited body information and due to the lack of global reference, it can be difficult to use this information to deduce the context of the activities. For example, an accelerometer can detect local motion, such as sitting down or walking about; however, it cannot tell if the subject is sitting on a chair or lying on the ground. It is also difficult to differentiate between low-impact sedate activities performed within the home due to the complexity in separating fine motion from the movement of the body as a whole. It is only by the use of ambient sensing that other essential information such as body posture can be obtained with effective discrimination.

Ambient sensing employs a large number of sensors that are ubiquitously placed in the environment such as video cameras [7, 8], infrared sensors, water flow and utility usage sensors, and pressure sensors mounted on furniture [9]. These systems can provide information about the location and activities of the subject within the environment and enable the detection of critical events such as falls. Nevertheless, pervasive systems based only on ambient sensors require the use of a large number of sensors, which involve complex and expensive deployments. They also suffer from other limitations such as the difficulty of deducing detailed changes in motion patterns and inability to detect vital signs related to the onset or progression of chronic disorders. Furthermore, the need for having large training data to be used for inferring activities hinders the practical use of ambient-only sensing frameworks.

By integrating the strengths of ambient and wearable sensing, it is possible to provide true pervasive systems that can be used to accurately infer subject condition based on activity and physiological parameters. In general, sensory data can be fused at signal, data, feature, or decision levels [10]. For instance, sensor signals can be combined by using simple hardware thresholds [5], whereas at the data level, pattern recognition methods such as Bayesian Networks [7], Hidden Markov Models (HMMs) [11] and Gaussian Mixture Models (GMMs) [12] are often used for data fusion and analysis. Furthermore, due to the large volume of sensing data, dimensionality reduction techniques such as Manifold Embedding [13], Principal Component Analysis (PCA) and

*Currently with the Centre for Informatics Sciences (CIS), Nile University, Egypt

feature selection [1] are often applied prior to applying actual activity classification procedures.

Another major challenge associated with pervasive system deployment is the integration and interoperability of diverse set of components including sensors, middleware, web services, databases as well as backend data mining and visualization tools for different user groups and with varying levels of security expectations. Additionally, the need to collect and operate on continuous sensor data streams introduces significant computational and storage loads that are exacerbated in case of a large number of users. Therefore, the development of efficient and scalable stream processing and management architecture is essential for practical pervasive system deployments. A number of light-weight software architectures for scalable data processing, transmission and storage have been recently introduced based on techniques such as wavelets, histograms, sketches, sub-sampling and synopsis data structures [14, 15]. These techniques optimize resource utilization and reduce memory usage, lower database access rate and enhance responsiveness for web clients.

The main contribution of this paper is the provision of a complete and integrated framework for the development of pervasive systems used in intelligent healthcare delivery and patient monitoring. We describe individual framework components that correspond to different phases from data acquisition platforms, stream management to fusion and analysis. We also present results of research experiments carried out while developing the proposed framework.

II. SYSTEM ARCHITECTURE

The software architecture for the proposed framework is based on the push style [16] message broker model [17] which is used by many established industrial systems such as J2EE [18] and several research frameworks [19]. Fig. 1 illustrates a schematic diagram outlining framework phases and data workflow. Patient activity information capture by wearable sensors, as well as data captured by ambient sensors, is streamed through authorised gateway devices to broker server(s) for logging and interface to databases. Further processing, data fusion and analysis for patient activity classification and behaviour profiling are carried out on dedicated servers or distributed among clusters. Brokers work as intermediary connectors that facilitate communication between heterogeneous and distributed system components for data acquisition, processing, storage and visualization.

Based on an asynchronous messaging paradigm, system components can publish data under a topic and/or subscribe to a particular topic or a category of topics. The broker delivers data to registered subscribers. It also enables for querying data-generating components that provide metadata in schema-conformant XML. User interaction with the system is implemented through web services and different stakeholders connect to the system in order to retrieve stored information or carry out specific data processing algorithms. The loose coupling of heterogeneous framework components allows for flexible and scalable pervasive system architecture that facilitate monitoring and knowledge discovery. Moreover, the use of secure connectivity and data abstractions is essential for protecting patient privacy.

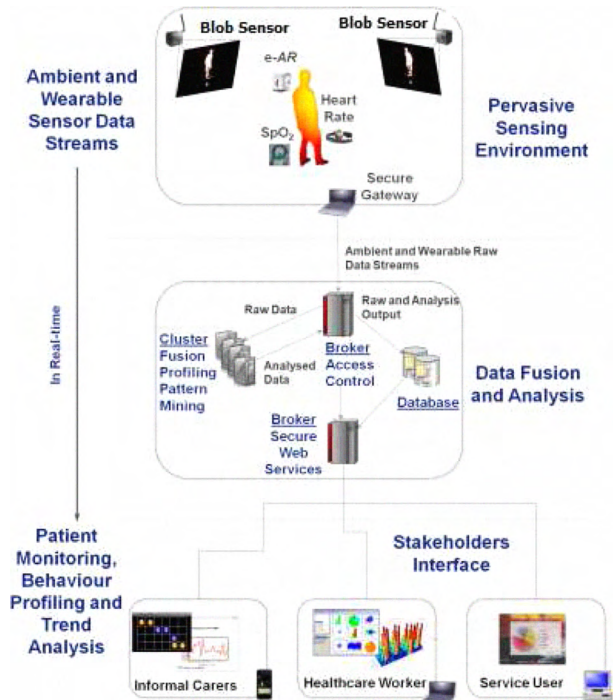


Figure 1. System architecture for pervasive systems where loose coupling between heterogeneous components ensures efficiency and scalability.

III. AMBIENT AND WEARABLE SENSING PLATFORMS

A. Ambient Visual Sensors

The deployment of ambient visual sensors, also referred to as “smart cameras”, is based upon a collaborative network of partly self-sufficient vision-based sensor nodes. The main idea behind ambient sensing is to distribute data processing on-board in order to reduce communications bandwidth and power consumption. For example, the Matrix Vision mvBlueLYNX smart camera has been demonstrated in the *SmartClassySurv* activity recognition framework [20]. This device embeds an imaging sensor, a FPGA for pre-processing, a PowerPC processor, an Ethernet interface for communication with the rest of the network, and is powered through power over Ethernet (PoE). In general, visual sensing nodes are composed of the following:

- Imaging sensor - represents the first phase in the overall image processing pipeline. The sensor is typically based on low-power CMOS technology and varies along the dimensions of captured image resolution, color depth, energy consumption and frame rate.
- Main processing board – deals with basic image processing tasks such as noise reduction and target segmentation. Processing boards can be implemented using general-purpose central processing units (CPUs), digital signal processors (DSPs) and/or field-programmable gate arrays (FPGAs).

- Communication module – for convenient deployment in healthcare environments, this module is typically based on wireless LAN transceivers such as the 802.11g/b Wi-Fi and 802.15.4 (Zigbee) with different channel capacity and energy consumption.

Fig. 2 shows a view of the ubiquitous sensing node developed at Imperial College London [8]. The node is composed of a VCSBC50 DSP camera and main processing board embedding a Texas Instrument MSP430 processor, 32KB RAM, 60KB Flash and 2.4GHz wireless communications (250Kb/s). It is packed in a slick casing for better environmental integration.



Figure 2. Blob sensor node.

Video data observed by the ambient sensors is processed on board in real-time and the sensor transmit only the derived signal metrics such as the silhouette, *i.e.* blob, of a moving object and its local motion in the form of optical flow. Through a network of ambient vision sensors, accurate tracking can be obtained which can improve the overall system robustness in behaviour profiling. Under this sensing paradigm, the binary blobs of monitored subject are first extracted from the video signal using a statistical background model, where every pixel is represented as a Gaussian mixture distribution maintained over time [21]. Incoming signals are compared with the existing background model, obtained by using a single Gaussian for the background colour model, and segmented into a binary map of foreground and background. Using two or more Gaussians in a Gaussian Mixture Model (GMM) requires the use of floating point operations to fuse them, which is not supported in most hardware and would be too computationally expensive.

Further optimisations are necessary to reach real-time performance. For example, the background model can only be updated every 20 frames. A fast and low-resolution foreground segmentation is performed beforehand to determine the region of interest (ROI). The full resolution segmentation and morphological filtering (erosion-dilation) are consequently performed only in the ROI bounding box, providing a substantial speed gain.

A couple of basic features can be computed on node from the binary blob image, such as the object centre, axis-aligned bounding box (AABB), and the eigenvectors of the object providing information on its orientation. In many situations, these features provide sufficient information for activity detection and classification. In such cases, only the features are transmitted through the network, thus reducing bandwidth usage and energy expenditure. It is worth noting that the blobs extracted by each ambient sensor do not carry any appearance information and no image data is transmitted to other devices. This is important for home care environments where privacy is of high priority. Fig. 3 illustrates example captured scene images and computed object blobs.

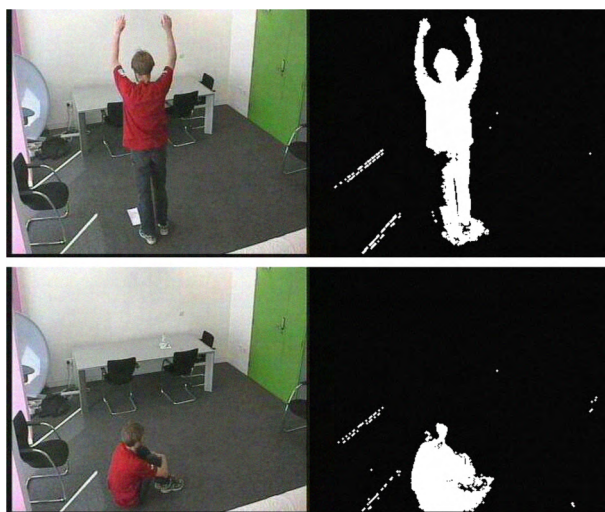


Figure 3. Example scene images (left) and extracted blobs (right).

B. Ear Worn Activity Recognition (*e-AR*) Sensor

A number of patient activities and physiological parameters can be acquired by using wearable sensors. These are small devices worn by the user and mainly composed of a miniaturised processing board, communication module and sensors for measuring body motion as well as other vital signs such as temperature, pulse and blood oxygen level. For the proposed framework, we have used the ear-worn Activity Recognition (*e-AR*) sensor [5] which is based on the Body Sensor Network (BSN) platform [22]. A basic BSN node comprises a Texas Instruments MSP430 16-bit ultra-low-power RISC processor, a Chipcon CC2420 radio transceiver, MCC ChipOX SpO₂ module, a temperature sensor and a 3-axis accelerometer. The BSN node runs TinyOS [23], which is a small, open-source and energy-efficient sensor board operating system. Fig. 4 shows the BSN node on a circuit board and *e-AR* sensor. During its use, the *e-AR* sensor periodically samples from the accelerometer and other sensors, and transmits this data to a nearby gateway station. Motion features such as head-tilt, mean, median and variance extracted from a limited time window (usually selected to be 1-2 seconds), are used to compute patient activity levels, *i.e.* activity indices. Alternatively, activity indices can be computed on BSN nodes in order to minimise transmission bandwidth, extend battery life and facilitate system scalability.




Figure 4. The e-AR sensor (a) with SpO₂ reader cord that clips to the ear (b) and Body Sensor Network (BSN) node on a circuit board (c).

IV. DATA ANALYSIS

Data analysis is an important part of pervasive sensing systems. It involves inferencing and learning for behaviour monitoring, context aware sensing, designing efficient databases that encompass different types of sensor data while optimising retrieval queries, and correcting for errors in data resulting from noise, interference and missing connections. Other important analysis applications include sensor fusion where different types of sensors are used to obtain better understanding of activity and pattern mining where patient activity is used to identify trends that occur over long period of time. The following subsections describe a number of data analysis techniques and their use in pervasive patient monitoring. Results obtained using these techniques within the proposed framework are presented in the next section.

A. Gaussian Mixture Models and Probabilistic Decision Level Sensor Data Fusion

Through the fusion of ambient and wearable sensor information, it is possible to achieve increased accuracy in activity inference, enhanced tolerance to sensor failures through the inclusion of both complementary and redundant data types, thus offering improved home monitoring systems with extended coverage of both local and global characteristics of activity. For example, most features derived from the ambient visual sensor depend on the relative position of the subject and the sensor, but this dependency can be alleviated by the wearable sensor remaining fixed in the subject's frame of reference. One approach for sensor fusion is by using a Gaussian Mixture Model (GMM) [12] which uses a Gaussian Bayes Expectation Maximisation (EM) classifier based on features extracted from wearable and ambient sensors. For the e-AR, these features include tilt and movement frequency spectrum whereas for the ambient sensor they include the aspect ratio and mean velocity of the object. The classification considers that different activities correspond to different

classes. It uses an EM iterative method to compute the maximum likelihood fitting [24] based on the assumption that the conditional probability density function for each of these classes is Gaussian. The expectation and maximization steps are performed iteratively until convergence and part of input data is used to evaluate the accuracy of the classifier based on the marginal probability of every activity with the highest probability chosen for the final classification.

For multi-dwelling housing environments, it is important to automatically identify related data streams before fusion can occur. This process is called sensor correlation. Probabilistic decision level fusion [25] is used for sensor correlation. In this case, feature sets are extracted such that they exhibit classification redundancy across sensing modalities. Consequently, activities can be detected accurately and independently by per-sensor classifiers at the same time. Feature set search space is firstly reduced by applying BFFS [1] to rank sensor-specific features which yield high accuracy when used singularly and independently. Subsequently, sets with the highest redundancy across sensors are automatically selected by using the multi-objective accuracy/redundancy score [26] to rank feature sets classification. With the appropriate e-AR and visual feature sets selected, they are exposed to pre-trained per-sensor activity classifiers. The results of these classifiers are then used to select those data streams that most likely result from the same subject. This technique enables automatic combination of ambient and wearable sensing data which improves overall activity classification for home healthcare monitoring.

B. Behaviour Profiling Using Hidden Markov Models (HMMs)

Observing patient activity and movement patterns over an extended period of time can be cumbersome especially for large number of subjects. A similarity based HMM technique [27] for clustering of location sequences, *i.e.* patient movement, can be used for representing the behaviour pattern of the patient and its temporal variation without explicitly defining activities, hence alleviating privacy concerns. Standard approaches to clustering with HMM comprise model training with a sequence, then using pair-wise distance based methods to perform the clustering. The proposed similarity based clustering approach uses a feature space that is generated using HMMs to express the similarity of sequences to each other. For this purpose, the features describing a sequence are calculated as similarity measures between that sequence and other reference sequences that are selected from the whole set as chosen by experts, or the whole dataset can be used. Behaviour profiling is achieved by observing the clustering of sequences in the new feature space. Assuming that certain clusters of behaviour sequences represent normal patient activity over a period of time, outlier can indicate deviations from normal behaviour patterns. If the deviation is large, additional data analysis can be performed to elucidate the potential causes of this deviation.

C. Pattern Mining for Routine Behaviour Discovery

Complementary to a model based approach for behaviour profiling is to discover patterns of activities using pattern mining algorithms. In [28] an algorithm is proposed for

constructing a compressed data structure specifically for describing routines by mining activity data patterns. Typically, when a user is wearing an *e-AR* sensor, an activity level is streamed periodically from the sensor. The activity level is the output of a classifier described in [29], and can take one of four values where the lowest level indicates almost no activity (during sleeping or sitting) and the highest level indicates a high-intensity activity such as running. While some activities may be described by a single activity level, most activities result in a sequence of activity levels. These combinations are discovered by the proposed pattern mining algorithm. The algorithm in [28] obtains a data structure called the routine tree, a picture of the user's routine showing patterns of activity at progressively finer time resolutions wherever there is more detail to uncover. Each node in the tree represents a time interval, for which frequent patterns are stored. A systematic, top-down tree construction method mines data at smaller durations, while avoiding mining in further detail where there is little structure to discover. The tree is then pruned by merging adjacent time intervals that have the same maximal frequency pattern to produce a more compact representation. Generated activity tree provides a representation of the composition of a user's routine which when combined with appropriate visualisation techniques, intuitive graphical views of behaviour patterns can be attained.

V. RESULTS

The sensing, management and analysis techniques described above were applied in patient monitoring experiments. For our studies we have designed a home healthcare laboratory where subjects can be observed performing a variety of tasks typically enacted within a living space. This laboratory has been augmented with support for both wearable and ambient visual sensors. By using the *e-AR* and visual sensors together we facilitate the extraction of local motion characteristics such as head tilt, sway as well as capturing overall pose and limb motion.

An initial study measuring the impact of fusing wearable and ambient sensing on the accuracy of activity classification has been carried out [12]. A set of 9 activities, shown in Table 1, were considered along with their recognition rates as illustrated in Fig. 5. The activities were carried out in sequence by 2 actors with each activity performed during approximately one minute. For all activity classes, a marked classification improvement can be observed with the exception of sitting which displays a minor decrease. Performance gains are particularly noticeable for classes where *e-AR* sensor data is ambiguous due to a lack of global information, as in classes 2, 3, 5 and 6. What's shown here is an increase in sensitivity of the aforementioned classes due to more discriminative data. Analysis of classification using the *e-AR* sensor alone would reveal that this increase in sensitivity is due to the correct classification of data items previously, and incorrectly, classified as 'Sitting' – thus whilst sitting has demonstrated a minor decrease to sensitivity, the specificity of this class has increased considerably. This validates our theory that extending sensor coverage to both local and global characteristics will result in accurate activity classification. However, because there is no system composed of a video-

based sensor and a single wearable inertial sensor is available to the authors' knowledge, it is not possible to compare these results.

TABLE I. THE SET OF 9 ACTIVITIES USED IN [12]

Activity	Class
walking	1
standing	2
standing (head tilted)	3
Sitting	4
Reading	5
Eating	6
Sitting (Sofa)	7
Lounging	8
Lying Down	9

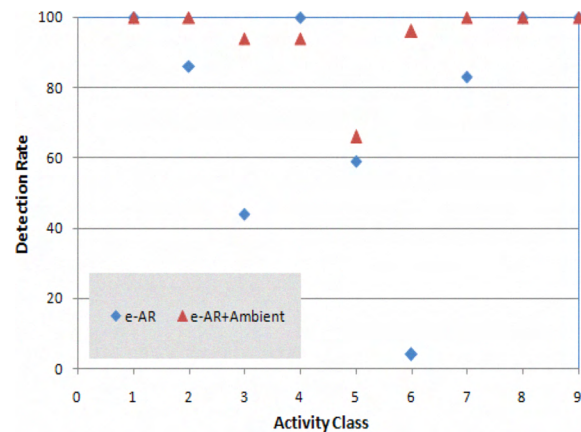


Figure 5. Activity recognition rates for 9 activity classes using the *e-AR* sensor alone and the *e-AR* sensor + visual sensor

In order to ensure the home monitoring systems can be readily adopted, there is a requirement that they should address multiple occupancy scenarios with minimal user interaction. Given a choice of objects visually extracted from the environment along with a set of *e-AR* signals, a probabilistic decision level fusion technique [25] can automatically match the correct *e-AR* sensor to the visually object extracted from its wearer. This has important consequences for environments such as care-homes, where many subjects may be under observation, requiring the correct signals to be matched before fusion can occur.

Fig. 6 shows the results obtained when the probabilistic decision level fusion technique is applied to a scenario with three people in a room. For this experiment, one real subject, wearing an *e-AR* sensor and monitored by a visual sensor, is used along with two streams simulating random blob information for two other occupants. The scenario consists of

one minute of different activities including: sitting, reading, eating, standing, standing with head tilted, walking, sitting on sofa, slouching and lying.

The three traces in Fig. 6 correspond to the percentage of time the e-AR/blob pairs were matched during the experiment. It can be seen that the proposed technique achieves a high degree of correlation accuracy matching e-AR/visual information for the real subject (blue trace) with low matching percentages for the randomly simulated blob information (yellow and purple traces). The activities returned from the blob and e-AR classifiers for the real subject are also presented at the top. It should be noted, however, that the technique is sensitive to the behaviour of the "erroneous" people in the room, *i.e.* people doing the same thing.

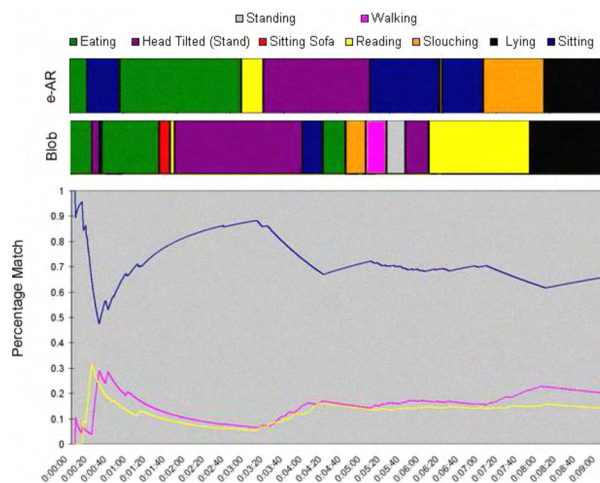


Figure 6. Traces correspond to percentage match of e-AR and visual sensors for same subject (blue) and simulated subjects (purple and yellow). The activities returned from the visual and e-AR classifiers for the real subject are also presented at the top.

Location sequences from ambient sensors can be used for modelling different behaviour patterns. It has been demonstrated [27] that with the use of wearable and vision based ambient sensors, it is possible to develop a practical visualization framework allowing the observation of daily activities in a homecare environment. To this end, an effective behaviour modelling method based on Hidden Markov Models is used for highlighting changes in activity patterns. Representation of sequences in a similarity space allows for clustering, detection of abnormalities and data-exploration. In this similarity space, it is possible to observe how close similar patterns are, and observe patients who change their behaviour. Fig. 7 illustrates the behaviour patterns of a number of individuals where each circle represents the location sequence of a person moving between several rooms. Principal Component Analysis (PCA) was applied to reduce the location sequences matrix to 2D space defined by the first and second PCA vectors. This information can be used to judge similarity between un-labelled sequences and outlier can indicate deviation from normal behaviour patterns.

The pattern mining technique [28] described in previous section was applied to two datasets. The datasets represent simulated activity levels for two types of users: office going and retired. Each set is compiled by using certain known activities for each type of user and a single activity can comprise a combination of one or more activity levels. Data for each activity was collected using an e-AR sensor under lab settings. This was then concatenated based on a simulated model for the daily routines of the two types of users. Fig. 8 demonstrates a visualisation of the obtained results when applying the proposed pattern mining technique to the simulated data. The structure of each day can be clearly seen and differences between users' daily routines in terms of carried out activities visualised. The 'Week' tree is a result of combining five daily trees, which captures the general picture of the simulated routine for one week. Long-term routine trends can also be captured by looking at certain weekdays over extended periods.

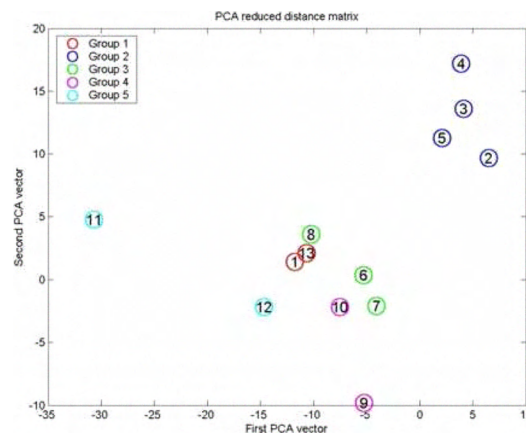


Figure 7. Each sequence of locations is represented as a circle in 2D space obtained by using Principal Component Analysis (PCA) to reduce the location sequences matrix to 2 dimensions [27]. The graph can be used to judge similarity between un-labelled sequences and outlier can indicate deviation from normal behaviour patterns.



Figure 8. Routine trees showing the behaviour changes in terms of carried out activities during the week for two types of users with different activity patterns. The bottom graph shows different activities where an activity may comprise one or more activity levels.

VI. CONCLUSIONS AND FUTURE WORK

The development of pervasive patient monitoring systems is an intricate job requiring the construction of appropriate sensing platforms, the integration of multiple hardware and software components and the utilisation of sophisticated data analysis algorithms. In this paper, we have presented a complete framework for pervasive patient monitoring applications. Practical wearable and privacy-preserving ambient sensing paradigms were described as well as scalable software architecture needed for sensor stream management and processing. Several potential data fusion and analysis techniques for pervasive monitoring were also presented along with some of the results obtained by applying these techniques. The framework comprises full data life cycle from acquisition to management to analysis in order to enable practical ubiquitous patient observation. However, it is worth noting that the framework is not fully automated, something we are planning to accomplish in the near future. The system can only identify one activity at a given time, and further hierarchical classification of multiple simultaneous activities is under consideration. Other future work areas include the development of algorithms for extracting and combining visual information in order to enhance ambient sensing in cluttered scenes and reduce dependency of acquired visual data on sensors orientation. Additional work is also required to develop improved data fusion techniques for more accurate activity detection and introduce robust autonomic behaviour profiling methods that minimise caretakers' workload and enable long term elderly monitoring. At the end, it is worth noting that the framework described in the paper has been developed as a part of the Smart and Aware Pervasive Healthcare Environments (SAPHE) project¹.

REFERENCES

- [1] G.Z. Yang, *Body Sensor Networks*, London: Springer-Verlag, 2006.
- [2] O. Aziz, L. Atallah, B. Lo, M. Elhelw, L. Wang, G. Z. Yang and A. Darzi, "A Pervasive body sensor network for measuring postoperative recovery at home," *Surg. Innov.*, vol. 14, no. 2, pp. 83-90, 2007.
- [3] E. Dishman, "Inventing wellness systems for aging in place," *IEEE Comp.*, vol. 37, no. 5, pp. 34-41, 2004.
- [4] K. V. Laerhoven, H. W. Gellersen, and Y. G. Malliaris, "Long term activity monitoring with a wearable sensor node," International Workshop on Wearable and Implantable Body Sensor Networks, 2006.
- [5] O. Aziz, B. Lo, R. King, A. Darzi and G.Z. Yang, "Pervasive body sensor network: an approach to monitoring the post-operative surgical patient," International Workshop on Wearable and Implantable Body Sensor Networks, 2006.
- [6] C. Park, H. Pai, Chou, Y. Bai, R. Matthews and A. Hibbs, "An ultra-wearable, wireless, low power ECG monitoring system," *IEEE BioCAS*. 2006.
- [7] D. Chen, R. Malkin, and J. Yang, "Multimodal detection of human interaction events in a nursing home environment," International Conference on Multimodal interfaces, 2004.
- [8] B. P. L. Lo, J. L. Wang, and G. Z. Yang, "From imaging networks to behavior profiling: Ubiquitous Sensing for Managed Homecare of the Elderly," International Conference on Pervasive Computing, 2005.
- [9] A. A. Reeves, J.W.P. Ng, S.J. Brown and N. M. Barnes, "Remotely supporting care provision for older adults," International Workshop on Wearable and Implantable Body Sensor Networks, 2006.
- [10] C. McCullough, B. V. Dasarathy and P. C. Lindberg, "Multi-level sensor fusion for improved target discrimination," *Proceedings of the 35th IEEE Decision and Control*, 1996.
- [11] L. R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Readings in speech recognition*, pp. 267-296: Morgan Kaufmann Publishers Inc., 1990.
- [12] J. Pansiot, D. Stoyanov, D. McIlwraith, B. Lo and G.Z. Yang, "Ambient and wearable sensor fusion for activity recognition in healthcare monitoring systems," *Proceedings of the 4th Int. Workshop on Wearable and Implantable Body Sensor Networks*, 2007.
- [13] A. Agovic, A. Banerjee, A. Ganguly and V. Protopopescu, "Anomaly detection in transportation corridors using manifold embedding," 1st International Workshop on Knowledge Discovery from Sensor Data (Sensor-KDD), 2007.
- [14] M. Modahl, I. Bagrak, M. Wolenetz, P. Hutto and U. Ramachandran, "Mediabroker: an architecture for pervasive computing," 2nd IEEE International Conference on Pervasive Computing and Communications, 2004.
- [15] Agarwal, C. and P. Yu, eds., *Data Streams: Models and Algorithms*, Kluwer Academic Publishers, 2007.
- [16] M. Hauswirth, and M.Jazayeri, "A component and communication model for push systems." 7th European Software Engineering Conference with 7th ACM SIGSOFT Symposium on the Foundations of Software Engineering, 1999.
- [17] F. Buschmann, R. Meunier, H. Rohnert, P. Sommerlad and M. Stal, *Pattern-Oriented Software Architecture*, vol. 1, John Wiley & Sons Ltd, 1996.
- [18] P. J. Perrone, V. S. R. R. Chaganti and T. Schwenk, *J2EE Developer's Handbook*, Prentice Hall, 2003.
- [19] C. Giannella, J. Han, J. Pei, X. Yan and P.S. Yu, *Mining Frequent Patterns in Data Streams at Multiple Time Granularities*, MIT Press, 2004.
- [20] S. Fleck, R. Loy, C. Vollrath, F. Walter and W. Strasser, "A smart camera network for distributed tracking and activity recognition and its application to assisted living." In First ACM/IEEE International Conference on Distributed Smart Cameras (ICDSC 07) 2007.
- [21] D. S. Lee, "Effective gaussian mixture learning for video background subtraction," *IEEE Trans. on Patt. Ana. and Mach. Intell.*, vol. 27, no. 5, pp. 827-832, 2005.
- [22] B. Lo, and G. Z. Yang, "Key technical challenges and current implementations of body sensor networks." *Proceedings of the 2nd International Workshop on Body Sensor Networks*, 2005.
- [23] TinyOS. 2008; www.tinyos.net/
- [24] A. P. Dempster, N. M. Laird and R. D. B, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society. Series B (Methodological)* Royal Statistical Society, 1977.
- [25] D. McIlwraith, J. Pansiot, S. Thiemjarus, B. Lo and G.Z. Yang, "Probabilistic decision level fusion for real-time correlation of ambient and wearable sensors," In *Proceedings of 5th International Workshop on Wearable and Implantable Body Sensor Networks*, 2008.
- [26] S. Thiemjarus, *A Framework for Contextual Data Fusion in Body Sensor Networks*, Ph.D Thesis, Imperial College London, 2007.
- [27] L. Atallah, M. ElHelw, J. Pansiot, D. Stoyanov, L. Wang, B. Lo and G.Z. Yang, "Behaviour profiling with ambient and wearable sensing," In *Proceedings of the 4th Int. Workshop on Wearable and Implantable Body Sensor Networks*, 2007.
- [28] R. Ali, M. ElHelw, L. Atallah, B. Lo and G.Z. Yang, "Pattern mining for routine behaviour discovery in pervasive healthcare environments," In *5th International Conference on Information Technology and Application in Biomedicine (ITAB)*, 2008.
- [29] B. Lo, L. Atallah., O. Aziz, M. ElHew, A. Darzi and G.Z. Yang, "Real-time pervasive monitoring for postoperative care," In *Proc. Of 4th International Workshop on Wearable and Implantable Body Sensor Networks*, 2007.

¹ www.saphe.info/