Interpreting streaming biosignals: in search of best approaches to augmenting mobile health monitoring with machine learning for adaptive clinical decision support

V.M. Jones[1], R. Batista[1], R.G.A. Bults[1], H. op den Akker[1,2], I. Widya[1], H. Hermens[1,2], R. Huis in't Veld[2], T. Tonis[2], M. Vollenbroek-Hutten[1,2].

[1] Faculty of Electrical Engineering, Mathematics and Computer Science, University of Twente, Drienerlolaan 5, 7522 NB Enschede, The Netherlands

[2] Roessingh Research and Development, Roessinghsbleekweg 33b, 7522 AH Enschede, The Netherlands

v.m.jones@utwente.nl

Abstract. We investigate Body Area Networks for ambulant patient monitoring. As well as sensing physiological parameters, BAN applications may provide feedback to patients. Automating formulation of feedback requires real-time analysis and interpretation of streaming biosignals and other context and knowledge sources. We illustrate with two prototype applications: the first is designed to detect epileptic seizures and support appropriate intervention. The second is a decision support application aiding weight management; the goal is to promote health and prevent chronic illnesses associated with overweight/obesity. We begin to explore extending these and other m-health applications with generic AI-based decision support and machine learning. Monitoring success of different behavioural change strategies could provide a basis for machine learning, enabling adaptive clinical decision support by personalising and adapting strategies to individuals and their changing needs. Data mining applied to BAN data aggregated from large numbers of patients opens up possibilities for discovery of new clinical knowledge.

Keywords Telemonitoring, mobile health, Body Area Networks, biosignals, Clinical Decision Support, data mining, machine learning, context awareness.

1 Introduction

We research the use of Body Area Networks (BANs) and wireless communications to provide remote monitoring and treatment services to patients. Depending on the requirements of the specific clinical application, these mobile applications may provide real-time feedback and advice to the (mobile) patient as well as performing sensing of physiological parameters. Recognition of alarm conditions and automatic formulation of feedback and advice may require real-time analysis and interpretation of biosignal data streams together with data from other context and knowledge sources. Section 2 below presents a short overview of our research on mobile health systems. Section 3 presents in more detail two contrasting applications involving analysis and interpretation of biosignal data streams together with data from other context and knowledge sources. Section 4 raises the question of how best to extend the approach with machine learning (to achieve adaptive CDS) and with data mining (to enable discovery of new clinical knowledge). Section 5 presents some discussion and future directions.

2 Health BANs for telmonitoring/teletreatment

At the University of Twente we have been researching mobile health systems based on Body Area Networks (BANs) since 2001 [1]. We define a health BAN as a network of communicating devices worn on, around or in the body which provides mobile health services to the user. In our generic architecture a BAN consists of an MBU (Mobile Base Unit, handling communication, storage and local processing) and a set of BAN connected devices (e.g. sensors, actuators, positioning devices). The MBU has been implemented on various PDA and smart phone platforms. Sensor data is collected, processed and transmitted to a remote healthcare location via the MBU. The generic architecture, a first prototype health BAN and a number of variants of the BAN for different clinical applications were prototyped and trialled during IST MobiHealth [2-3]. In MobiHealth BAN data was transmitted over GPRS and UMTS to hospitals and a health call centre. The nine trials included telemonitoring of patients with cardiac arrhythmias, COPD patients, pregnant women and casualties in trauma care. BAN development continued and new variants of the BAN for different clinical applications including epilepsy and chronic pain were developed in the Dutch FREEBAND Awareness project [4-5] European eTen project HealthService24 [6], the European eTen project MYOTEL [7-8] and the Dutch project FOVEA [9]. The MobiHealth BAN applications simply transmitted and displayed biosignals remotely whilst Awareness introduced analysis and interpretation of biosignals in the light of context data. By including feedback loops, BAN telemonitoring was augmented in Myotel with teletreatment services to provide information, advice, coaching and (bio)feedback to patients.

Each clinical application requires a specific set of sensors as well as development of application-specific software and user interfaces. Sensors which have been integrated into the BAN to date include electrodes for measuring ECG and EMG, pulse oxymeter, motion sensors (step counters, 3D accelerometers), temperature and respiration sensors. Apart from sensors, other devices which have been incorporated into different variants of the BAN include positioning devices, alarm buttons and a multimodal biofeedback device.

3 Two Example Applications

The first example is a healthcare application for telemonitoring of patients with temporal lobe epilepsy. This was one of three prototype applications developed during the Awareness project. The second example is a wellbeing application designed to support weight management with the goal of preventing chronic illnesses associated with overweight and obesity. This prototype is under development in the FOVEA project. These contrasting examples are selected to illustrate the wide range of clinical applications and resulting BAN designs which our approach covers whilst being based on a generic architecture and common middleware. The epilepsy and weight management applications represent a chronic healthcare application and a health and wellbeing application respectively. The former includes emergency scenarios while the latter does not. They use different devices and have different application functionality and different interface and dialogue requirements. Both involve analysis and interpretation of biosignals in combination with other knowledge and context sources.

3.1 Epilepsy monitoring

The epilepsy scenario involves processing and analysis of biosignals and context information, including positioning of patient and carers, to identify medical emergencies (seizures) and facilitate appropriate response. The scenario was used primarily to explore the possibilities of the technology and to experiment with adding context awareness to BAN applications. An experimental seizure detection algorithm was designed to run on the BAN. The algorithm applies data fusion to changes in heart rate and posture and activity information in order to attempt to discriminate between heart rate changes due to epileptic seizure and those due to other causes including physical activity.

Fig. 1. shows the components of the Epilepsy BAN. It incorporates an Xsens MT9-B inertial sensor sensing 3D acceleration, three electrodes (Ag/AgCl contact electrodes) measuring ECG and the Mobi8-MT9 sensor front-end. The MBU is implemented on a smart phone (in this case an HTC P3600). The electrodes are placed on the patient one centimetre below the right extremitas sternalis, on the fourth left rib below the armpit and at the spinal cord at C7 (reference electrode). In this design, simple rule based decisions are made on the basis of biosignal data. When the patient has a zero or very low activity level and heart rate increase reaches a predefined threshold, the event is labelled as a possible seizure. If posture is lying, or changes to lying, the probability that the patient is having a seizure is revised upwards. Heart rate is derived from the ECG signal, sampled at 1024 Hz, by RTop detection (algorithm adopted from [10]). The beat to beat heart rate is converted to heart rate change by calculating the difference between the mean heart rate in two moving time windows

of 10 and 120 seconds. 3D accelerometer data is sampled at 128 Hz. Activity level [11] and posture (lying or not, detected by reference to the earth's gravitational field) are calculated every 10 seconds. The internal GPS device of the HTC P3600, together with cell-ID information, is used for location determination so that appropriate assistance can be dispatched if a seizure is detected. The specialist is notified in case of a detected seizure and can view the patient's biosignals and location.

Fig. 1. Epilepsy BAN: electrodes and activity sensor

Fig. 2. shows the display of biosignals on the m-health portal. Three traces are displayed against a time axis: ECG, activity level and heart rate (derived from ECG). At the right hand side numeric readouts for activity and heart rate are displayed.

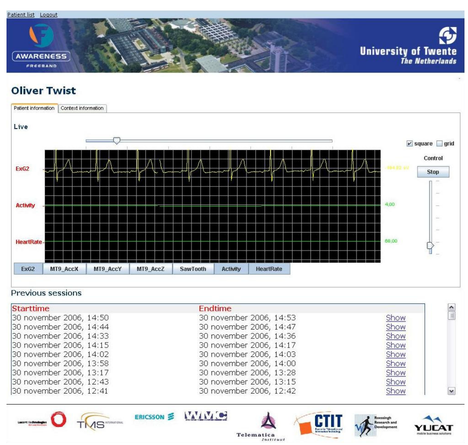

Fig. 2. The m-health portal displaying three biosignal traces: ECG, activity and heart rate

The professional can also view location information. Fig. 3. shows the map display showing the position of the patient and the positions and status (availability) of all the informal carers who are registered for that patient. The nearest available carer can be dispatched and guided to the patient using GPS positioning and a map display on the carer's mobile device.

Fig. 3. Awareness Epilepsy application: location display when a seizure is detected

This BAN application was trialled (on healthy volunteers only) during Awareness in order to evaluate the technical performance of the system. Algorithms for derivation of heart rate and activity level were implemented on the BAN but the seizure detection algorithm could only be tested offline (on data from healthy subjects) due to computational limitations of the PDA. The detection algorithm still needs to be tested on data from epileptic patients before specificity, sensitivity etc. can be established. However the experiments with healthy subjects showed some false positives.

3.2 Sustainable weight management

Overweight and obesity are associated with increased morbidity and mortality and are on the increase around the world. Many governments, HMOs and other organizations are attempting to promote healthier lifestyles, with weight management as a major goal. However, even when public awareness of the health consequences of unhealthy eating and drinking has been raised, this knowledge has had little effect on consumer behavior in terms of lasting change in dietary and exercise patterns. Achieving and maintaining weight management goals requires more than intellectual recognition; interventions need to be informed by behavioural change theory as well as nutrition education theory. How best to effect sustainable behavioural change, critical to the success of many health promotion initiatives, has been an area of research for some years.

Partners in the Dutch project FOVEA study how to change consumer behaviour in the direction of a healthier lifestyle by applying behavioural theory with support from ICT, including ambulatory monitoring technology. The aim is to support sustainable behavioural change with respect to food and drink consumption and exercise in order to improve health and wellbeing and prevent chronic illness. The approach is to provide real-time personalised feedback and advice at the point of decision making. The advice is tailored to the individual's weight management goals and stage of change. As a use case, we target the *inclined abstainer* who is an *external eater* in the *Action* stage of the *Stages of Change* model [12-13]. The Restaurant of the Future (RoF) [14], a company restaurant in Wageningen, The Netherlands, provides an instrumented environment which is used in this and other projects as a testbed for interactive research in a real life setting. The RoF insfrastructure includes video cameras for behavioural observation, weight-scales at the checkouts and automatic registration of individuals' food and drink consumption at the point of sale terminals. It also offers possibilities for altering the ambient environment in order to investigate effects on physiology and behaviour of subtle changes in environmental factors.

A prototype of the FOVEA system, integrating components from the RoF infrastructure with the food database of the canteen supplier, has been developed. The FOVEA system includes a mobile system designed to give real time monitoring and personalized feedback. The University of Twente is responsible for developing this mobile system. The FOVEA mobile system conforms to our BAN architecture; in this case the BAN consists of a smart phone and a single sensor: the smart phone's on board accelerometer. The FOVEA system, including the mobile system, is being trialled in 2011 at the RoF on 60 trial subjects selected from regular visitors to the RoF who have BMI 25.00-29.99 (WHO classification "overweight, pre-obese" [15]). A dietician formulates a plan with each user, including up to five lunch compositions (each with up to five food items) to match the user's personal goals and preferences.

The mobile part of the FOVEA system is implemented as a smart phone application on a Samsung Galaxy S running Android. Food and drink consumption is registered, physical activity is monitored and feedback is given in real time. The mobile application detects the different self service buffets (using Bluetooth or UMTS) in the RoF. By means of indoor positioning, a map of the layout of the RoF, connection with the food database and knowledge of the individual's targets, the mobile application is able to guide the user away from less healthy options (using a cue avoidance strategy) and towards healthier options and balanced meal compositions. Fig. 4 shows example screenshots, in this case a user profile and the layout of the RoF. Icons indicate the entrance (red arrow), food buffets (e.g. bread, hot meals, salads, soup etc.) and the point of sale terminals (which are linked into the RoF information systems).

The FOVEA mobile system stores the user's profile and keeps track of their energy balance throughout the day. In the RoF it can be used to display the food and beverage selections on offer that day, allowing the user to check an item before making a selection. "Good" and "bad" selections are highlighted to help the user make healthy consumption decisions. Fig. 5 shows more example screenshots, here displaying the list of buffets and the options for warm beverages. The "compliant" and "non compliant" items (for this user) are highlighted in green and orange respectively. Compliant means that this item is part of the lunch composition selected by the user on this occasion. In this case we see that *koffie verkeerd* (coffee made with milk) and *thee* (tea) are compliant with the current lunch composition.

Energy values of food and beverage items in the RoF are stored in the food database and are used to track total kilocalories of items selected. Registration of food and drink consumption on the phone enables real time estimation of energy intake and helps the user to manage their daily energy budget. Fig. 6 shows the energy screen, and the impact of consuming one item (a *mueslibol*) on energy balance. (A mueslibol is a sweet bread roll containing muesli.) The screen shows that a mueslibol contains 150 KiloCalories and is not compliant with the current lunch composition.

Real-time measurement of physical activity (using the smart phone's on-board 3-axis accelerometer) is used to estimate energy expenditure (EE) in real time. The EE

Fig. 4. FOVEA mobile system screenshots: user profile and restaurant layout.

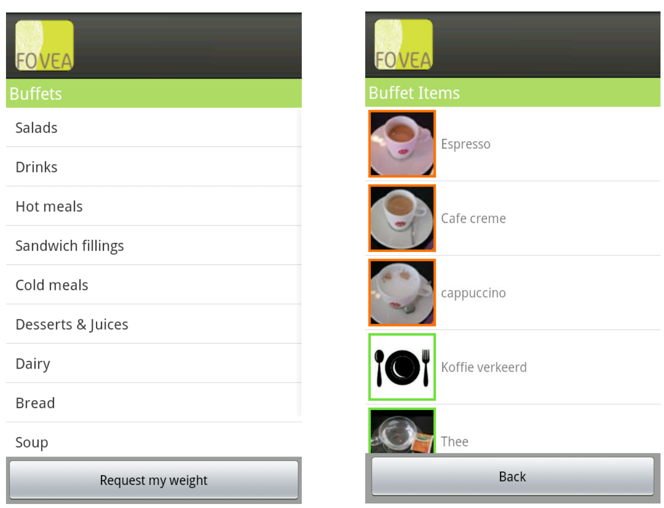

Fig. 5. Screenshots: buffet list and warm beverage options

algorithm was developed during the project and is being calibrated against validated methods including a commercially available step counter (the OMRON Walking

Style Pro). Two different methods of calculating basal metabolic rate (BMR) are used: Harris Benedict (HB) and Miffin and St Jeor (MJ). Hence two versions of the Temporary Energy Balance are displayed. The yellow smiley in Fig. 6 indicates that at least one of the two versions of Temporary Energy Balance will go from positive to negative if the user consumes this item (i.e. cumulative energy intake will exceed cumulative energy expenditure if this item is consumed). If both Temporary Energy Balances would remain above zero the smiley would be green. In case at least one of the previous energy balances was already below zero the smiley would be red. Even if the item is not compliant with the predetermined lunch composition, or sends the Temporary Energy Balance negative, the user is free to purchase the item and register their selection by pressing *Add*.

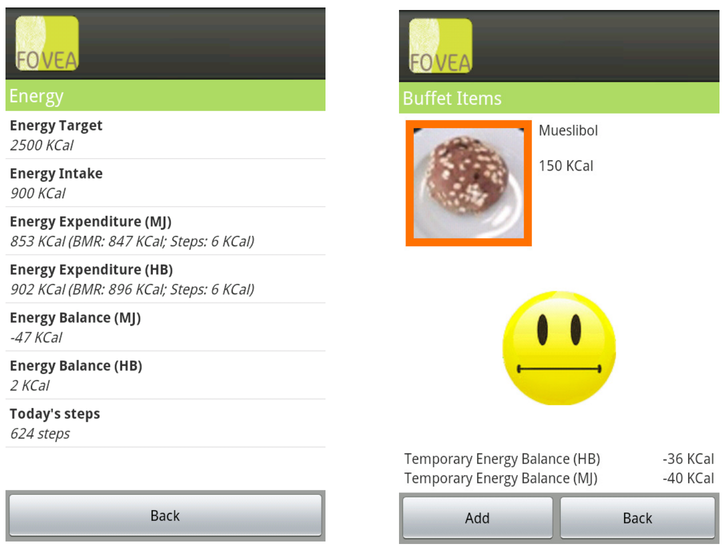

Fig. 6. Screenshots: energy screen, and impact of one item on energy balance. HB - Harris Benedict method of estimating BMR; MJ - Miffin and St Jeor method.

The activity monitoring and feedback mechanisms developed in FOVEA could also be applied in a number of other applications, e.g. optimizing daily physical activity patterns in COPD, cardiac rehabilitation and chronic pain management.

4. Machine learning and data mining for adaptive CDS and knowledge discovery

Two contrasting applications, for healthcare and health and wellbeing respectively, are presented as illustrative examples of mobile monitoring and feedback applications. Both involve analysis and interpretation of (streaming) physiological signals in combination with other knowledge and context sources.

FOVEA is an example of a mobile monitoring and feedback application which would benefit by being augmented with more intelligent decision support. We believe the system can be more effective if it can monitor and learn from the effectiveness of various behavioural change strategies and thus adapt to the individual user, so improving chances of success in reaching and maintaining personal lifestyle goals.

We propose that in future the FOVEA mobile application be extended to monitor individual performance in terms of adherence to weight management guidelines and (changing) personal diet and exercise plans, and so be able to adapt to those strategies which prove to be more successful for the individual in both the short and longer term. We expect that the resulting augmented application will provide more intelligent and personalised adaptive decision support.

Furthermore, data mining techniques applied to aggregations of BAN data from large numbers of patients opens up possibilities for discovery of new clinical knowledge. Epileptic seizures are rare but serious clinical events which moreover have far reaching consequences for patients' daily living. Routine collection of biosignal and context data on a large population over time by means of ambulatory monitoring would yield a large dataset which could serve medical research purposes. We surmise that the epilepsy case could possibly benefit from applying machine learning and data mining to accumulated biosignal and context data from many patients, possibly uncovering new medical knowledge which could potentially improve the detection of ongoing seizures and ability to predict upcoming seizures.

5. Discussion and Future Work

The ability in general to automate delivery of feedback and advice implies theory-based clinical decision support functionality which in turn relies on real-time interpretation and analysis of streaming biosignals together with other context and knowledge sources. The examples of monitoring and feedback applications discussed above are offered as a basis for discussion of possible approaches for extending health BAN applications with theory-based adaptive clinical decision support delivered on a mobile platform in an ambulatory setting. In many clinical applications monitoring the success of different behavioural change strategies in the short and longer term can provide a basis for machine learning, enabling personalisation and adaptation of strategies to the individual and their changing needs over time.

The early prototype of the FOVEA mobile system is a proof of concept and applies some pragmatic solutions. It can be extended and improved in numerous ways; for example by connecting to other environments to give the user decision support, bolster motivation and apply cue management in other contexts, e.g. at home, in other restaurants or when shopping for food. This early prototype can be said to deliver decision support, however this is a behavioural description; the system does not (yet) apply AI techniques associated with (clinical) decision support. In this respect it resembles some of the early expert systems which were classed by some user communities as "expert" systems on the grounds that their behaviour emulated some aspects of human "expert" behaviour (ie on grounds of *outcome*). The AI community however focuses on developing underlying models and mechanisms, for example for knowledge representation and reasoning, to support the *process* of expert reasoning and decision making. Our future ambitions therefore also include the intention to develop a generic framework for (clinical) decision support, including data mining and machine learning mechanisms, which can be implemented on mobile platforms so that future BAN applications are well-grounded in generic and sound AI-based approaches.

One knowledge-based approach to augmenting mobile monitoring and treatment with adaptive real time CDS is visualized in Fig. 7. (Many other approaches are possible.) The red parts of Fig. 7 show the points of innovation over classical expert systems-based DSSs, namely: inclusion of streaming biosignals and context data; task knowledge relating to biosignal analysis and interpretation; ML applied to BAN data to effect individual adaptation; and data mining over data from large populations over time to enable discovery of new clinical knowledge. Subtracting the red parts of the figure reduces it to a depiction of a classical expert systems architecture.

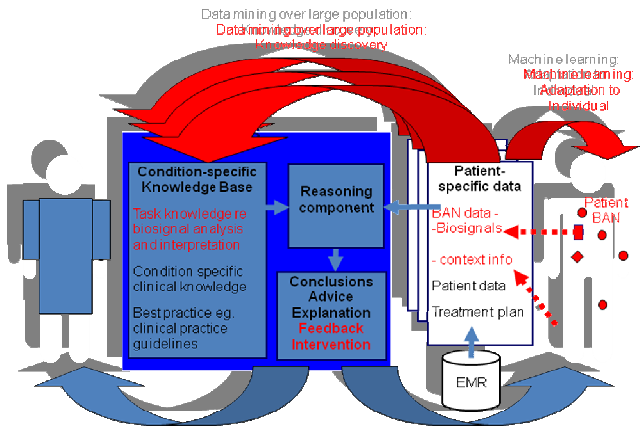

Fig. 7. One approach to adaptive CDS adding streaming data, ML and DM

We have begun to experiment with an ML approach to improving compliance to feedback. By applying classical machine learning techniques to activity data and behavioural data gathered by mobile monitoring of various patient populations with chronic illnesses, Akker et al. [16] attempted to predict, based on personal historical usage data, a user's reaction to a given feedback message. From activity and feedback data for 95 patients, feedback compliance to each individual message was calculated and relevant context features were extracted. The features include timing information (e.g. time of day, day of week), information about the type of message sent to the patient, weather related features and features related to the history of usage. Compliance to individual feedback messages ('yes' or 'no') was the class. A number of supervised classifier schemes were tested; the rule-based Ridor classifier proved best over all in terms of accuracy. Results compared to baseline (ZeroR classification) were promising and a further analysis of optimal feature sets using Genetic Algorithms gives insight into the most useful context information. Using this method, 86% classification accuracy was achieved on average over all patients that were included. During operation, whenever the machine learner or CDS component decides that compliance to feedback at this time is unlikely, the message output is delayed, reducing unnecessary burden on the patient by the system. This method translates well to any application in which the system initiates the interaction with the user, especially if the system is designed to be used for sustained periods of time, as is necessary when a change in user behaviour is the ultimate goal. Wieringa et al [17] also demonstrated that other AI-inspired methods are promising for generating the content of messages that are intended to promote behavioural change in patients or users. In this experiment, message content was generated by traversing a structured ontology of applicable motivational messages. Parts of the ontology were pruned by Boolean functions attached to the ontology entities. For example, the subtree of the ontology that specifies messages suggesting the user goes outdoors for some exercise has a Boolean function "isWeatherGood()" attached; that subtree is pruned if this function evaluates to false during operation. User preferences for specific messages or types of message are learned by storing user compliance to system-generated messages and using this information during traversal of the ontology to increase the chance of selecting message-branches that had positive reactions in the past.

Adding intelligence to these ambulant applications is made possible by the continuing miniaturization of hardware and increases in processing power available on smart phones and PDAs. Learning algorithms no longer need a bulky mainframe or server in order to run. Lightweight machine learning algorithms, and powerful support for multithreading built into current Android platforms, for example, make it feasible for mobile applications to run ever more sophisticated machine learning processes in the background.

Amongst the immediate challenges and opportunities arising, we mention:

- Incorporation of real time input and automated analysis of streaming biosignals and context data into a clinical decision support system;
- Selection of the best (generic) technical approaches and mechanisms for implementing adaptive clinical decision support on a mobile platform;
- Distribution of coherent clinical decision support functionality across a complex fixed-and-mobile distributed environment; and
- Maintenance of consistency of knowledge and beliefs in the distributed environment.

As Computer Science and Biomedical Engineering researchers involved in mobile healthcare and health and wellbeing applications we offer these ideas to the workshop with the expectation of stimulating discussion with experts in data stream mining and machine learning on best approaches to developing a sound AI-based framework for next generation mobile health applications with adaptive clinical decision support.

Acknowledgements

Grateful thanks go to our Awareness and FOVEA partners, funding bodies and trial subjects. Freeband AWARENESS (http://awareness.freeband.nl) was sponsored by the Dutch government under contract BSIK 03025. FOVEA (http://www.fovea-project.org/) is funded by the Dutch Ministry of Economic Affairs and the Province of Gelderland under the programme: Pilot Gebiedsbericht Innovatiebeleid.

References

1. V. M. Jones, R. G. A. Bults, D. M. Konstantas, and P. A. M. Vierhout, "Healthcare PANs: Personal Area Networks for trauma care and home care," presented at Fourth International Symposium on Wireless Personal Multimedia Communications (WPMC), Aalborg, Denmark, 2001.
2. MobiHealth project: http://www.mobihealth.org/
3. Val Jones, Aart van Halteren, Ing Widya, Nikolai Dokovsky, George Koprinkov, Richard Bults, Dimitri Konstantas and Rainer Herzog (2006). MobiHealth: Mobile Health Services based on Body Area Networks, IN M-Health: Emerging Mobile Health Systems, Robert H. Istepanian, Swamy Laxminarayan, Constantinos S. Pattichis, Editors, Springer. pp.219-36. 624 p. ISBN: 0-387-26558-9.
4. AWARENESS project, http://www.freeband.nl/project.cfm?id=494&language=en.
5. V.M. Jones, H. Mei, T. Broens, I. Widya, and J. Peuscher, Context Aware Body Area Networks for Telemedicine, Proceedings Pacific-Rim Conference on Multimedia (PCM) 2007, City University of Hong Kong, Hong Kong, 11-14 December 2007.
6. eTen HealthService 24 project, http://www.healthservice24.com.
7. MYOTEL project: http://www.myotel.eu/
8. V.M. Jones, R. Huis in't Veld, T. Tonis, R.B. Bults, B. van Beijnum, I. Widya, M. Vollenbroek-Hutten, H. Hermens. Biosignal and Context Monitoring: Distributed Multimedia Applications of Body Area Networks in Healthcare. Proceedings of the 2008 IEEE 10th Workshop on Multimedia Signal Processing, (MMSP2008), 8-10 October 2008, Cairns, Australia, Editors: David Feng, Thomas Sikora, W.C. Siu, Jian Zhang, Ling Guan, Jean-Luc Dugelay, Qiang Wu, Wanqing Li, pp. 820-825. IEEE Catalog Number CFP08MSP-CDR, ISBN 978-1-4244-2295-1, Library of Congress 2008902132. Copyright © 2008 IEEE
9. http://www.fovea-project.org/
10. Christov, I, Real time electrocardiograrm QRS detection using combined adaptive threshold, Biomedical Engineering Online, 2004, 3:28
11. C.V. Bouten, W.P. Verboeket-van de Venne, et al, "Daily physical activity assessment: comparison between movement registration and doubly labeled water", J. Appl. Physiol, 1996, 81, 1019-1026.
12. Velicer, W.F., Prochaska, J.O., Fava, J.L., Norman, G.J., Redding, C.A.: Smoking cessation and stress management: Applications of the Transtheoretical Model of behavior change. Homeostasis 38, 216-233 (1998)
13. I. Widya, R.G.A. Bults, R. de Wijk, B. Loke, N. Koenderink, R.J. Mendes Batista, V.M. Jones, H.J. Hermens, Requirements for a Nutrition Education Demonstrator, 17th. Intl. Working Conference on Requirements Engineering: Foundation for Software Quality (RefsQ2011), Essen, Germany, pp. 48-53, March 28-30, 2011. Eds. D. Berry and X. Franch, LNCS 6606, Springer-Verlag, Berlin.
14. http://www.restaurantvandetoekomst.wur.nl/UK/
15. http://apps.who.int/bmi/
16. Harm op den Akker, Val Jones, Hermie Hermens, Predicting Feedback Compliance in a Teletreatment Application, Proceedings 3rd International Symposium on Applied Sciences in Biomedical and Communication Technologies: Sustainable Communication and Biomedical Technologies, (ISABEL 2010), November 7-10, 2010, Rome, Italy, pp.1-5.
17. Wilko Wieringa, Harm op den Akker, Val M. Jones, Rieks op den Akker, and Hermie J. Hermens, Ontology-Based Generation of Dynamic Feedback on Physical Activity, Accepted for: 13th Conference on Artificial Intelligence in Medicine (AIME'11), July 2-6, 2011, Bled, Slovenia.