

Smart Environments and Context-Awareness for Lifestyle Management in a Healthy Active Ageing Framework

Davide Bacciu¹, Stefano Chessa^{1,2(✉)}, Claudio Gallicchio¹, Alessio Micheli¹,
Erina Ferro², Luigi Fortunati², Filippo Palumbo², Oberdan Parodi², Federico Vozzi²,
Sten Hanke³, Johannes Kropf³, and Karl Kreiner⁴

¹ Department of Computer Science, University of Pisa, Largo Pontecorvo 3, Pisa, Italy

² CNR-ISTI, PISA CNR Research Area, via Moruzzi 1, 56126 Pisa, Italy

stefano.chessa@unipi.it

³ Health and Environment Department, AIT Austrian Institute of Technology GmbH,
Vienna, Austria

⁴ Safety and Security Department, AIT Austrian Institute of Technology GmbH,
Vienna, Austria

Abstract. Health trends of elderly in Europe motivate the need for technological solutions aimed at preventing the main causes of morbidity and premature mortality. In this framework, the DOREMI project addresses three important causes of morbidity and mortality in the elderly by devising an ICT-based home care services for aging people to contrast cognitive decline, sedentariness and unhealthy dietary habits. In this paper, we present the general architecture of DOREMI, focusing on its aspects of human activity recognition and reasoning.

Keywords: Human activity recognition · E-health · Reasoning · Smart environment

1 Introduction

According to the University College Dublin Institute of Food and Health, three are the most notable health promotion and disease prevention programs that target the main causes of morbidity and premature mortality: malnutrition, sedentariness, and cognitive decline, conditions that particularly affect the quality of life of elderly people and drive to disease progression. These three features represent the target areas in the DOREMI project. The project vision aims at developing a systemic solution for the elderly, able to prolong the functional and cognitive capacity by stimulating, and unobtrusively monitoring the daily activities according to well-defined “Active Ageing” lifestyle protocols. The project joins the concept of prevention centered on the elderly, characterized by a unified vision of being elderly today, namely, a promotion of the health by a constructive interaction among mind, body, and social engagement.

This work has been funded in the framework of the FP7 project “Decrease of cOgnitive decline, malnutRition and sedEntariness by elderly empowerment in lifestyle Management and social Inclusion” (DOREMI), contract N.611650.

© Springer International Publishing Switzerland 2015

F. Pereira et al. (Eds.) EPIA 2015, LNAI 9273, pp. 54–66, 2015.

DOI: 10.1007/978-3-319-23485-4_6

To fulfill these goals, food intake measurements, exergames associated to social interaction stimulation, and cognitive training programs (cognitive games) will be proposed to an elderly population enrolled during a pilot study. The DOREMI project is going further with respect to the current state of the art by developing, testing, and exploiting with a short-term business model impact a set of IT-based (Information Technology) services able to:

- Stimulate elderly people in modifying dietary needs and physical activity according to the changes in age through creative, personalized, and engaging solutions;
- Monitor parameters of the elderly people to support the specialist in the daily verification of the compliance of the elderly with the prescribed lifestyle protocol, in accordance with his/her response to physical and cognitive activities.
- Advise the specialist with different types and/or intensities of daily activity for improving the elderly health, based on the assigned protocol progress assessment.
- Empower aging people by offering them knowledge about food and physical activity effectiveness, to let them become the main actors of their health.

To reach these objectives, the project builds over interdisciplinary knowledge encompassing health and artificial intelligence, the latter covering aspects ranging from sensing, machine learning, human-machine interfaces, and games. This paper focuses on the machine learning contribution of the project, which applies to the analysis of the sensor data with the purpose of identifying users' conditions (in terms of balance, calories expenditure, etc.) and activities, detecting changes in the users' habits, and reasoning over such data. The ultimate goal of this data analysis is to support the user who is following the lifestyle protocol prescribed by the specialist, by giving him feedbacks through an appropriate interface, and by providing the specialist with information about the user lifestyle. In particular, the paper gives a snapshot of the status of the project (which just concluded the first year of activity) in the design of the activity recognition and reasoning components.

2 Background and State of the Art on Machine Learning

Exploratory data analysis (EDA) analyzes data sets to find their main features [1], beyond what can be found by formal modeling or hypothesis testing task. When dealing with accelerometer data, features are classified in three categories: time domain, frequency domain, and spatial domain [2]. In the time domain, we use the standard deviation in a frame, which is indicative of the acceleration data and the intensity of the movement during the activity. In the frequency domain, frequency-domain entropy helps the distinction of activities with similar energy intensity by comparing their periodicities. This feature is computed as the information entropy of the normalized Power Spectral Density (PSD) function of the input signal without including the DC component (mean value of the waveform). The periodicity feature evaluates the periodicity of the signal that helps to distinguish cyclic and non-cyclic activities. In the spatial domain, orientation variation is defined by the variation of the gravitational

components at three axes of the accelerometer sensor. This feature effectively shows how severe the posture change can be during an activity.

Other EDA tasks of the project concern unsupervised user habits detection aimed at finding behavioral anomalies, by retrieving heterogeneous and multivariate timeseries of sensor data, over long periods. In the project, these tasks are unsupervised to avoid obtrusive data collection campaign at the user site. For this reason, we focus on motif search on sensory data collected in the test by exploiting the results obtained in the field of time series motifs discovery [3,4]. Time series motifs are approximately repeated patterns found within the data. The approach chosen is based on *stigmery*. Several works used this technique in order to infer motifs in time series related to different fields, from DNA and biological sequences [5,6] to intrusion detection systems [7].

Human activity recognition refers to the process of inferring human activities from raw sensor data [8], classifying or evaluating specific sections of the continuous sensors data stream into specific human activities, events or health parameters values. Recently, the need for adaptive processing of temporal data from potentially large amounts of sensor data has led to an increasing use of machine learning models in activity recognition systems (see [9] for a recent survey), especially due to their robustness and flexibility. Depending on the nature of the treated data, of the specific scenario considered and of the admissible trade-off among efficiency, flexibility and performance, different supervised machine learning methods have been applied in this area.

Among others, Neural Network for sequences, including Recurrent Neural Networks (RNNs) [10], are considered as a class of learning models suitable for approaching tasks characterized by a sequential/temporal nature, and able to deal with noisy and heterogeneous input data streams. Within the class of RNNs, the Reservoir Computing (RC) paradigm [11] in general, and the Echo State Network (ESN) model, [12,13] in particular, represent an interesting efficient approach to build adaptive non-linear dynamical systems. The class of ESNs provides predictive models for *efficiently* learning in sequential/temporal domains from heterogeneous sources of noisy data, supported by theoretical studies [13,14] and with hundreds of relevant successful experimental studies reported in literature [15]. Interestingly, ESNs have recently proved to be particularly suitable for processing noisy information streams originated by sensor networks, resulting in successful real-world applications in supervised computational tasks related to AAL (Ambient Assisted Living) and human activity recognition. This is also testified by some recent results [16,17,18,19,20], which may be considered as a first preliminary experimental assessment of the feasibility of ESN to the estimation of some relevant target human parameters, although obtained on different and broader AAL benchmarks.

At the reasoning level, our interest is for hybrid approaches founded on static rules and probabilistic methods. Multiple-stage decisions refer to decision tasks that consist of a series of interdependent stages leading towards a final resolution. The decision-maker must decide at each stage what action to take next in order to optimize performance (usually utility). Some examples of this sort are working towards a degree, troubleshooting, medical treatment, budgeting, etc. Decision trees are a useful mean for representing and analyzing multiple-stage decision tasks; they support decisions learned from data, and their terminal nodes represent possible consequences [21].

Other popular approaches, which have been used to implement medical expert systems, are Bayesian Networks [22] and Neural Networks [23], but they require many empirical data to train the algorithms and are not appropriate to be manually adjusted. On the other hand, in our problem the decision process must be transparent and mainly requires static rules based on medical guidelines provided by the professionals. Thus, the decision trees are the best solution since they provide a very structured and easy to understand graphical representation. There also exist efficient and powerful algorithms for automated learning of the trees [24,25,26]. A decision tree is a flow-chart-like structure in which an internal node represents the test on an attribute, each branch represents a test outcome and each leaf node represents a class label (decision taken after computing all attributes). A path from root to leaf represents classification rules. Decision trees give a simple representation for classifying examples. In general, as for all machine learning algorithms, the accuracy of the algorithms increases with the number of sample data. In applications in which the number of samples is not large, a high number of decisions could lead to problems. In these cases, a possible solution is the use of a Hybrid Decision Tree/Genetic Algorithm approach as suggested in [27].

3 Problem Definition and Requirements

Our main objective is to provide a solution for prolonging the functional and cognitive capacity of the elderly by proposing an “Active Ageing” lifestyle protocol. Medical specialists monitor the progress of their patients daily through a dashboard and modify the protocol for each user according to their capabilities. A set of mobile applications (social games, exer-games, cognitive games and diet application) feedback the protocol proposed by the specialist and the progress of games to the end user. The monitoring of each user is achieved by means of a network of sensors, either wearable or environmental, and applications running on personal mobile devices. The human activity recognition (HAR) measures characteristics of the elderly lifestyle in the physical and social domains through non-invasive monitoring solutions based on the sensor data. Custom mobile applications cover the areas of diet and cognitive monitoring. In the rest of this section, we present the main requirements of the HAR.

By leveraging environmental sensors, such as PIRs (Passive InfraRed) and a localization system, the HAR module profiles user habits in terms of daily ratio of room occupancy and indoor/outdoor living. The system is also able to detect changes in the user habits that occur in the long-term. By relying on accelerometer and heartbeat data from a wearable bracelet, the HAR module provides time-slotted estimates, in terms of calories, of the energy expenditure associated with the physical activities of the user. Energy consumption can result from everyday activities and physical exercises proposed by the protocol. The system also computes daily outdoor covered distance, the daily number of steps and detects periods of excessive physical stress by using data originated by the accelerometer and the heartbeat in the bracelet. Finally, a smart carpet is used to measure the user weight and balance skills, leveraging a machine learning classification model based on the BERG balance assessment test.

The HAR assesses the social interactions of the user both indoor and outdoor. In particular, in the indoor case, HAR estimates a quantitative measure of the social interactions based on the occurrence and duration of the daily social gatherings at the user house. Regarding the outdoor socialization, the system estimates the duration of the encounters with other users by detecting the proximity of the users' devices.

The Reasoner uses the data produced by the HAR, the diet, and the games applications to provide an indicator of the user protocol compliance and protocol progress in three areas: social life, physical activity and related diet, cognitive status. These indicators, along with the measured daily metrics and aggregate data, support medical specialists on providing periodical changes to the protocol (i.e.: set of physical activities and games challenges, diet). The Reasoner is able to suggest changes to the user protocol by means of specialist-defined rules.

The HAR module and the Reasoner are, therefore, core system modules, bridging the gap between sensors data, medical specialists, and the end user.

4 An Applicative Scenario

We consider a woman in her 70s, still independent and living alone in her apartment (for the sake of simplicity, we give her the name of Loredana). She is a bit overweight and she starts forgetting things. Recently, the specialist told her that she is at risk for cardiovascular disease, due to her overweight condition, and that she has a mild cognitive impairment. For this reason, Loredana uses our system as a technological support to monitor her life habits and to keep herself healthier and preventing chronic diseases. In a typical day, Loredana measures her weight and balance by means of a smart carpet, which collects data for the evaluation of her BERG scale equilibrium and her weight. The data concerning the balance is used to suggest a personalized physical activity (PA) plan, while the data about the weight give indications about the effectiveness of the intervention in terms of a personalized diet regimen and PA plan. During the day, Loredana wears a special bracelet, which measures (by means of an accelerometer) her heart rate, how much she walked, and how many movements she did with physical exercises. These data are used, by the system developed in the project, to assess her calories expenditure and to monitor the execution of the prescribed physical exercises. The bracelet is also used to localize her both indoor (also collecting information about the time spent in each room) and outdoor (collecting information about distance covered). Furthermore, the bracelet detects the proximity of Loredana with other users wearing the same bracelet, while machine-learning classification models based on environmental sensors deployed at home (PIRs and door switches) detect the presence of other people in her apartment to give indication about the number of received visits. These data are used as an indicator of her social life.

Loredana also uses a tablet to interact with the system, with which she performs cognitive games and inserts data concerning her meals, which are converted by the system (under the supervision of her specialist) in daily Kcalories intake and food composition. She is also guided through the daily physical exercises and games that are selected by the system (under the supervision of her specialist) based on the evolution

of her conditions (in terms of balance, weight, physical exercises etc.). All the data collected during the day are processed at night to produce a summary of the Loredana lifestyle, with the purpose of giving feedbacks to Loredana in terms of proposed physical activity, and presenting the condition of Loredana to the specialist on a daily basis.

5 Activity Recognition and Reasoning

5.1 High Level Architecture and Data Flows

The high-level system architecture is presented in Fig. 1, highlighting the data flow originated at a pilot site (in terms of sensor data), through the data processing stages (preprocessing, activity recognition and reasoner subsystems), and then back to the user (in terms of feedbacks) and to the specialist (in terms of information about the user performance).

In particular, Fig. 1 shows that the data processing system contains five main subsystems running on the server (the grey rectangles in the figure), three databases (RAW, HOMER [28] and KIOLA [29]), plus a middleware that uploads the sensors' data in the RAW database, whose description is out of the scope of this paper. At night, a synchronization mechanism (shown as a clock in Fig. 1), sequentially activates these five subsystems. In turn, these subsystems pre-process data (preprocessing subsystem in the Fig. 1), configure the predictive activity recognition tasks (task configurator subsystem), process the daily pre-processed data through the predictive human activity recognition subsystem (HAR subsystem), perform the exploratory data analysis (EDA subsystem), and refine and aggregate the results of these stages (Reasoner subsystem).

Along with the sensors' data, the RAW DB also stores the intermediate data produced by the pre-processing subsystem. The HOMER DB stores configuration information (e.g. regarding sensors deployment) that is used by the task configurator to retrieve the tuning parameters for the different pilot sites. The refined data produced by the HAR and EDA subsystems is then stored in the KIOLA DB, where the Reasoner reads them. The Reasoner outputs feedbacks for the user in terms of suggestions about her lifestyle, and data about her compliance to the suggested lifestyle protocol for the specialist (or caregiver) through the dashboard.

The data processing stages deal with three flows of data, related to the user diet, social relationships, and sedentariness, respectively. The dietary data flow relies on data produced by the smart carpet (pressure data and total weight), the bracelet (heart rate and accelerometers data), and data about the food composition provided by the user himself through an interface on his tablet. In particular, the data produced by the bracelet pass through the HAR subsystem that estimates the user physical activity. The data flow about the user's social relationships relies on a number of environmental sensors, which detect the contacts of the user with other people. To this purpose, the HAR and EDA subsystems detect the user encounters and the proximity of the user with other people by fusing information produced by presence sensors, user's localization and door switches. This data flow also relies on the user's mood

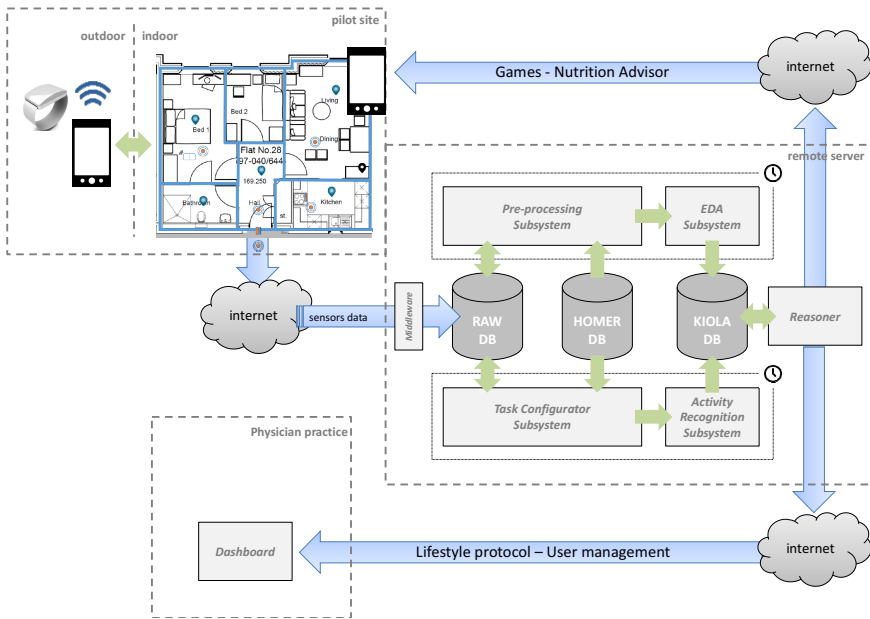


Fig. 1. High-level architecture and deployment of a typical installation, with data flowing from a pilot site to the remote Activity Recognition and Reasoning system.

information, which the user himself asserts daily through an app on his tablet. Finally, the sedentariness data flow exploits data produced by the bracelet (heart rate, user's localization, movements, and step count), the smart carpet, and data about the use of the application that guides the user through the daily physical activity. The HAR subsystem processes the data of the smart carpet to assess the user balance according to the BERG scale. The HAR and EDA subsystems also process data from the bracelet to assess the intensity of the physical effort.

The Reasoner fuses all these data flows at a higher level than that of HAR and EDA. In a first step, it exploits rules extracted from clinical guidelines to compute specific parameters for each of the three data flows. For example, it uses the physical activity estimation in the sedentariness data flow to assess the compliance of the user to the prescribed lifestyle protocol (as the medical expert defines it). In the second step, the Reasoner performs a cross-domain reasoning on top of the first step, allowing a deeper insight in the well-being of the patient. Note that the empirical rules needed to define the second-level reasoning protocol are yet not available, as they are the output of medical studies on the data collected from the on-site experimentation that will be concluded in the next year of the project activity. Hence, at this stage of the project, the second level reasoning is not yet implemented.

Note that the Reasoner subsystem operates at a different time-scale with respect to the HAR and EDA subsystems. In fact, the aim of HAR and EDA is the recognition of short-term activities of the user. These can be recognized from a sequence of input sensor information (possibly pre-processed) in a limited (short-term) time window.

All short-term predictions generated across the day are then forwarded to the Reasoner for information integration across medium/long time scales. Medium term reasoning operates over 24h periods (for example, to assess the calories assumption/consumption balance in a day). Long-term reasoning, on the other hand, shows general trends by aggregating information on the entire duration of the experimentation in the pilot sites (for example to offer statistical data about the user, which the medical experts can use to assess the overall user improvement during the experimentation).

5.2 Human Activity Recognition and Exploratory Data Analysis Subsystems

The goal of the activity recognition subsystems (HAR and EDA) is to evaluate the user parameters (referred to as *predictions*) concerning his short-term activities performed during the day. It exploits the *activity recognition configuration*, which is the result of an off-line configuration phase aimed at finding the final setting for the pre-processing and for the activity recognition subsystems (both HAR and EDA) that are deployed to implement the activity recognition tasks.

The EDA subsystem analyses pre-processed data in order to profile user's habits, to detect behavioral deviations of the routine indoor activities, and to provide aggregated values useful to the Reasoner in the sedentariness area, such as user habits, daily outdoor distance covered, daily steps and information about outdoor meetings with other users. The actual nature of the data streams processed by the EDA depends on the particular sensors originating them. For example, in the case of BERG score prediction concerning the user balance, the data produced by the smart carpet have a high and variable frequency (~100 Hz). These data are normalized and broken into fixed frequency time series segments of 200 ms, from which the preprocessing stage extracts statistical features consisting in mean, standard deviation, skewness and kurtosis. The resulting features time series, presenting a lower frequency of 5 Hz, are the input of the EDA subsystem. A similar pre-processing stage is applied to data coming from the other sensors.

The unsupervised models used for the EDA do not rely on a long-term, invasive and costly ground truth collection and annotation campaigns, which may be not acceptable by the users. Rather, EDA is designed to detect symptoms of chronic diseases (which are most relevant for the project purposes), characterized by a gradual, long-term deviation from the user typical behavior, or by critical trends in the user's vital parameters. For example, EDA features a module for the detection of abnormal deviations in the user habits (based on motif discovery and stigmergy algorithms) that relies on the locations of the user (at room-level) during his daily living activities.

Concerning the HAR subsystem, its internal architecture is shown in Fig. 2. It is composed of two main subsystems, which are the task configurator and the activity recognition subsystem. The former handles the retrieval of configuration information and pre-processed data for the tasks addressed in the system and forwards it to the latter subsystem, which is responsible for performing the actual activity recognition tasks. The core of the activity recognition system is given by the HAR scheduler (which activates the activity recognition tasks when all sensor information is consolidated and pre-processed in the RAW DB), and by the pool of activity recognition

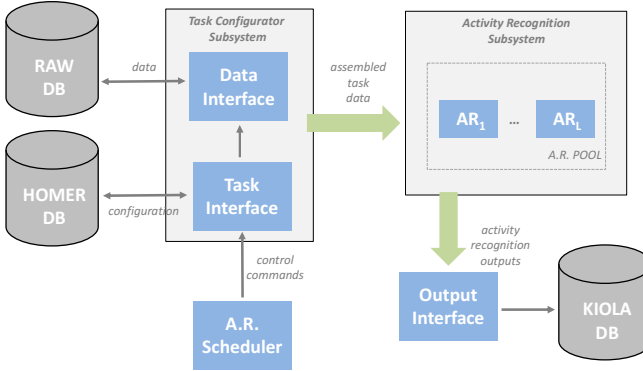


Fig. 2. Detailed architecture of the supervised HAR system

components (based on predictive learning models), one for each specific task. These components implement the trained predictive learning model obtained from a preliminary validation phase, and they produce their predictions by computing the outputs of the supervised learning model in response to the input data.

5.3 Reasoning Subsystem and Dashboard

Fig. 3 shows the Reasoner, the high-level database and the dashboard with the other system components. The high-level database receives the data from three sources: the activity recognition subsystem (data about physical activity, calories consumption, balance, user sociality etc.), the diet application on the tablet (nutritional data inserted by the users themselves), and the application for serious games on the tablet (statistics about the performance of the user in cognitive games).

The Reasoner compares all these data with the clinical protocol the person should follow based on the pre-sets from the medical experts. To this purpose, it adopts a rule based with hierarchical decision trees, where the rules will be created according to the actual medical guidelines. Based on this, a general overview, as well as some calculated data relations, is presented to the specialist's dashboard. The Reasoner settings can be modified by the medical experts to change the protocol according to a certain user behavior (for example he can change the food composition or reduce the overall caloric intake), or the Reasoner itself may adapt the protocol according to pre-defined rules, when some known conditions occur. For example, an improvement in the physical activity assessed by the heart rate response to exercise may result in a progressive increase in the intensity of proposed exercises. The Reasoner gives feedbacks to the user by means of applications on the user tablet (namely, the nutritional and physical activity advisors and the cognitive games).

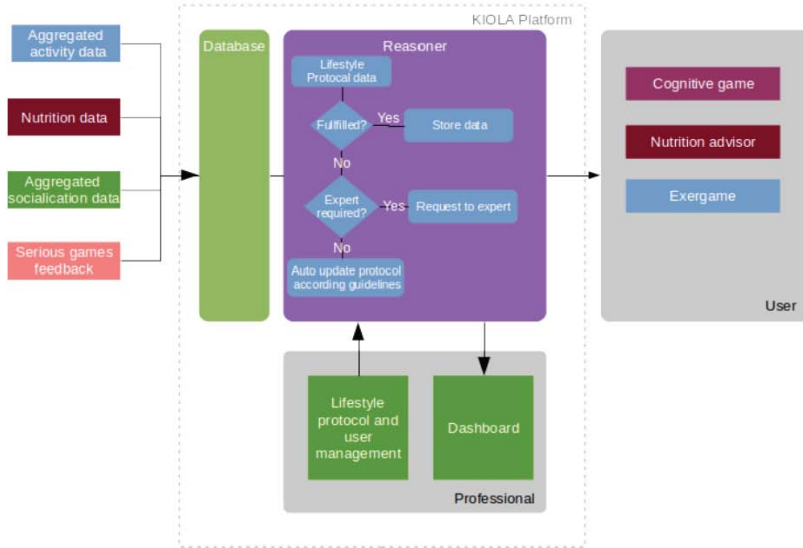


Fig. 3. Architecture of the Reasoner and its relationship with the other system components

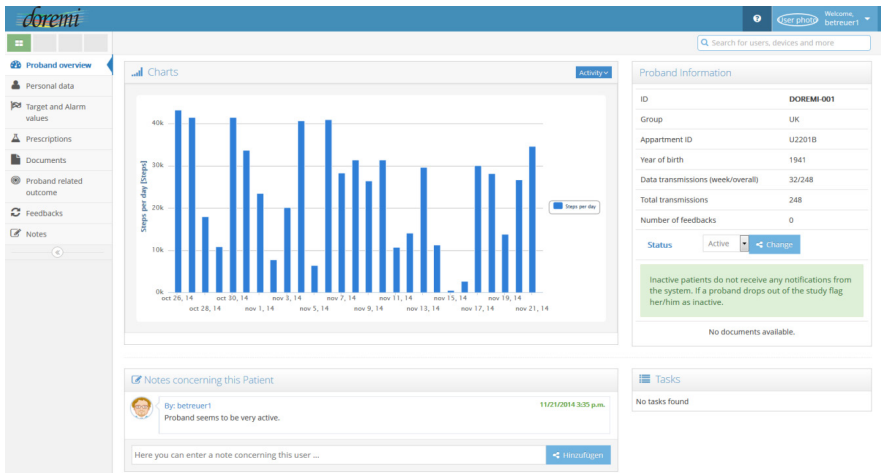


Fig. 4. Personal dashboard showing activity data of an end-user

The reasoning module and the dashboard are integrated in the KIOLA modular platform, suitable for clinical data and therapy management. It is built on top of the open-source web-framework Django¹ and uses PostgreSQL 9.4 as primary data storage. KIOLA has two groups of components: *core components* (that provide data

¹ <http://www.djangoproject.com>

models for receiving and storing external sensor data, rule-based reasoning on observations, and messaging services to communicate results of the reasoning to external systems), and *frontend components* (a dashboard for the specialists, an administrative interface, and a search engine for all data stored in KIOLA). In particular, the dashboard provides specialists with the possibility to review and adjust clinical protocols online, and it is designed for both mobile devices and computers. The dashboard can provide either an overview of all end-users to which the specialist has access, or a detailed view of a specific end-user. Charts are used to visualize all observations in the area of social, physical games, and dietary data (see Fig. 4). A task module on the dashboard is also used to notify specialists when the reasoning system suggests an adoption of the clinical protocol. Here, the specialist can then approve or disapprove the recommendation, and he can tune the parameters of the protocol by himself.

6 Conclusions

The DOREMI project addresses three important causes of morbidity and mortality in the elderly (malnutrition, sedentariness and cognitive decline), by designing a solution aimed at promoting an active aging lifestyle protocol. It envisages to provide an ICT-based home care services for aging people to contrast cognitive decline, sedentariness and unhealthy dietary habits. The proposed approach builds on activity recognition and reasoning subsystems, which are the scope of this paper. At the current stage of the project such components are being deployed, and they will be validated in the course of the year by means of an extensive data collection campaign aimed at obtaining the annotated datasets. These datasets, that are currently being collected over a group of elderly volunteers in Pisa, in view of the experimentation in the pilot sites planned by the beginning of 2016.

References

1. Tukey, J.W.: Exploratory data analysis, pp. 2–3 (1977)
2. Long, X., Yin, B., Aarts, R.M.: Single-accelerometer-based daily physical activity classification. In: Engineering in Medicine and Biology Society, EMBC 2009. Annual International Conference of the IEEE. IEEE (2009)
3. Fernández-Llatas, C., et al.: Process Mining for Individualized Behavior Modeling Using Wireless Tracking in Nursing Homes. *Sensors* **13**(11), 15434–15451 (2013)
4. der Aalst, V., Wil, M.P., et al.: Workflow mining: A survey of issues and approaches. *Data & Knowledge Engineering* **47**(2), 237–267 (2003)
5. Yang, C.-H., Liu, Y.-T., Chuang, L.-Y.: DNA motif discovery based on ant colony optimization and expectation maximization. In: Proceedings of the International Multi Conference of Engineers and Computer Scientists, vol. 1 (2011)
6. Bouamama, S., Boukerram, A., Al-Badarneh, A.F.: Motif finding using ant colony optimization. In: Dorigo, M., Birattari, M., Di Caro, G.A., Doursat, R., Engelbrecht, A.P., Floreano, D., Gambardella, L.M., Groß, R., Şahin, E., Sayama, H., Stützle, T. (eds.) ANTS 2010. LNCS, vol. 6234, pp. 464–471. Springer, Heidelberg (2010)

7. Cui, X., et al.: Visual mining intrusion behaviors by using swarm technology. In: 2011 44th Hawaii International Conference on System Sciences (HICSS). IEEE (2011)
8. Bao, L., Intille, S.S.: Activity Recognition from User-Annotated Acceleration Data. In: Ferscha, A., Mattern, F. (eds.) Pervasive 2004. LNCS, vol. 3001, pp. 1–17. Springer, Heidelberg (2004)
9. Lara, O.D., Labrador, M.A.: A survey on human activity recognition using wearable sensors. *Communications Surveys & Tutorials*, IEEE **15**(3), 1192–1209 (2013)
10. Kolen, J., Kremer, S. (eds.): *A Field Guide to Dynamical Recurrent Networks*. IEEE Press (2001)
11. Lukoševičius, M., Jaeger, H.: Reservoir computing approaches to recurrent neural network training. *Computer Science Review* **3**(3), 127–149 (2009)
12. Jaeger, H., Haas, H.: Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science* **304**(5667), 78–80 (2004)
13. Gallicchio, C., Micheli, A.: Architectural and markovian factors of echo state networks. *Neural Networks* **24**(5), 440–456 (2011)
14. Tino, P., Hammer, B., Boden, M.: Markovian bias of neural based architectures with feedback connections. In: Hammer, B., Hitzler, P. (eds.) *Perspectives of neural-symbolic integration*. SCI, vol. 77, pp. 95–133. Springer-Verlag, Heidelberg (2007)
15. Lukoševičius, M., Jaeger, H., Schrauwen, B.: Reservoir Computing Trends. *KI - Künstliche Intelligenz* **26**(4), 365–371 (2012)
16. Bacciu, D., Barsocchi, P., Chessa, S., Gallicchio, C., Micheli, A.: An experimental characterization of reservoir computing in ambient assisted living applications. *Neural Computing and Applications* **24**(6), 1451–1464 (2014)
17. Chessa, S., et al.: Robot localization by echo state networks using RSS. In: *Recent Advances of Neural Network Models and Applications*. Smart Innovation, Systems and Technologies, vol. 26, pp. 147–154. Springer (2014)
18. Palumbo, F., Barsocchi, P., Gallicchio, C., Chessa, S., Micheli, A.: Multisensor data fusion for activity recognition based on reservoir computing. In: Botía, J.A., Álvarez-García, J.A., Fujinami, K., Barsocchi, P., Riedel, T. (eds.) *EvAAL 2013*. CCIS, vol. 386, pp. 24–35. Springer, Heidelberg (2013)
19. Bacciu, D., Gallicchio, C., Micheli, A., Di Rocco, M., Saffiotti, A.: Learning context-aware mobile robot navigation in home environments. In: *5th IEEE Int. Conf. on Information, Intelligence, Systems and Applications (IISA)* (2014)
20. Amato, G., Broxvall, M., Chessa, S., Dragone, M., Gennaro, C., López, R., Maguire, L., McGinnity, T., Micheli, A., Renteria, A., O’Hare, G., Pecora, F.: Robotic Ubiquitous COgnitive network. In: Novais, P., Hallenborg, K., Tapia, D.I., Rodríguez, J.M. (eds.) *Ambient Intelligence - Software and Applications*. AISC, vol. 153, pp. 191–195. Springer, Heidelberg (2012)
21. Lavrac, N., et al.: Intelligent data analysis in medicine. *IJCAI* **97**, 1–13 (1997)
22. Chae, Y.M.: Expert Systems in Medicine. In: Liebowitz, J. (ed.) *The Handbook of applied expert systems*, pp. 32.1–32.20. CRC Press (1998)
23. Gurgun, F.: Neuronal-Network-based decision making in diagnostic applications. *IEEE EMB Magazine* **18**(4), 89–93 (1999)
24. Anderson, J.R., Machine learning: An artificial intelligence approach. In: Michalski, R.S., Carbonell, J.G., Mitchell, T.M. (eds.) vol. 2. Morgan Kaufmann (1986)
25. Hastie, T., et al.: *The elements of statistical learning*, vol. 2(1). Springer (2009)
26. Murphy, K.P.: *Machine learning: a probabilistic perspective*. MIT Press (2012)

27. Carvalho, D.R., Freitas, A.A.: A hybrid decision tree/genetic algorithm method for data mining. *Information Sciences* **163**(1), 13–35 (2004). [EDA1] Tukey, J.W.: Exploratory data analysis, pp. 2–3 (1977)
28. Fuxreiter, T., et al.: A modular platform for event recognition in smart homes. In: 12th IEEE Int. Conf. on e-Health Networking Applications and Services (Healthcom), pp. 1–6 (2010)
29. Kreiner, K., et al.: Play up! A smart knowledge-based system using games for preventing falls in elderly people. *Health Informatics meets eHealth (eHealth 2013)*. In: Proceedings of the eHealth 2013, OCG, Vienna, pp. 243–248 (2013). ISBN: 978-3-85403-293-9