Workshop Proceedings of the 8th International Conference on Intelligent Environments
J.A. Botía et al. (Eds.)
IOS Press, 2012

doi:10.3233/978-1-61499-080-2-4

Activity, Behavior and Context: The ABC of Pervasive Healthcare Research

Jesus FAVELA
Computer Science Department, CICESE, Mexico

Abstract. Pervasive healthcare aims at applying mobile and ubiquitous computing technologies to address some of the major issues facing healthcare, such as the aging of the population, an increase in the prevalence of chronic diseases and escalating costs. One of the main characteristics of pervasive healthcare systems is their reliance on contextual information such as using the location of a nurse in a hospital ward to provide her with information relevant to the task at hand, or reminding a patient to take his medication when he returns home and his pillbox has not been opened during the day. Advances in activity recognition have made possible a second generation of pervasive healthcare solutions using activity-aware computing to identify the activities of daily living being performed by an elder and assist her with the task, or to monitor the intensity of the physical activity performed by a patient undergoing treatment. Moving forward this paper describes a new wave of pervasive healthcare applications, which are becoming possible due to progress in the estimation of user behavior. These behavior-aware systems can be used to measure the effects of medication in patients suffering from mental disorders, provide feedback in rehabilitation therapies or to perceive abnormal behaviors that provide early signs of an ailments such as Alzheimer's disease.

Keywords. Behavior-aware computing, pervasive healthcare, activity recognition

Introduction

With the aging of the population, a significant increase in the prevalence of obesity and chronic diseases, and escalating costs, the provision of healthcare services faces major challenges in the coming years. These challenges demand patients to play a more active role in the management of their disease and rehabilitation process, and from the general population to be more conscious of their health habits and wellbeing. Recent advances in ICT, and particularly in mobile and ubiquitous computing, promise to address some of these issues by assisting in the monitoring of patients for disease management, providing opportune information to caregivers and healthcare professionals for timely treatment, and deploying the core infrastructure for people to have information that allows them to make healthier decisions and play a more active role in their own healthcare and wellbeing.

Pervasive healthcare refers to the application of pervasive or ubiquitous computer technologies for the diagnosis, treatment and prevention of disease, illness and other physical and mental impairments. This area has become one of the most active fields of research and application within ubiquitous computing. Advances in this field include smart home technologies that assist elders in performing activities of daily living [1]; mobile wellbeing applications that use the sensors in mobile phones to measure energy expenditure and persuade users to improve their eating habits [2, 3]; and pervasive hospital environments that provide physicians with the latest laboratory results when approaching the bed of a patient [4].

In this paper I describe the evolution of Pervasive Healthcare considering three main periods characterized by the underlying information inferred and used by these systems: Context, Activity and Behavior, arguably the main building blocks of pervasive healthcare systems. The first period is characterized by the use of contextual information directly measured from sensors or inferred from low-level data. An example of this first-generation of pervasive healthcare solutions is the use of location information in a hospital information system to provide hospital staff with information relevant to the task at hand. The second generation of pervasive healthcare research has been made possible by advances in activity recognition. Examples of activity-aware applications include assisting elders with activities of daily living or detecting activity periods to monitor the amount exercise performed by a user. These applications rely on the estimation of the user's activity, such as whether the person is walking versus running.

We are at the dawn of a new wave of pervasive healthcare solutions that can take advantage of behavior information inferred from users and communities. An example of this type of application is recognizing wandering behavior in a person suffering from dementia to notify a caregiver. I describe some potential healthcare applications of behavior-aware computing and some of the main technical challenges that need to be tackled in pervasive healthcare research in order to realize these solutions. Both activity and behavior can be considered context, and thus, applications that adapt to changes in these parameters are, strictly speaking, context-aware. Yet, there has been a clear evolution over the last few years from ubiquitous applications that use lower-level context, such as location, towards those that rely the recognition of activities, such as walking, and gradually towards applications that depend on the recognition of human behaviors, such as wandering.

1. Context-aware computing in healthcare

Context-aware computing refers to an application's ability to adapt to changing circumstances and respond based on the context of use [5]. A context-aware system has the capacity to perceive and capture the world that surrounds the user, and to adapt its behavior to provide information and services that are useful and relevant to that place and time [6]. In this sense, context is any information that can be used to characterize the situation of a given entity, which can be a person, place or object that is considered relevant for the interaction between a user and the application [6]. Context can be characterized at different levels of abstraction. Low-level context is deduced directly from sensors; were a 'sensor' refers not only to sensing hardware but also to every data source that provides usable contextual information, such as temperature, luminosity, sound, movement, air pressure, environmental humidity, user profile, preferences and schedule. Whereas high-level context is abstract and inferred from low-level information, for example, activity, location or user mood [7].

Context is difficult to use for several reasons [8]. First, capturing low-level contextual information requires the use of sensors and computing devices. Context must be abstracted to make sense for the application. For instance, the ID of a mobile user must be abstracted into the user's name or role. Finally, context is dynamic, i.e., a location-based system must update its display as the user moves, which require tracking the user's location by gathering information from multiple sensors, and using techniques that estimate the user's location or guess the route that the user will follow.

The design and development of context-aware applications faces two major challenges. First, it is important to identify which information is worth capturing and how contextual information could be automatically detected to support the user's tasks. Second, it is essential to provide adequate methods for the presentation of information and services in support of the user's ongoing activities.

The incorporation of robust location estimation technology in mobile phones has made location-based systems ubiquitous. Location information can be used to provide services and information to mobile users. For instance, in a hospital setting where highly mobile hospital staff executes activities depending on their location, location-based systems can provide opportune information, such as, access to patient's records when the user is near to the patient's bed or inform the user of the location of an electrocardiograph.

Context-aware computing has been applied to a variety of healthcare environments, such as supporting communication among hospital workers [9, 10], aiding patients to adhere to their medication regimes [11], or alerting about risks of drug-drug interaction [12].

To illustrate the use of context-aware computing in healthcare consider the smart hospital environment (iHospital) described in [13, 14]. The iHospital is a vision of a highly interactive workplace, where hospital staff can access relevant medical information trough a variety of heterogeneous devices, and can collaborate with colleagues taking into account contextual information. To realize this vision several pervasive healthcare applications were developed that altogether use contextual information to: provide awareness of the location of hospital staff and artifacts [10]; support collaboration and opportunistic encounters through context-aware communication [15] (see Figure 1); adapt and personalize information to the user's or artifact's context (e.g., role, location, laboratory studies status) [15]; and support the peripheral monitoring of patients [17].

Key to the development of the iHospital is accurate and inexpensive indoor location estimation. This was achieved using a backpropagation neural network that maps the strength of radiofrequency signals received by mobile devices to the location of the user carrying the device [18]. The approach uses information about the neighbors surrounding the location of the user to simulate previous time instant guesses and alleviate the hopping trajectories of users. The results obtained using this approach average 1.3m distance error during continuous motion, which favorably compares with the 2m average error obtained when not using information of neighbor locations.

Figure 1. A medical intern receives a context-aware message and collaborates with a nurse on a public display.

2. Activity-aware computing for healthcare

Activity-aware technologies aim at allowing "smart environments" to respond proactively to the needs and intentions of their users by automatically inferring the activity being performed by the user [19]. Activity-aware applications thus need intelligent capabilities to be adaptive and reactive to users' tasks and learn from user's actions in order to provide opportune and reliable services based on their goals.

Human activities are complex and dynamic. Activities at different levels of granularity can be performed concurrently and they might be interwoven. Hence, the development of activity-aware applications requires appropriate representations of computational activities that are relevant to the services provided by the system. An additional challenge for the development of activity-aware computing is the development of robust approaches for activity recognition.

COACH (Cognitive Orthosis for Assisting with aCtivites in the Home) is an example of an activity-aware application aimed at assisting in the completion of activities of daily living [1]. The system aims at assisting people who suffer from dementia perform tasks, such as handwashing. The system uses a video camera to identify the position of the hands and task objects such as the soap or towel, to infer the state of the activity and provide appropriate prompting. A planning system determines the current step of the activity and decide the appropriate response of the system. Knowledge about the user's abilities is used to adapt the prompting.

Activity recognition is currently one of the more active areas of research in ubiquitous computing. Numerous projects have explored algorithms, methods and devices for the automatic recognition of users' activities. These approaches to inferring people's activities range from estimating low-abstraction activities (such as chewing or sitting) to estimating more complex activities such as whether the person is leaving for work later than usual. For instance, detecting the manipulation of certain objects using RFID tags has been used to estimate activities of daily living (ADL), such as preparing a drink or taking a medication [20]. The SEER (Selective Perception Architecture for Activity Recognition) system used a hidden markov model (HMM) to estimate activities at a higher level of abstraction (such as whether a user is attending a meeting) [21]. The inputs to SEER include audio, video, and computer interactions captured through sensors distributed in an office environment.

One of the trends pushing activity recognition research is the significant increase in available data, due to the incorporation of sensors such as GPS, accelerometers and microphones, in mobile phones. The ubiquity of these devices also makes them ideal platforms for the deployment of pervasive healthcare solutions. Relevant examples to healthcare include the use of accelerometer data to reliably determine if a person is standing, walking or running, and to estimate energy expenditure [22]. The automatic estimation of these activities is relevant for wellness applications that measure periods of moderate or intense activity. In addition, the audio captured by microphones in mobile devices can also be used to detect the semantic location of a person (e.g. at a restaurant, on a bus) and infer activities such as “eating”, “transportation”, “conversation”, etc. This field of research known as computational auditory scene recognition can also be used to measure the intensity and frequency of social interaction [23, 24]. Data from multiple sources can be used to achieve higher recognition precision.

Decision trees and Hidden Markov models have been used to combine evidence from heterogeneous sources. Figure 2. shows an example of a two-level HMM used to recognize hospital staff performed activities [25]. The model was trained with data gathered from close to 200 hours of observation of hospital activities and interactions. The data includes information about the location of the hospital worker, the artifacts s/he uses to perform a task, and the roles of the people with whom s/he interacts (e.g., nurse, patient) [26]. The model achieved an accuracy of 92.6% distinguishing among 6 different activities such as patient care, clinical case assessment, and information management. Interestingly, the HMM attained better performance than human experts who were asked to recognize the activity with the same data.

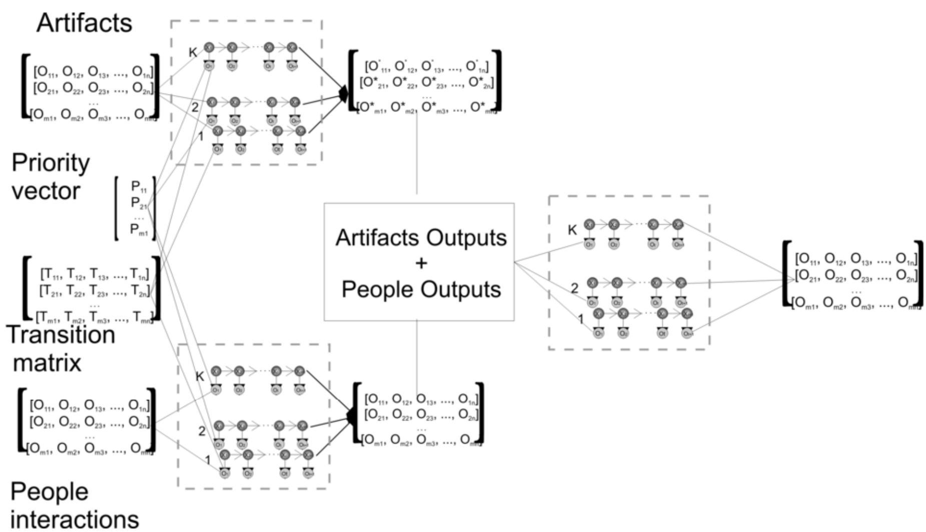

Figure 2. Architecture of HMM used to estimate the activities performed by hospital staff [25].

Different activity-aware applications tolerate different recognition accuracies for providing adequate activity-aware responses. An approximate estimation of energy expenditure using accelerometer data might be more than appropriate if the alternative is self-report on the amount of daily exercise performed by a user. Yet, some activities might be more difficult to classify than others. For instance, an accelerometer worn at the waist greatly underestimates activity counts during cycling. With regards to the recognition of the activities performed by hospital workers an application might be more tolerant to errors while inferring the current availably of a nurse in order to decide whether to interrupt her or no, than to errors registering that a nurse gave a medication to a patient. The accuracy achieved by the HMM described above might be adequate for the first application, but not for the second one. Thus, the selection of the activity recognition strategy should be closely associated to the design of the activity-aware application.

3. Towards behavior-aware computing

Behavior is the response of a system or organism to various stimuli in its environment. In sociology, behavior is considered as not being directed at other people, thus being the most basic of human action. Analyzing behavior is relevant to healthcare since different pathologies are a consequence of behavior or are shaped by it. For instance, unhealthy eating habits and/or lack of exercise lead to obesity, which in turn might lead to hypertension and diabetes. On the other hand, mental disorders such as Alzheimer's disease, Parkinson's disease or autism are characterized by behaviors such as wandering, hand tremor, or failure to socialize. Indeed, the recognition of these behaviors often leads to their early detection and diagnosis.

Over the last few years there have been important advances in automatic behavior recognition. ALEX Pentland has explored the use of "honest signals" that can to reliably predict human behavior [27]. These signals, which are often perceived unconsciously by a person, can be reliably estimated using sensors such as accelerometers and microphones. He has proposed four honest signals for which he has found significant correlation with a variety of behaviors. These signals are: *mimicry*, the reflective copying of one person by another; *activity*, the movement that indicates interest and excitement; *influence*, the extent to which a person matches the other's speech; and *consistency*, the fluidity of speech and motion, which is perceived by others as expertise [27]. Additional research has explored the feasibility of recognizing other behaviors relevant to healthcare such as stress [28], mood [29], or personality traits [30]. Conscientiousness, for instance, has been found to be linked to longevity and to have wide-ranging effects on health-relevant activities [31]. Thus, a behavior-aware medication reminder might be designed to use different persuasion mechanisms for conscientious patients than for a user who tends to be careless.

In the rest of this section I describe some of the challenges facing the development of behavior-aware computing and potential applications in healthcare and wellbeing.

3.1. Challenges in behavior-aware computing

The research challenges of behavior-aware computing are similar to the ones faced in the development of context and activity-aware applications. These are: indentifying meaningful behavior-aware applications; representing behavior information; determining what behavior information is relevant to the application and how it should adapt to changes in behavior; and recognizing behavior. Behavior-aware computing research can be motivated by application requirements (top-down), such as detecting problematic or abnormal behavior in persons with dementia to notify their caregiver, or by technological advances (bottom-up) that make feasible inferring certain types of behavior.

3.2. Applications of behavior-aware computing in healthcare

There are numerous healthcare applications for systems that can adapt to changes detected in the behavior of the user.

3.2.1. Detecting influenza epidemics

A study that collected data from 70 students in a University campus for a period of 4 months provides evidence that data obtained from mobile phones could be used to detect changes in behavior associated to symptoms of influenza [32]. In particular, subjects that suffered symptoms of influenza experienced lower activity and entropy levels and a decrease in communication frequency at different times of the day. Detecting these changes in behavior could be used to detect when a person has a flue.

Understanding the dynamics of social networks can also be used to predict the spread of diseases. The Reality Mining project has shown how social relationships can be automatically uncovered with information gathered from the sensors in mobile phones [33]. Research being conducted to predict mobility patterns and known locations can be used by public health authorities to implement appropriate responses to epidemic breaks.

3.2.2. Risky or abnormal behaviors: Wandering

It has been estimated that around 60 percent of those suffering from dementia wander at some point, and half of those not found within a day suffer serious injuries or death. Wandering behavior is unique to a given person in a particular context or situation [34]. The factors that affect wandering include those related to the background of the person such as health status, personality, cognitive decline, and socio-demographic characteristics, as well as factors related to the context of the situation including personal and physical needs and social aspects of the environment [34].

A straightforward approach to detect wandering behavior would be to recognize random walking from location data. This would work for some instances of wandering. However, some of the most problematic episodes of wandering are due to direct walking, when the wanderer aims to the place were they used to work or live, without feeling lost nor disoriented.

LaCasa (Location and Context-Aware Safety Assistant) is an early prototype of a system that incorporates background and contextual information to detect wandering behavior and provide appropriate assistance to the person with dementia and/or her caregiver [35]. The system is based on a partially observable Markov decision process, or POMDP, a stochastic model of a temporal dynamic process that can handle partial obsevability. The model includes a set of preferences encoded as utilities for the persons involved, and computes a policy of assistance, such as prompting the user or calling the caregiver, that optimizes over these preferences given the state and the dynamics of the world in the long term. The data used by the POMDP is obtained thru opportunistic sensing using a mobile phone (e.g. known locations and common routes), as well as background information gathered from interviews with the person with dementia and her caregiver.

3.2.3. Early diagnosis of dementia or frailty

Frailty is as a multidimensional syndrome, which involves loss of physical and cognitive reserves and ultimately an increase of vulnerability [36]. Frailty is commonly assessed with a collection of instruments and tests. Several of these tests involve self-reporting to questions such as: "during the last week how many days you walked continuously for at least 10 minutes?" or "how frequently you speak with your friends?". More accurate responses to these questions could be obtained from the analysis of data obtained from the sensors available in mobile phones. For instance, voice activity detection could be used to assess the strength of the social network of an older adult or the nature of a conversation [24]. Further analysis of the audio signal and the amount of time each participant talks could be used to infer depression or anxiety.

Fatigue is another factor associated to frailty. Unobtrusively monitoring gait speed over a period of time, for instance in a route frequently followed by the user, could provide early evidence of fatigue. In [37] the use of ambient videogames is proposed to measure muscle fatigue, which has been associated with sensation of fatigue, which is considered a predictor of disability and even mortality.

4. Conclusions

Inferring contextual information and recognizing the activity being performed by a person have been two of the cornerstones of ubiquitous computing research. Numerous advances in these areas have been motivated by applications in healthcare and wellbeing, giving rise to the field of Pervasive Healthcare. A new wave of pervasive healthcare research is being made possible by the ubiquity of mobile phones that incorporate sensing technology and advances in behavior recognition. Thus, making behavior, along with context and activity, the third building block of pervasive computing research.

References

[1] J. Boger, J. Hoey P. Poupart, C. Boutilier, G. Fernie, A. Mihailidis. A planning system based on Markov decision processes to guide people with dementia through activities of daily living. *IEEE Trans Inf Technol Biomed*. 10(2):323-33. 2006.

[2] K.H. Connelly, A.M. Faber, Y. Rogers, K.A. Siek, T. Toscos. Mobile applications that empower people to monitor their personal health. *e & i Elektrotechnik und Informationstechnik*, 123(4):124–128, 2006.

[3] K.A. Siek, *et al*. When Do We Eat? An Evaluation of Food Items Input into an Electronic Food Monitoring Application. *In Pervasive Healthcare 2006*, pp. 1–10, 2006.

[4] M.D. Rodríguez, J. Favela, E.A. Martínez, M..A. Muñoz. Location-aware access to hospital information and services. *IEEE Transactions on Information Technology in Biomedicine* 8(4): 448-455. 2004.

[5] A.K. Dey. Understanding and Using Context. *Personal and Ubiquitous Computing* Vol. 5, No.1, 4-7. 2001
[6] G.D. Abowd, M., Ebling, G., Hunt, H., Lei & M-H. Gellersen, H.-W. Context-aware computing. *IEEE Pervasive Computing*, 1(3), 22-23. 2002.
[7] P. Prekop, M. Burnett. Activities, Context and Ubiquitous Computing. *Computer Communications*, Vol 26, No. 11, 1168-1176. 2003.
[8] A.K. Dey. Providing Architectural Support for Building Context-Aware Applications. PhD. Thesis. Georgia Institute of Technology, 2000.
[9] J.E. Bardram. Applications of context-aware computing in hospital work: examples and design principles, In *Proc. of the ACM symposium on Applied computing*. 2004, pp. 1574 - 1579.
[10] M. Munoz, *et al*. Context-aware mobile communication in hospitals. *IEEE Computer*. Vol. 36, No. 8, pp. 60–67. 2003
[11] D. Asai, *et al.* Context-aware reminder system to support medication compliance. *Proc. IEEE SMC*, Anchorage, Alaska 2011: 3213-3218
[12] J.D. Duke and D. Bolchini. A successful model and visual design for creating context-aware drug-drug interaction alerts. *AMIA Annu Symp Proc*. 2011;2011:339-48.
[13] J. Bardram, H. Baldus, J. Favela. Pervasive Computing in Hospitals, Chapter No. 3, In *Pervasive Computing in Healthcare*. Bardram, J., Mihailidis, A., Wan, D. (eds). CRC Press. Pp. 49-77. 2006.
[14] J. Favela, A.I. Martínez, M.D. Rodríguez, V.M. González. Ambient Computing Research for Healthcare: Challenges, Opportunities and Experiences. *Computación y Sistemas*.Vol. 12, No. 1, pp. 109-127. 2008.
[15] A.D. Mejía, J. Favela, A.L. Moran. Understanding and Supporting Lightweight Communication in Hospital Work. *IEEE Trans. on Information Technology in BioMedicine*, Vol. 14, No. 1, pp 140-146. 2010.
[16] J. Favela, M.D. Rodríguez, A. Preciado, V.M. Gonzalez. Integrating Context-aware public displays into a Mobile Hospital Information System, *IEEE Trans. on Information Technology in BioMedicine*. 2004. Vo., 8, No. 3. pp. 279-286.
[17] M. Tentori, J. Favela, D. Segura. Chapter 8: Monitoring hospital patients using ambient displays. In *Mobile Health Solutions for Biomedical Applications*. Medical Information Science Reference. 2009.
[18] L.A. Castro and J. Favela. Reducing the Uncertainty on Location Estimation of Mobile Users to Support Hospital Work. *IEEE Transactions on systems, man and cybernetics--Part C*, Vol. 38, No. 6: 861-866. 2008.
[19] M. Tentori, J. Favela. Activity-aware computing for healthcare. *IEEE Pervasive Computing*, 2008. 7(2): p. 51-57.
[20] M. Philipose, K.P. Fishkin, M. Perkowi. Inferring Activities from Interactions with Objects, *IEEE Pervasive Computing*, vol. 3, No. 4, 2004, pp. 50–57.
[21] N. Oliver, A. Garg, E. Horvitz. Layered Representations for Learning and Inferring Office Activity from Multiple Sensory Channels, *Computer Vision and Image Understanding*, vol. 96, no. 2, 2004, pp. 163–180.
[22] M. Mathie, A. Coster, N. Lovell, B. Celler. Detection of daily physical activities using a triaxial accelerometer. *Medical and Biological Engineering and Computing*. Vol. 41, No. 3, 296-301. 2003.
[23] V. Peltonen, *et al.*, Computational Auditory Scene Recognition, In *IEEE International Conference on Acoustics, Speech, and Signal Processing*. - Orlando, Florida, USA ,2002. - May 13 - 17.
[24] D. Zhang, et al., Modeling Individual and Group Actions in Meeting with Layered HMMs, *IEEE Trans. Multimedia*, vol. 8, no. 3, 2006, pp. 509–520.
[25] D. Sanchez, J. Favela, and M. Tentori, Activity recognition for the Smart Hospital. *IEEE Intelligent Systems*, 2008. 23(2): p. 50-57.
[26] J. Favela *et al.*, Activity Recognition for Context-Aware Hospital Applications: Issues and Opportunities for the Deployment of Pervasive Networks, *Mobile Networks and Applications*, vol. 12, no. 2-3 2007, pp. 155–171.
[27] A. Pentland. Honest signals: how social networks shape human behavior. In *Proc. ACM Multimedia 2011: 583-584*. 2011
[28] J. Hernandez, R.R. Morris, R.W. Picard. "Call Center Stress Recognition with Person-Specific Models," In Proceedings of the Affective Computing and Intelligent Interaction, Memphis, October 9-12, 2011.
[29] A. Matic, A. Papliatseyeu, A. Gabrielli, V. Osmani, O. Mayora-Ibarra. Happy or moody? why so? In: Proc. Ubi-Health (2010)
[30] G. Chittaranjan, J. Blom, D. Gatica-Perez. Who's Who with Big-Five: Analyzing and Classifying Personality Traits with Smartphones. In Prox. of ISWC 2011: 29-36
[31] H.S. Friedman, *et al*. Childhood conscientiousness and longevity: Health behaviors and cause of death. Journal of Personality and Social Psychology, Vol 68(4): 696-703. 1995.

[32] A. Madan, M. Cebrián, D. Lazer, A. Pentland. Social sensing for epidemiological behavior change. In *Proc. UbiComp 2010*: 291-300. 2010

[33] N. Eagle, A. Pentland. Reality mining: sensing complex social systems. *Personal and Ubiquitous Computing* 10(4): 255-268. 2006.

[34] D. Algase, C. Beel-Bates, E., Beattie. Wandering in Long-Term Care. *Annals of Long-Term Care: Clinical Care and Ageing* 2003. Vol 11, No. 1, pp 33-39.

[35] X. Yang, J. Hoey, E. Quintana, and J. Favela. LaCasa Location and Context-Aware Safety Assistant. *Pervasive Healthcare 2012*. San Diego, CA., May 2012.

[36] K. Rockwood, *et al*., A global clinical measure of fitness and frailty in elderly people, *CMAJ*, vol. 173, no. 5, pp. 489-495, 2005.

[37] I. Zavala-Ibarra and J. Favela. Ambient videogames for assessing muscle disease related to aging. *Pervasive Healthcare 2012*. San Diego, CA., May 2012.