The GrowMeUp Project and the Applicability of Action Recognition Techniques

Gonçalo S. Martins, Luís Santos, and Jorge Dias

Institute for Systems and Robotics, University of Coimbra
3030-790, Coimbra, Portugal
{gmartins,luis,jorge}@isr.uc.pt
http://ap.isr.uc.pt/

Abstract. In recent years, an increased life expectancy and decreased birth rate have contributed to the phenomenon known as population aging, resulting simultaneously in an increased number of elderly people and in a decreasing number of young caretakers.
In this work, we will introduce the GrowMeUp research project. This project aims to build a scientific and technological solution able to assist the elderly at home, delaying the need for specialized full-time care. Its success will rely on, amongst other key innovations, a highly adaptable services paradigm, within which we briefly discuss the importance of action recognition as key context information.
The system will create a model of its user and will adapt its interaction accordingly. The knowledge of such model is updated from the robot-human interaction, through the observation of the elderly person's routine and also by sharing information with other robots that are connected to a cloud of knowledge.
We will discuss the motivation and context surrounding the project, the underlying software and hardware architecture, its scientific goals, its implementation strategy and how Action Recognition, which will be the focus of a significant part of future work, will be used in the context of the project.

Keywords: Adaptable Services, Assisted Living, Action Recognition, Cloud Computing, Research Project

1 Introduction

Population aging, *i.e.* the tendency of the overall age of a population to grow, resulting in a larger number of senior citizens, is a well-documented phenomenon prevalent in most, if not all, developed countries [3].

This phenomenon has been fueled by several causes. First of all, advances in medical technology allow for the early diagnosis and effective treatment of diseases, to the point where it is now possible to live an almost normal life carrying

This work was developed in the context of the GrowMeUp project, funded by the European Union's Horizon 2020 Research and Innovation Programme - Societal Challenge 1 (DG CONNECT/H) under grant agreement N° 643647.

very serious diseases [15]. Ailments that were previously almost certainly fatal have today a very significant chance of cure, given an early diagnosis. The evolution, refinement and widespread use of hygienic techniques has also resulted in a reduction of the impact of infectious diseases [21]. In extreme cases, it is possible to contain and treat patients suffering from extremely contagious diseases without further spread of the pathogen [9].

These advances have led to our current never-before-seen life expectancy, stretching to around 80 years [18], making it possible to live longer, healthier lives with higher quality standards. Today's elderly citizens have the best, most effective medical techniques, procedures and chemicals ever developed working in tandem to improve and maintain their health.

Life expectancy itself, however, is not solely responsible for the generalized aging phenomenon; an equivalent increase in birth rates would probably be able to maintain a balanced age distribution among the population, but despite the recent unyielding growth of life expectancy, birth rates have been declining severely [13], most noticeably in the latest years. The combination of these factors leads to a growing deviation in population age distribution, resulting in the gradual worsening of the population aging phenomenon.

Senior citizens suffer from a myriad of specific problems, such as illnesses exclusive to – or more prevalent on – old age. But the challenges faced by this section of society stretch beyond those of physical health. Social deprivation, either due to the death of the members of their social circles, gradual loss of contact, geographical isolation, and others, is a particularly prevalent issue, leading to feelings of loneliness and depression.

As a person ages, and their physical and psychological capabilities diminish, the need for constant monitoring of the person tends to increase. In most cases, the person must be moved from their preferred environment in order to ensure their safety and comfort. This leads to the loss of the person's independence, since they must either move in with a family member or into a caretaking institution.

These issues threaten the person's mental stability, self-image and general quality of life, and are some of the issues tackled within this project.

Given how technology is now an inherent part of human life and indissociable from our daily routines, a technological solution to these problems would be most welcome; this is the main goal of the GrowMeUp research project. In the context of this project, we are interested in applying and further developing Human Behavior Analysis techniques, such as Action Recognition and Emotion Recognition, striving for a technological solution for the problems faced by the older portion of the population.

The GrowMeUp solution will be embodied in a previously-developed robotic platform [2] designed specifically for usage with elderly people, striking a balance between functionality, aesthetics and cost. This platform will be the main enabling technology behind the project, and the target of most of the development that will take place therein.

The project aims to advance the state-of-the-art in Social Robotics, and to build a system able to provide assistance and companionship to elderly people, and also assurance to their family members and caretakers. This system will allow the elderly person to live for longer in their own environment without losing contact with their social circles, staying socially active either via teleconference or through the social facilities provided by the robot itself. Additionally, the behavior analysis components of the system will allow it to detect and report any emergencies, to respond to the person's non-verbal cues and commands, and to provide the level of monitoring needed. The system will provide support and encouragement as the person deals with daily life.

The system will be modular by design and will evolve with usage, progressively adapting to its user's needs. A cloud infrastructure will provide for constant communication between the several platforms deployed at a given time, which will share data amongst themselves with the goal of augmenting and improving the services the each are able to render. The robot will be deployed as an "infant" and will "grow" as time goes on. As the user's capabilities diminish, the robot's abilities and knowledge will improve, enabling it to provide a wider range of personalized services.

GrowMeUp is neither the first effort in this area, nor the first solution aimed at providing better living conditions for the elderly. Projects such as CompanionAble [4], ASTRO [5], Accompany [26], MobiServ [17], SRS [20], among others, have acknowledged the issue at hand, and have each contributed their own solution, or part thereof.

Despite their usefulness and innovative value, these generally revolve around the provision of physical services, such as monitoring and physically helping the person perform a particular task. Alternatively, the GrowMeUp project will focus on the *social* aspect of the senior citizen's daily life, providing encouragement, guidance and support instead of physical help.

Other projects, such as GiraffPlus [8] and VictoryaHome, have created integrated AAL solutions where a network of sensors is deployed to monitor the user, and where a telepresence robot is used to enable the user to communicate with their social circles. These previous works are of the utmost relevance, and will be taken into account in the development of our own system. Our solution differs from these previous works in mainly two points: firstly, our solution intends to create a robot that is not only capable of telepresence, but that *itself* becomes a member of the person's social circle; secondly, our goal is not only to develop a solution, but to develop one that evolves and adapts to its user's needs, as we will see in the following sections.

The GrowMeUp project aims to innovate in five key areas: behavior understanding using context analysis; intelligent dialogue management; decision making and learning using contextual information; ICT-based services for assisted living and service robots; usage of a cloud-based knowledge base for increased reliability and performance in service robotics.

In the following sections, we will discuss the scientific background that sustains the GrowMeUp project, as well as its main goals, techniques and challenges.

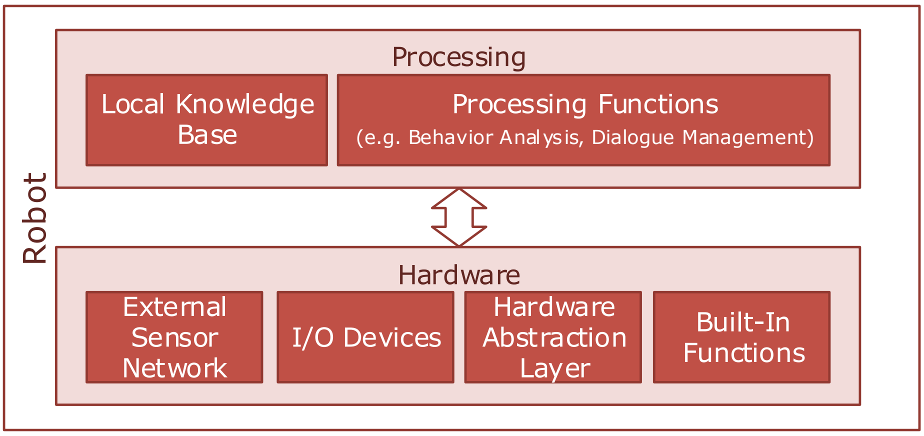

Fig. 1. A simplified pictorial representation of the robotic system.

2 The GrowMeUp Solution

2.1 A Learning Robot

As the aging process progresses, it is predictable that any given person will start, sooner or later, showing signs of light physical and cognitive difficulties. These gradually increasing difficulties give rise to changes in the needs and behavior of elderly people. As their problems progress, the support that the persons are given has to evolve as well.

This is currently achieved by caregivers as part of routine interaction with the patients. By informal conversation, such as asking how they have been, and by observation, for instance by checking if all the medication has been taken according to prescription, caregivers assess the current condition of the person, and adapt their actions accordingly.

The GrowMeUp approach will aspire to alleviate the burden of these caregivers by providing a technological solution that can monitor and interact with the elderly person at any time. Like the caregivers, who learn about the person's condition through interaction and observation, so will the robot be able to interact and observe the person, learning and adapting as they change.

Once developed, the system will be deployed with little knowledge of its user and environment. As time progresses, the robot will interact with its user both actively, for instance by asking questions, and passively, by observing, gaining knowledge with each interaction. The robot will observe its user, analyzing their motion and emotional cues, and will ask questions such as "What are you doing?" or "Have you taken your medication?" when pertinent. Through this interaction, and through information entered by care professionals by other means, the robot will build a profile of its user that it will take into account when making decisions. Each interaction with the user has the potential to provide information that the robot can use to refine its knowledge of the user, thus causing it to learn and evolve as it interacts.

This intended functionality will mainly be supported by three scientific fields: Behavior Analysis, Dialogue Management and Machine Learning.

Behavior Analysis is a very active field of research, with a substantial amount of recent work, and focuses mainly on the extraction of useful information from the behavior of people or other entities. This field deals with problems such as the recognition of actions present in sensor data, or the analysis and measurement of the characteristics of social interaction [19]. The application of behavior analysis techniques in this context will constitute one of the greater scientific challenges of the project, as we will see in Section 3.

Machine Learning acts as the supporting framework of many action recognition techniques, and is also one of the basic concepts of the field of Artificial Intelligence. In the context of the project, Machine Learning techniques will be used to enable our solution to incorporate new data as it is received and analyzed into the several models it will maintain, be it models of the person's behavior, of their routine, of the environment the robot is embedded in, of the person's favorite objects, *etc.* These techniques will be at the heart of the solution.

The robot is also expected to be able to conduct itself in a perceivably intelligent conversation, namely for acquiring data on its user. For instance, the robot may know the person's medical prescriptions, and should be able to verbally query the user as to its fulfillment, *e.g.* by asking "Have you taken you heart pill today?" and analyzing the answer. The analysis of the answer will go beyond the recognition of the words employed; the robot should be able to "suspect" the person by correlating the data gathered by the speech recognition system with the person's remaining status, such as emotional condition and non-verbal cues. In this aspect, the project will develop a new hierarchical and agent-based dialogue management paradigm, exploiting the behavior understanding models and contextual information used by the robot, as well as previous interactions with the user, to further support and enrich the dialogue. This issue has been the focus of a significant research effort throughout the years, and some solutions exist already, namely those based on Finite-State Machines [14] or, more recently, Hierarchical Task Networks [28].

By combining, and improving on, the knowledge from these fields, we will develop, implement and demonstrate the system summarily described in Fig. 1. This system will be composed, essentially, of a robotic platform, described in Section 2.4, and a software system that will incorporate a local knowledge base, a set of services implemented in a Service-Oriented Architecture [24], and a cloud infrastructure; the latter will be described in the following sections.

2.2 Service Adaptability and Context Awareness

The functionality provided by the GrowMeUp solution will be based on the concept of *service*. A service, in the context of the GrowMeUp architecture, is a set of instructions related to the fulfillment of a given expectation. For example, the "fall prevention" service would include instructions for the robot to stay

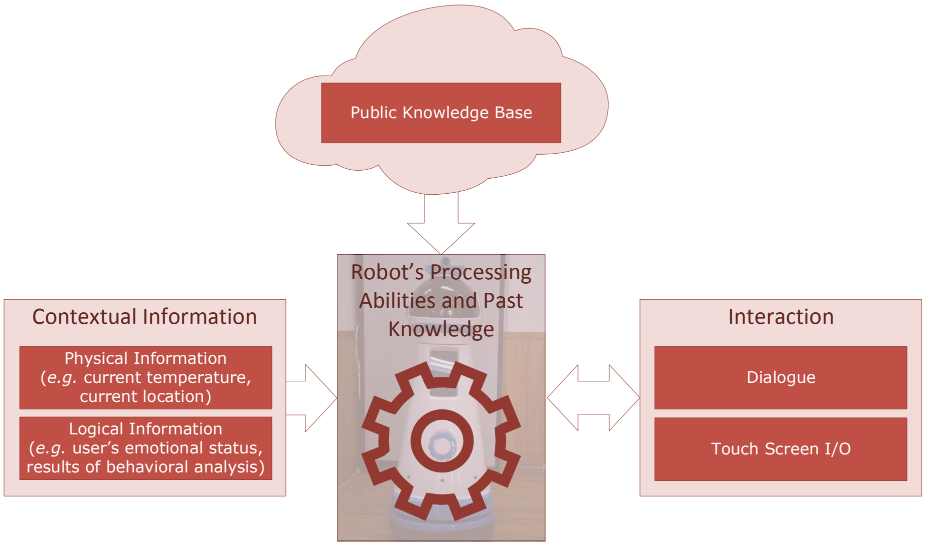

Fig. 2. A simplified pictorial representation of the various sources of data, and how it flows through the system. Data is received from both the environment and from the cloud, which is then used to refine the robot's model of its environment and user. Interaction with the user is also an important source of information.

vigilant of how the person is walking, if they are doing so with visible difficulty, and to be highly alert for signs of falling, so that it can then call for help. The services will be encoded in XML[1], so that they are both human- and machine-readable.

Context analysis will be one of the system's main features, and the target of a significant portion of our future research effort. Specifically, the system should be able to *adapt its services* based on the context that surrounds it.

To achieve this goal, new context-aware behavior models will be developed within the project. These models will take into account all the contextual information deemed relevant, which will include both physical and logical information, with physical information being that which concerns solely the physical world, such as the current temperature, the robot's localization and the objects that surround it, and logical information being that which concerns the user's status, such as their emotional status, current ailments and possible motivations, as briefly summarized in Fig. 2.

To further extend the functionality of the robot, we will make use of a localization technology, such as the iBeacon technology, to enrich the contextual information available to the robot. Beacons will be placed on certain objects, which will then be of particular interest to the robot. These beacons can then

[1] *Extensible Markup Language*, a widely-used format for the storage of data in both human- and machine-readable format. More information can be fount at `https://en.wikipedia.org/wiki/XML`.

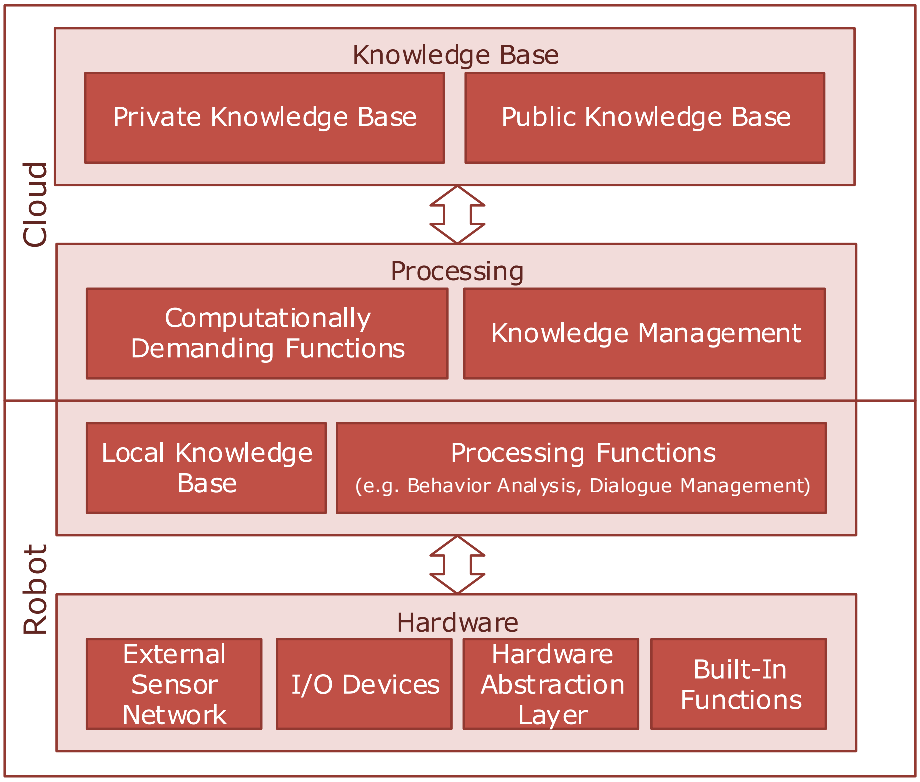

Fig. 3. A simplified pictorial representation of the GrowMeUp architecture, and of the way processing power and data are distributed through the system.

be used, for example, for the robot to know the location of some of the user's belongings, enabling it to remind the person of where they are.

The development of these models, from the specification of the characteristics of the environment that are indeed relevant, to their actual implementation and demonstration, will also pose a significant challenge, and will be the focus of a significant portion of the scientific work performed within the project, as illustrated in Section 3.

2.3 Information Sharing Using a Cloud Infrastructure

As mentioned before, we plan to deploy the robot into to the user's home with little to no previous knowledge. As the robot interacts with its user, it "grows", learning about its user and its environment. This paradigm is inspired by the way a human child learns during its very first years of life, knowing very little when taken home for the first time, and then learning by interacting with adults and other children.

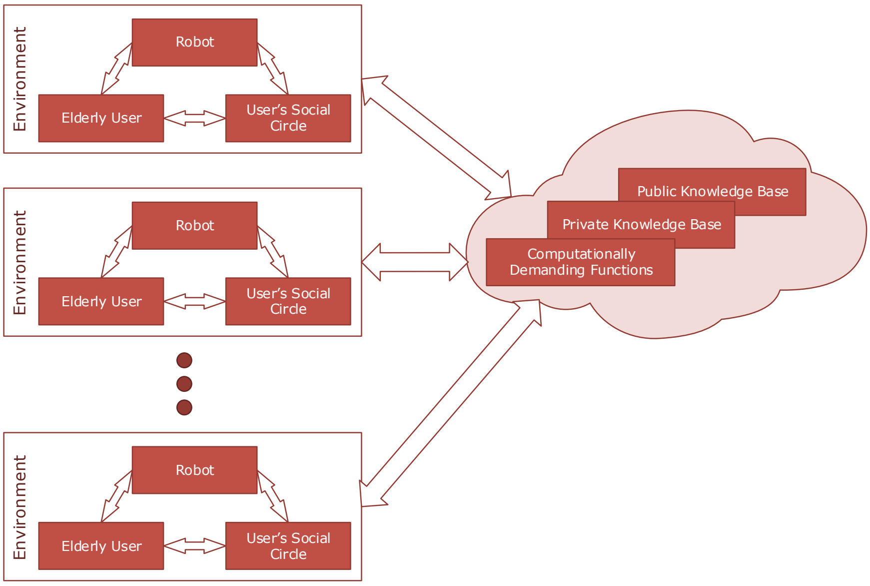

Fig. 4. A simplified functional diagram of the full GrowMeUp solution, and of the many communication channels that exist throughout the system.

The interaction with other children is of particular inspiration for the cloud component of the system. While the robots' amount of knowledge on deployment is relatively small, the robots will be connected to each other through a cloud infrastructure, sharing knowledge and potentially growing together.

As robots gather knowledge on their users and environment, they will build a local, private knowledge base, of which a copy will be kept on the cloud. They will also assist in the creation of a public knowledge base, containing information that can then be used by other robots, deployed in the same or in other environments, to enrich their own knowledge and to assist them in the fulfillment of their duties. This public knowledge base will serve as a repository of knowledge that is *public* only in the sense that it is freely accessible by other robots; access to this data will be strictly controlled, and will naturally not be available to the general public.

As pictured in Fig. 3, the cloud will feature functionality beyond that of holding data for future access and use. Since the robots are only equipped with consumer-grade computational power, the cloud will also provide processing power for heavy tasks that with lighter time constraints.

Furthermore the cloud infrastructure will allow multiple robots to access information, enabling each robot to gain information from the experience of its peers, as pictured in 4. Once prompted for services, the robot will be able to leverage the knowledge contained in the cloud, and reason upon it. For instance,

(a) (b)

Fig. 5. A pair of pictures of the robot's current external appearance.

the robot may be prompted for a recipe for pancakes on the first day of its stay with a user. At this point, the robot does not know what pancakes are, and queries the public knowledge base to gather additional information. If the recipe exists in the public knowledge base, the robot will receive it and present it to the user. The information the robot receives from the cloud can also be used as part of the reasoning process. If a robot that has been living with its user for a longer time receives the same request, it will consider the user's dietary needs and restrictions before presenting, and may even present healthier alternatives, if apposite.

This network of robots will also boost the users' social interaction by allowing them to explicitly interact with each other, via teleconferencing or by promoting face-to-face interaction.

Finally, the cloud infrastructure, by providing a means for the robots collect data on their users and their users' social circles, will enable the robots to make informed decisions, namely when it comes to classifying the user's current activity, an issue that will be further explored in Section 3.

2.4 Implementation and Testing

The robotic platform that will be used for development, pictured in Fig. 5, was designed during the Social Robot[2] project in a cooperative effort composed of both academic and industrial partners. The platform is described in more detail in [2].

This platform was designed with the intent of striking a balance between three extremely important factors: functionality, aesthetics and cost.

The robot is divided into two main parts: the top and bottom sections. The bottom section houses the drive system, power supply, processing unit and navigational sensors, including a laser range finder and sonar sensors. The drive system is composed of four wheels: two powered, differential drive wheels at the front, and two un-powered, omni-directional at the back.

The top section contains the user-interface components, such as a 10-inch touchscreen, an LED-based avatar, two RGB cameras and a depth sensor (Asus Xtion Pro Live and Microsoft LifeCam Studio). Aside from the powered wheels, the robot contains no actuators, a design decision that reflects the project's focus on social interaction.

The robot's shape and outer shell were designed with aesthetics in mind. The robot's relatively short stature (125cm) makes it a socially-recognizable presence, but still a peaceful, non-imposing one.

Lastly, cost-effectiveness was also an important factor in the robot's design. The robot's systems are mostly comprised of consumer-grade components, except where these would hinder the robot's ability to perform as intended, thus achieving a cost-effective design. Despite this, the robot itself is expected to suffer improvements and enhancements throughout the duration of the project, as it is prepared for interaction with the end user.

Software will be developed following Software Engineering good practices and, as such, unit tests will be thoroughly performed for each module that is developed, and the system will be tested in a laboratory setting. Furthermore, towards the end of the project, pilot tests will be carried out, during which the robot will be tested in a real-world scenario, interacting with real people and environments.

The pilot tests are one of the most important steps of the project, and represent a significant portion of the total workload. Two main trials will be conducted, so that the information gathered from the first trial can be used to improve the system, which will then be finally validated on the second trial.

Trials will be conducted in two locations, under the supervision of two specialized organizations. Both organizations are experienced in caring for elderly citizens, and also in participating in scientific projects, namely for the testing of innovative technologies. Users will be trained and familiarized with the system beforehand, in an attempt to maximize the value of the data that will be extracted from these trials.

[2] More information can be found at `http://ap.isr.uc.pt/?w=project_information&ID=55`.

These pilot tests will give rise to numerous ethical questions, such as how datasets may be published and stored. These issues, while not yet solved, are the focus of a part of the project, and all efforts will be made to protect the privacy and rights of the persons involved in these trials.

3 Context-Supported Action Recognition Within the GrowMeUp Solution

As stated before, much of the robot's intended functionality is supported by the ability to understand the actions the user is taking. As such, the analysis of the user's behavior will constitute one of the project's main scientific goals, as well as of the authors of this work. Essentially, we will need to be able to answer the question "What is the user doing?"

The answer to the question lies in the field of Action Recognition, which will be the focus of a significant part of our future work. It consists of recognizing the actions taken by people in sets of data gathered by sensors, a recognition that can be performed through multiple mathematical mechanisms, namely by using Machine Learning techniques or Bayesian Inference [25].

The problem of automatically recognizing actions taken by humans has been thoroughly studied, and there already exist a few comprehensive surveys on this field of research [16] [1], some focusing on slightly different sub-fields, such as the recognition of full-body motion, the recognition of hand gestures[6][22], the recognition of actions captured by multiple cameras [10] or the recognition of actions captured solely by depth data [30]. A brief look into these surveys hints us towards the already vast amount of work performed on this area. We can also observe that, as time went on, researchers have flocked to this field, resulting in a great number of recent publications. As can be observed in the mentioned surveys, some works have achieved very low error rates on many datasets, a clear indication of the field's maturity.

We are particularly interested in *multimodal* action recognition, *i.e.* the recognition of actions existent in data extracted from a heterogeneous array of sensors, such as the combination of a camera and a microphone, or of a camera and a depth sensor. The work in this area is less abundant, and a large portion of it is dedicated to the analysis of data gathered from affordable sensors such as the Microsoft Kinect [29][27]. However, some other examples can be found. For instance, in [7], the authors present a multimodal technique based on the fusion of two sensory inputs: depth images and inertial data, namely acceleration measurements taken using six accelerometers.

In their daily livelihood, the user is likely to interact with other people. Thus, we are interested in being able to identify, even if at a relatively simple level, the nature of the interaction taking place between the user and a third person. There is previous work on this matter [11][23][12], and we intend to develop this idea and integrate it into the project.

The identification of the nature of the user's interaction with other people, or with objects, will allow us to better classify the action that is being taken.

This information will be used in tandem with information gathered from other sources, such as the user and robot's current location, the user's current emotional and health status, and the localization of nearby objects of interest. This will constitute the current *context* of the action that is taking place, which will be an important part of the classification process, mainly by modulating the way the system distributes its confidence. For instance, if the action "cooking" is being detected, but the robot knows, through the data that is stored in its knowledge base, that the person is bed-ridden, then there is cause for doubting the result that we have obtained.

Additionally, recognizing and analyzing its current context will allow the robot to adapt the services it provides to the user in order to better suit them to the current situation. Parameters such as the user's current emotional and health status, the proximity to certain objects of interest, information contained in the cloud infrastructure, as well as others yet undefined, will provide the necessary contextual information for the robot to better analyze the user's *expectation*, and to better fulfill it.

The applicability of this adaptability can fit many and varied use cases. For instance, we can take into account the emotional status of the user. If a person that is calm and performing a routine action – such as cooking – asks for help, the robot may offer to present a new recipe, or to call a friend of the person to better explain the procedure needed to cook the dish at hand. However, if a person in distress asks for help, perhaps even using the same words as before, it is up to the robot to discern that the person is in an altered, perhaps endangered state, and that they may need medical help.

4 Conclusions

In this paper, we have presented the GrowMeUp research project, its main scientific challenges and goals, and its general vision and ambitions.

The project aims to develop a social service robot capable of interacting with its user, and of learning from that interaction, adapting its functionality as it "grows" and studies its user. The GrowMeUp solution will include a cloud infrastructure, that will enable the robots to share data amongst themselves, and to learn from each other's experiences. The ethical issues that plague this functionality will also receive attention from the project, and all measures will be taken to respect the privacy of the users.

Action recognition and context-aware models will be developed using Bayesian approaches, which will allow for models to be easily updated with new data. Variable spaces will relate routines, objects, emotions, recognized actions in a probabilistic graphical model that will be used to create adaptable user profiles. Services will use this data as input arguments to suit specific needs of the users.

We have also presented our views on the relevance of action recognition techniques within the confines of the project, namely on the analysis of contextual information, including the current action being taken by the user, in order to enhance and modulate the services provided by the robot. These views and the topics therein will constitute the main theme of our future work, which will mainly be directed at developing a context-aware action recognition framework that can provide useful information to support the system's intended adaptability.

References

1. Jake K. Aggarwal and Michael S. Ryoo. Human activity analysis: A review. *ACM Computing Surveys (CSUR)*, 43(3):16, 2011.
2. Paulo Alvito, Carlos Marques, Paulo Carriço, and J. Freire. A robotic platform for the social robot project. In *Proc. of the 23rd IEEE International Symposium on Robot And Human Interactive Communication (RO-MAN 2014), Workshop on Interactive Robots for aging and/or impaired people, Edinburgh*, 2014.
3. Gerard F Anderson and Peter Sotir Hussey. Population aging: a comparison among industrialized countries. *Health affairs*, 19(3):191–203, 2000.
4. Atta Badii. Companionable-integrated cognitive assistive & domotic companion robotic systems for ability & security. In *We are pleased to present the Workshop Proceedings of the 32nd Annual Con-ference on Artificial Intelligence (KI 2009), which is held on September 15-18 in Paderborn. This year the volume includes papers or abstracts of ten workshops: 3rd Workshop on Behavior Monitoring and Interpretation-Well Being, Complex*, page 164. Citeseer, 2009.
5. Filippo Cavallo, Michela Aquilano, Manuele Bonaccorsi, Raffaele Limosani, Alessandro Manzi, Maria Chiara Carrozza, and Paolo Dario. On the design, development and experimentation of the ASTRO assistive robot integrated in smart environments. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 4310–4315. IEEE, 2013.
6. Ankit Chaudhary, Jagdish Lal Raheja, Karen Das, and Sonia Raheja. Intelligent approaches to interact with machines using hand gesture recognition in natural way: a survey. *arXiv preprint arXiv:1303.2292*, 2013.
7. Chen Chen, R. Jafari, and N. Kehtarnavaz. Improving human action recognition using fusion of depth camera and inertial sensors. *Human-Machine Systems, IEEE Transactions on*, 45(1):51–61, Feb 2015.
8. Susanne Frennert and Britt Östlund. Review: seven matters of concern of social robots and older people. *International Journal of Social Robotics*, 6(2):299–310, 2014.
9. Charles E. Hill, Eileen M. Burd, Colleen S. Kraft, Emily L. Ryan, Alexander Duncan, Anne M. Winkler, John C. Cardella, James C. Ritchie, and Tristram G. Parslow. Laboratory test support for ebola patients within a high-containment facility. *Lab Medicine*, 45(3):e109–e111, 2014.
10. Alexandros Iosifidis, Anastasios Tefas, and Ioannis Pitas. Multi-view human action recognition: A survey. In *Intelligent Information Hiding and Multimedia Signal Processing, 2013 Ninth International Conference on*, pages 522–525. IEEE, 2013.
11. Tian Lan, Leonid Sigal, and Greg Mori. Social roles in hierarchical models for human activity recognition. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 1354–1361. IEEE, 2012.
12. Matthias Luber and Kai Oliver Arras. Multi-hypothesis social grouping and tracking for mobile robots. In *Robotics: Science and Systems*, 2013.
13. Joyce A. Martin, Brady E. Hamilton, Stephanie J. Ventura, M. J. Osterman, Sharon Kirmeyer, T. J. Mathews, and Elizabeth C. Wilson. Births: final data

for 2009. *National vital statistics reports: from the Centers for Disease Control and Prevention, National Center for Health Statistics, National Vital Statistics System*, 60(1):1–70, 2011.
14. Michael F McTear. Software to support research and development of spoken dialogue systems. In *Eurospeech*, 1999.
15. Amanda Mocroft, B. Ledergerber, C. Katlama, O. Kirk, P. Reiss, A. d'Arminio Monforte, B. Knysz, M. Dietrich, A. N. Phillips, Jens Dilling Lundgren, et al. Decline in the AIDS and death rates in the EuroSIDA study: an observational study. *The Lancet*, 362(9377):22–29, 2003.
16. Thomas B Moeslund, Adrian Hilton, and Volker Krüger. A survey of advances in vision-based human motion capture and analysis. *Computer vision and image understanding*, 104(2):90–126, 2006.
17. Maria Nani, Praminda Caleb-Solly, Sanja Dogramadzi, Tina Fear, and Herjan van den Heuvel. MOBISERV: an integrated intelligent home environment for the provision of health, nutrition and mobility services to the elderly. 2010.
18. Jim Oeppen, James W. Vaupel, et al. Broken limits to life expectancy. *Science*, 296(5570):1029–1031, 2002.
19. Jukka-Pekka Onnela, Benjamin N. Waber, Alex Pentland, Sebastian Schnorf, and David Lazer. Using sociometers to quantify social interaction patterns. *Scientific Reports*, 4, 2014.
20. Lucia Pigini, David Facal, Alvaro Garcia, Michael Burmester, and Renzo Andrich. *The proof of concept of a shadow robotic system for independent living at home*. Springer, 2012.
21. Didier Pittet, Stéphane Hugonnet, Stephan Harbarth, Philippe Mourouga, Valérie Sauvan, Sylvie Touveneau, Thomas V. Perneger, et al. Effectiveness of a hospital-wide programme to improve compliance with hand hygiene. *The Lancet*, 356(9238):1307–1312, 2000.
22. Siddharth S Rautaray and Anupam Agrawal. Vision based hand gesture recognition for human computer interaction: a survey. *Artificial Intelligence Review*, 43(1):1–54, 2015.
23. Kamrad Khoshhal Roudposhti. *Probabilistic-based Human Behaviour Analysis using Hierarchical Framework*. PhD thesis, University of Coimbra, Portugal, 2014.
24. Luís Santos, C. Christophorou, E. Christodoulou, J. Dias, and George Samaras. On the development strategy of an architecture for e-health service robots. In *6th International Conference on e-Health*, 2014.
25. Luís Santos, Kamrad Khoshhal, and Jorge Dias. Trajectory-based human action segmentation. *Pattern Recognition*, 48(2):568–579, 2015.
26. Joe Saunders, Nathan Burke, Kheng Lee Koay, and Kerstin Dautenhahn. A user friendly robot architecture for re-ablement and co-learning in a sensorised home. *Assistive Technology:*, 2013.
27. A. Shahroudy, Gang Wang, and Tian-Tsong Ng. Multi-modal feature fusion for action recognition in rgb-d sequences. In *Communications, Control and Signal Processing (ISCCSP), 2014 6th International Symposium on*, pages 1–4, May 2014.
28. Cameron Smith, Marc Cavazza, Daniel Charlton, Li Zhang, Markku Turunen, and Jaakko Hakulinen. Integrating planning and dialogue in a lifestyle agent. In *Intelligent Virtual Agents*, pages 146–153. Springer, 2008.
29. Jaeyong Sung, Colin Ponce, Bart Selman, and Ashutosh Saxena. Human Activity Detection from RGBD Images. *plan, activity, and intent recognition*, 64, 2011.
30. Mao Ye, Qing Zhang, Liang Wang, Jiejie Zhu, Ruigang Yang, and Juergen Gall. A survey on human motion analysis from depth data. In *Time-of-Flight and Depth Imaging. Sensors, Algorithms, and Applications*, pages 149–187. Springer, 2013.