Adaptivity as a Service (AaaS): Personalised Assistive Robotics for Ambient Assisted Living

Ronnie Smith

Edinburgh Centre for Robotics, Edinburgh, UK
ronnie.smith@ed.ac.uk

Abstract. The need for personalised Ambient Assisted Living (AAL) solutions is widely recognised. However, many existing solutions lack flexibility in terms of long-term user-adaption, a natural effect of working within the constraints of lab-based evaluation. As such, few approaches directly address a potentially desirable heterogeneous future where AAL is accessible to all: people should be able to create solutions that work for them through bricolages of assorted *off-the-shelf* (OTS) devices, including robots. For AAL to succeed at scale, these devices must share experiences of user interactions within and outwith the home. *Adaptivity as a Service* (AaaS) sandboxes adaptivity into its own research domain and posits that adaptation should be managed by a highly specialised adaptivity service that 'fits in' with different AAL solutions.

Keywords: Personalisation, Adaptive Robotics, Ambient Assisted Living (AAL), Human-in-the-Loop (HITL), Digital Twin

1 Introduction

In AAL, *adaptivity* refers to how solutions adapt to changing habits, situations, individual preferences and evolving needs of users. In practice, adaption might mean: changing how to respond to user commands and/or activities; how robots or other devices interact with humans physically and socially; how human intentions are perceived; and even modification of interaction modalities themselves.

Adaptivity can readily be viewed a machine learning problem, dependent on high quality training data. Hence, much early work on *Human Activity Recognition* (e.g. [1]), which provides vital context-awareness needed to offer real-time assistance in AAL, is based on supervised learning. Likewise, interaction adaptivity is often based on *Partially Observable Markov Decision Processes* (POMDPs). In 'GrowMeUp', social robot decision making (e.g. next action) is tuned for each user using a POMDP trained based on the positive/negative impact of actions on the user's state [2]. Historically, such data-driven approaches effectively start afresh with each new environment/user, since they do not transfer what they have learned previously to speed up the induction phase.

Models are essential in many AAL approaches. For instance, in [3] activities are recognised by comparing sensor events over sliding (fixed duration) time windows with templates of *Activities of Daily Living* in an ontological model.

These templates specify criteria such as: sensor events, duration, conditions to be met, and objects involved in an activity. Models are also used to incorporate predefined templates of potential users, such as "dependent, assisted, at risk, and active" used in the approach described in [4] to modify system behaviour. However, rigid modelling approaches are ultimately time consuming to design, and are consequently unlikely to capture the diversity of elderly care requirements.

A third category of approaches are 'hybrid' solutions, where data- and knowledge-driven (model-based) sources are combined to overcome the disadvantages of each. Work in [5] exemplifies the concept, in the context of HAR: initial knowledge-engineered activity 'seeds' (templates) enable initial activity detection, while newly discovered activities are recorded and grouped for later labelling, often with the active participation of the end user. This expands the initial pool of activities correctly recognised by the system. 'GrowMeUp' represents another example of a hybrid approach, for the way it incorporates dynamic profiling mechanisms: pre-defined profile schemas (which specify *what* can be learned) are fleshed out using knowledge gathered over time by a social robot companion [6]. While hybrid approaches may ultimately be the key to long-term adaptivity, existing approaches have not taken advantage of the wealth of useful data that can be directly extracted from a population of individuals.

There are also efforts to increase personalisation in the health care industry more widely. For example, IBM and collaboration partners are currently working on 'IBM Watson Health' to harness the potential of a wealth of existing patient health data in a cloud-based cognitive system [7]. The implication of this for patients of the future is improved treatment aided by artificial intelligence and the consolidation of their health data that can be shared with their designated healthcare professionals throughout their life.

The hypothesis advanced in this work is that all the dimensions of adaption in AAL could be delivered more effectively by adopting an *Adaptivity as a Service* (AaaS) approach.

2 A Proposal for Adaptivity as a Service in AAL

AaaS intends to address three types of long-term adaptation of AAL systems, namely: (1) context-awareness to account for predicted physical and mental decline, based on an individual's known ageing-related conditions; (2) adaption in assistive functionality, including interaction modalities, to fit individual needs/wants; and (3) quick adaption to new users, based on experience.

AaaS addresses these three types of adaption by using local and global levels to separate individual service delivery from centralised processing and storage. The intention is to slot into AAL platforms as a high-level intermediary between local control/decision making and devices such as robots, where it can intercept communication between the two. User models reside primarily in the global level, where the users are collocated in a population of many comparable users receiving adapted AAL. The collocation of this data seeks to accelerate learning potential and the amount of useful knowledge that can be extracted.

Hybrid models should be employed in AaaS for user modelling and personalisation, with their implementation intrinsically linked. Unlike in HAR, the hybrid model in the proposed AaaS framework will be realised as a *Digital Twin* (DT), i.e. a digital representation of the user and high-level *personal policies*.

AaaS relies on *Human-in-the-Loop* (HITL) learning to modify initial templates where the parameters of what can be modified are previously defined as part of knowledge engineering effort. The Digital Twin is a user model generated from given profile information. When a new user is created, a set of initial personal policies are generated and assigned to the DT that can be used to adapt system behaviour, based on prior experience with other users of similar profile.

Personal policies are nested high-level plans originally added by humans to a global policy bank. They describe: (1) assistive plans for execution in a robotic AAL environment, and (2) possible modifications to external plans to accommodate specific user wants/needs, including tunable parameters for (1) and (2). These high-level plans rely not on specific hardware, but on command execution via commonly available APIs, e.g. via *Robot Operating System* (ROS).

The DT consolidates provided and learned user information (demographic, health, care needs, preferences, etc.) into a single place. This will enable AaaS to continuously evolve its understanding of users over time and to harness that understanding for the benefit of all. The expectation is that such an approach will maximise useful knowledge retention and opportunities thereof, in order to improve both long-term acceptability and scalability of AAL solutions.

Personal policies should be updated using state of the art HITL methods, including *Reinforcement Learning* (RL). Granular policy/scenario linkage would benefit from a spatiotemporal aspect: each policy execution is a standalone instance, which can be later grouped with like instances, since the same policy may apply under different circumstances and preferences. Over time, policies for specific scenarios settle to suit user *wants*—emphasis over *needs*. This has short- and long-term forecasting implications: short-term, AaaS should be able to assess whether actions planned by the local AAL system (which may not *actively* adapt) will suit the user by testing hypothesis against the DT, allowing for plan modification or suppression; long-term, it becomes possible to predict behaviour and health patterns in relation to health conditions specified in the user profile.

Feeding back of individual adaptions should enable automatic policy selection for new users, while enabling updates of global beliefs about personal traits and the impact of functional and cognitive impairments, via aggregation of real-world experiences. At scale, there should be opportunities to investigate merging of similar feedback to best reflect what has been learned, while eliminating outliers.

3 Research Questions

Although the vision for AaaS is rather extensive, research will initially focus on the underpinning scientific issues that must be addressed in order for it to become feasible. A number of key research questions are as follows:

- How can personalisation and adaptivity plans be represented and encoded at a high level?
- How to create an interface for AaaS that enables third-party devices/platforms to most easily integrate?
- How can low-level granularity in policy/scenario linkage be achieved given the spatiotemporal aspect of human activities and daily routines?
- How can feedback from various users and sources be merged and generalised to best reflect learning from multiple similar users, while eliminating outliers?
- How does AaaS handle "unhealthy" feedback, where a policy has moulded to unhealthy individual wants?

4 Conclusion

This paper has established the fundamental principles of AaaS and highlights its novel HITL approach of a distributed architecture with local instances (individual homes) and the cloud, whereby an individual can both benefit from and contribute to a wider network of adaptivity. Future research will stem from determining the best approaches to meet key goals of AaaS.

Acknowledgements. This work is supported by the Engineering and Physical Sciences Research Council (grant EP/L016834/1).

References

1. Martins, G. S., Santos, L. & Dias, J. User-Adaptive Interaction in Social Robots: A Survey Focusing on Non-physical Interaction. en. *International Journal of Social Robotics* **11,** 185–205 (Jan. 2019).
2. Martins, G. S., Al Tair, H., Santos, L. & Dias, J. aPOMDP: POMDP-based user-adaptive decision-making for social robots. en. *Pattern Recognition Letters* **118,** 94–103. ISSN: 01678655. (2020) (Feb. 2019).
3. Chen, L., Nugent, C. D. & Wang, H. A Knowledge- driven Approach to Activity Recognition in Smart Homes. *IEEE Transactions on Knowledge and Data Engineering* **24,** 961–974 (2011).
4. Stavrotheodoros, S., Kaklanis, N. & Tzovaras, D. A personalized cloud-based platform for AAL support to cognitively impaired elderly people. *IFMBE Proceedings* **66,** 87–91. ISSN: 16800737 (2018).
5. Chen, L., Nugent, C. & Okeyo, G. An ontology-based hybrid approach to activity modeling for smart homes. *IEEE Transactions on Human-Machine Systems* **44,** 92–105 (2013).
6. Martins, G. S., Santos, L. & Dias, J. *BUM: Bayesian user model for distributed social robots* en. in *RO-MAN* UM (IEEE, Lisbon, Aug. 2017), 1279–1284. ISBN: 978-1-5386-3518-6. (2020).
7. Ahmed, M. N., Toor, A. S., O'Neil, K. & Friedland, D. Cognitive Computing and the Future of Healthcare: The Cognitive Power of IBM Watson Has the Potential to Transform Global Personalized Medicine. *IEEE Pulse* **8,** 4–9. ISSN: 2154-2317 (May 2017).