Fraunhofer
IAO

FRAUNHOFER INSTITUTE FOR INDUSTRIAL ENGINEERING IAO

KATHRIN POLLMANN (EDS.) | DANIEL ZIEGLER (EDS.)

PROCEEDINGS OF THE WORKSHOP ON BEHAVIORAL PATTERNS AND INTERACTION MODELLING FOR PERSONALIZED HUMAN-ROBOT INTERACTION 2020

PREFACE

HOW CAN WE DESIGN SOCIAL ROBOTS THAT REALLY MATCH OUR NEEDS AND EXPECTATIONS?

It is not surprising that we are probably more inclined to accept a robot in our daily lives and even our home, if we have the feeling that it is a perfect match. A robot that behaves in a way that we find pleasant and appropriate, that considers our individual characteristics and preferences. The concept of personalization has been introduced, in order to create such personal, tailored human-robot interactions (HRI). Personalizing HRI means that robot takes into account individual user characteristics and can adjust its behavior to the situation and the human interaction partner.

We believe that personalized HRI can help us to create more acceptable social robots in the future. Previous research indicates that personalization can have positive effects on the user experience during HRI as well as the user's attitudes towards and perceptions of the robot. Still, there are many open questions regarding the tools, methods and processes we can use to assess relevant user characteristics and translate them into personalized robot behaviors and interactions.

By organizing the »Workshop on Behavioral Patterns and Interaction Modelling for Personalized Human-Robot Interaction«[1] at HRI 2020, we wanted to bring together researchers from different disciplines such as psychology, interaction design, ethics, software and hardware engineering to discuss the value of and exchange ideas about the concept of personalization for HRI. We received 12 position papers covering diverse perspectives of personalized HRI. After the review process, six submissions have been accepted for presentation and discussion at the workshop.

Unfortunately, the main conference and all workshops had to be canceled due to the worldwide spread of COVID-19. While it is certainly impossible to capture the interactive nature of the planned workshop, we hope the publication of the position papers compiled in these proceedings will help to yield discussions and new ideas to support the advancement of personalized HRI.

Kathrin Pollmann and Daniel Ziegler
Workshop Organizers

1 Kathrin Pollmann and Daniel Ziegler. 2020. *PersonalHRI 2020: Workshop on Behavioral Patterns and Interaction Modelling for Personalized HRI.* In Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20). Association for Computing Machinery, New York, NY, USA, 654–655. DOI: https://doi.org/10.1145/3371382.3374846

TABLE OF CONTENTS

Social Human-Robot Interaction is Personalized Interaction

Kathrin Pollmann
kathrin.pollmann@iao.fraunhofer.de
Fraunhofer IAO, Institute for Industrial Engineering
Stuttgart, Germany

Daniel Ziegler
daniel.ziegler@iao.fraunhofer.de
Fraunhofer IAO, Institute for Industrial Engineering
Stuttgart, Germany

ABSTRACT

Previous research has indicated that robots that show social behavior when interacting with humans are better accepted. In this paper we argue that the robot behavior does not only need to be social, but also tailored to individual user characteristics, in order to be experienced as positive by the user in long-term use. To realize such personalized human-robot interactions, we need to identify differences among users, represent them in a user model and develop a number a design variants for the robot behavior. We propose to augment the conventional human-centered design process to focus on obtaining the relevant user information and creating matching design variants. We describe the different phases of the HCD4Personalization and illustrate them with examples from the NIKA project which is aimed at developing acceptable and positive interaction strategies for social robots that support older adults in their homes.

KEYWORDS

social robots, personalization, human-centered design, methodology

1 SOCIAL HUMAN-ROBOT INTERACTION

In the past decades we observed a shift from applying robots in industrial context to putting them in everyday life situations where they are expected to provide assistance and service for people. It has been shown that in such scenarios, it is important that the robot shows some social interaction skills to ensure a natural and successful interaction. Thus, the term of *social robots* has been introduced, describing robots that can interact and communicate with humans by following the behavioral norms expected by the people with whom the robot is intended to interact [1].

In the past years, much research has been conducted to improve the technical features and interactivity of robots to make them appear more social. However, their application in real-world scenarios is still limited, and insufficient answers have been provided for how to ensure the acceptance and long-term use of social robots, especially in private spaces such as our homes.

In this paper we argue that, while social robot behavior is widely regarded as a fundamental prerequisite for successful human-robot interaction (HRI), it is not sufficient to guarantee acceptance and long-term use of a robot. Existing approaches often neglect how the user experiences the interaction with the robot and that this experience is very subjective and personal, depending on individual characteristics, needs and abilities. When designing social robots for long-term private use it is hence crucial to consider the question of how we can provide a personalized positive interaction experience for the individual user.

Workshop on Behavioral Patterns and Interaction Modelling for Personalized Human-Robot Interaction 2020, March 23–26, 2020, Cambridge, United Kingdom

2 WHAT IS PERSONALIZATION AND WHY DOES IT MATTER?

Personalization has been discussed in the field of human-technology interaction as a design approach to create technical products that can be tailored or tailor themselves to individual user characteristics [3]. In the context of HRI this means that the robot takes into account individual user characteristics as well as the situational context and automatically adapts its behavior to them. Different studies have shown positive effects of personalized robot behavior on the user experience (UX), perceptions [2] and long-term acceptance of the robot [7].

To realize personal HRI, it is first necessary to create a user model that contains all user characteristics and attributes relevant to the process of personalization [6]. Later, individual user profiles can be generated for single users based on this model. The user model should best be created based on user data gather during an extensive user research.

In addition, it is not sufficient to design only one designated interactive behavior for the robot. On the contrary, it is required to have different combinable variants of the behavior, one for each different type of user profile. Based on the individual user type the matching behavior can then be selected and displayed by the robot. Although these steps are based on the general idea of the human-centered design process (HCD) [5], they cannot be fully addressed by existing user research and design methods.

3 HCD4PERSONALIZATION: A METHODOLOGICAL FRAMEWORK FOR PERSONALIZED INTERACTION DESIGN

Planning the HCD in particular includes the selection of suitable methods for the activities in the subsequent phases according to the requirements of the specific project. Especially, the selected methods have to be aligned to the project's schedule and the availability of human and technical resources, not to forget the availability of the potential target user groups. The conventional HCD and its methods are centered around identifying similarities between users and building design solutions that address those as much as possible. The concept of personalization, on the other hand, builds upon the differences between users and the ideas to provide different design solutions for different user types. Therefore, when planning the HCD to design personalized HRI there are additional questions that need to be taken into account:

- Which methods are required in order to identify and implement relevant potentials for personalization that match the users' needs and requirements?
- How to select the right samples out of the target user groups for user involvement activities in order to be able to identify relevant differences and meaningful user characteristics?

We developed a methodological framework that augments the well-know HCD with new research questions and methods to bring user differences in the focus of design and to develop personalized social robot interaction experiences. The framework is called HCD4Personalization and contains the following four steps that are integrated into the four iterative phases of the HCD: analysis, interpretation, design and evaluation.

The following sections describe the methods and research questions applied in each step, in order to focus on personalization. The description is illustrated by findings from the NIKA project. The project is aimed at developing interaction strategies for social robots that promote the independent living of older adults in their home. In the project, we applied the HCD4Personalization process to identify different user types and develop the matching behavioral variants. For demonstration purposes, in this paper, we focus in one particular use case of the project: brain training. To maintain the users' mental health and keep them active, the NIKA robot can initiate a quiz game. To make playing the game a personalized experience for the user, it is necessarily to create user model with the relevant characteristics and needs as well as design variants of the robot's behavior during the game.

3.1 Identify relevant user characteristics

The goal of the analysis phase in HCD is to understand the context of use in which the social HRI will take place. This context investigation includes specific behaviors, views and attitudes of the potential users. The methods applied in this phase usually focus on the identification of similar characteristics across the target group to be able to derive user requirements that apply to the whole user group.

However, to allow for a personalized design of social HRI, the analysis needs to put special emphasis on users' individual needs and requirements. In particular, the activities in this phase should be based on the following questions:

- How do users differ?
- Which differing characteristics are relevant for the personalization of HRI?
- How can relevant user characteristics be collected and stored in a user model?

In the NIKA project, the target group are older adults that live in their own homes and are still quite active and not cognitively or motorically impaired. To analyze the context, we conducted contextual inquiries (observations combined with interviews) in the homes of eight older adults and gathered insights about their daily routines, characteristics and psychological needs. In this phase, the adaptation of the HCD lies not so much in the applied methods, but rather in the type of information recorded and questions asked during the interview, as well as in the processing of the gathered data.

Apart from traditional user research results like personas, a key result for personalization is a candidate user model. It defines which characteristics should be known about each individual user and how these characteristics will be stored by the system (see [6]). The definition of a candidate user model allows to refer to a well-defined set of characteristics when specifying the personalization logic in further steps of the design process. The NIKA candidate user model contains attributes that represent personality traits according to the Big Five model [8] (such as extroversion, agreeableness and neuroticism) as well as a subset of user needs taken from the UXellence® framework [4] (such as competence, stimulation and competition).

3.2 Derive personalization hypotheses

To specify user requirements in the second HCD phase, the findings from the context of use analysis need to be interpreted and further processed, in order to derive ideas for personalizing the robot behavior. In the context of personalization this interpretation again needs to be focused on the variability of user characteristics and the corresponding variability of the design space:

- Which specific properties comprise the HRI design space?
- Which combinations of user characteristics necessitate HRI solutions with different properties?

In the NIKA project, we used the concept of personalization hypotheses to structure the interpretation process and document its results. Each personalization hypothesis describes the possible variation of a specific design dimension as well as its relation to specific user characteristics. To reduce initial complexity, we focused on extremes of the design space in a first step. To provide an example, the design dimension of *user motivation* was derived from the analysis in phase 1. This means that the user's performance and willingness to do the brain training can be increased by creating specific robot behavior to motivate the user during the quiz game. This might be realized in different ways. Opening up the design space for this *user motivation* dimension we came up with the following two extreme ideas for design variants: In one extreme the robot might behave supportive and extensively praise the user for her performance and effort. In the other extreme the robot might challenge the user e.g. by questioning if she knows the correct answer to the next question. Based on our prior user research findings we would expect users with a low level of neuroticism or a negative basic mood to prefer the supportive behavior, while users with a strong need for competition most likely experience the challenging behavior as more positive.

3.3 Create HRI design variants

In the design phase, the concepts outlined in the personalization hypotheses are transformed into concrete design solutions, with the following questions in mind:

- What different variants of HRI need to be designed?
- How can the concepts for the variants be represented by prototypes that can be experienced and evaluated by users?

In general, a personalized system needs to be designed in a modular way that allows to keep parts of the interaction the same for all users and change some parts depending on the individual user profile. For example, in the quiz game all users basically get the

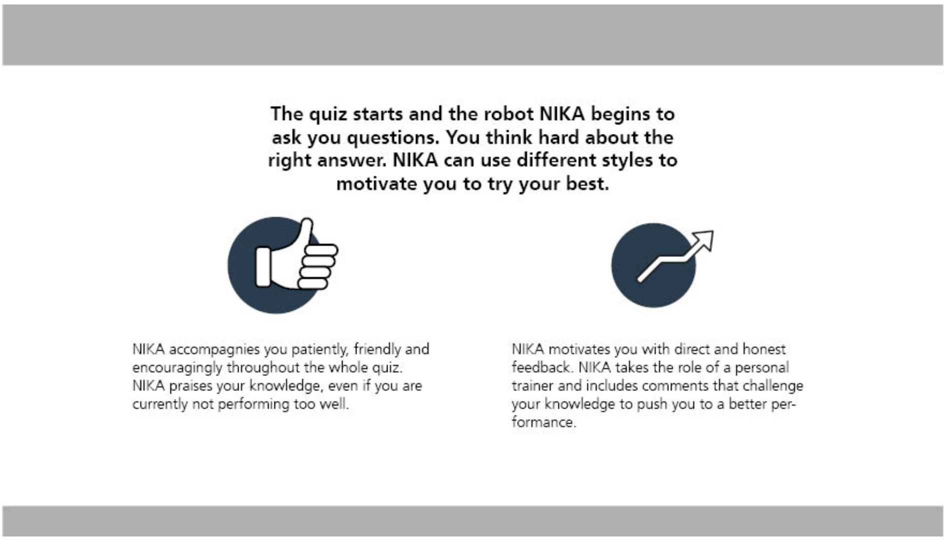

Figure 1: Prototypes for two different motivation styles the robot can show during the quiz game.

same course of the game, while aspects like the motivational behavior of the robot might vary between individual users. Thus, in the NIKA project, we created different variants for the robot behavior, building upon the design space defined through the personalization hypotheses. As the effort to implement behavior on a actual robot is often high, in this early stage of the project we worked with low fidelity prototypes of the behavioral variants. Concretely, we created a textual description for each variant which was accompanied by a rather abstract illustration. Figure 1 depicts an example of the prototypes for the different motivation styles mentioned before.

3.4 Evaluate personalized HRI

In the evaluation phase, the prototypes developed in the design phase are evaluated together with users. When evaluating different design variants for personalized HRI the following questions need to be addressed:

- Are the designed variants for the robot behavior the right ones?
- Have the design variants been assigned to the right user types or combinations of user characteristics?
- How can the design variants be optimized to better meet the variation of preferences within the target group?

In the NIKA project the evaluation was conducted as an online study. In this study 101 participants aged 60 or higher provided their feedback on the prototypes presented as text and illustration by rating how much they liked them on a 5-point Likert scale. In addition, they filled in a questionnaire to assess their stampings on the personality traits and psychological needs included in the candidate user model. The ratings of the prototypes were used to determine whether the variants were correctly chosen as a basis for personalized HRI. Variants whose ratings show a broad statistical dispersion are suitable to address different user types and are thus good candidates for personalized HRI. If statistical dispersion is low, this does, on the other hand, indicate that this behavior is evaluated equally by all users and that this variant is hence not suitable for personalization. Moreover, the questionnaire results were correlated with the ratings for the design variants to identify connections between certain combinations of user characteristics and preferences for specific design variants. Thus, we can conclude whether we assembled the right characteristics in the candidate user model and whether our hypotheses how certain combinations of characteristics are related to specific design variants are correct.

3.5 Iterative process

As the conventional HCD, the HCD4Personalization needs to be conducted in multiple iterations, in order to arrive at a final user model, definition of the design space and suitable design variants. In this paper, we only presented results and experiences from a first iteration conducted within the NIKA project. For the next iteration, the candidate user model and design variants will be refined based on the results of the evaluation. In addition, the fidelity of the prototypes need to be gradually increased to finally implement the behavioral variants on an actual robot and collect user feedback during the real-life interaction with this robot.

4 SUMMARY AND CONCLUSION

In this position paper, we argued that HRI needs to be personalized in order to be truly social. Concretely, this means that the interactive behavior needs to be tailored to the individual user's characteristics and needs. We proposed the HCD4Personalization, a methodological framework for developing personalized HRI which is based on the conventional human-centred design process, but was augmented with relevant research questions and methods to focus on personalization. We use an example from the NIKA project to illustrate a first iteration of gathering relevant characteristics for a user model, develop personalization hypotheses, translate them into design variants and test these variants and how they can be mapped to specific user profiles (combinations of user characteristics).

ACKNOWLEDGMENTS

The NIKA project was funded by the German Federal Ministry of Education and Research (BMBF 16SV7941).

REFERENCES

[1] Christoph Bartneck and Jodi Forlizzi. 2004. A design-centred framework for social human-robot interaction. In *RO-MAN 2004. 13th IEEE International Workshop on Robot and Human Interactive Communication (IEEE Catalog No. 04TH8759)*. 591–594.

[2] Nikhil Churamani, Paul Anton, Marc Brügger, Erik Fließwasser, Thomas Hummel, Julius Mayer, Waleed Mustafa, Hwei Geok Ng, Thi Linh Chi Nguyen, Quan Nguyen, Marcus Soll, Sebastian Springenberg, Sascha Griffiths, Stefan Heinrich, Nicolás Navarro-Guerrero, Erik Strahl, Johannes Twiefel, Cornelius Weber, and Stefan Wermter. 2017. The Impact of Personalisation on Human-Robot Interaction in Learning Scenarios. In *Proceedings of the 5th International Conference on Human Agent Interaction (HAI '17)*. ACM, New York, NY, USA, 171–180. https://doi.org/10.1145/3125739.3125756

[3] Haiyan Fan and Marshall Scott Poole. 2006. What Is Personalization? Perspectives on the Design and Implementation of Personalization in Information Systems. *Journal of Organizational Computing and Electronic Commerce* 16, 3-4 (2006), 179–202.

[4] Nora Fronemann and Matthias Peissner. 2014. User Experience Concept Exploration: User Needs as a Source for Innovation. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: fun, fast, foundational*. ACM, New York, NY, 727–736.

[5] ISO. 2010. Ergonomics of human-system interaction: Part 210: Human-centred design for interactive systems.

[6] Alfred Kobsa. 2001. Generic User Modeling Systems. *User Modeling and User-Adapted Interaction* 11, 1/2 (2001), 49–63. https://doi.org/10.1023/A:1011187500863

[7] Maarten H. Lamers and Fons J. Verbeek (Eds.). 2011. *Human-Robot Personal Relationships - Third International Conference, HRPR 2010, Leiden, The Netherlands, June 23-24, 2010, Revised Selected Papers*. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, Vol. 59. Springer. https://doi.org/10.1007/978-3-642-19385-9

[8] Sebastiaan Rothmann and Elize P Coetzer. 2003. The big five personality dimensions and job performance. *SA Journal of Industrial Psychology* 29, 1 (2003), 68–74.

Human-centered HRI Design – the More Individual the Better?

Chances and Risks of Individualization

Miriam Funk
Human Factors, Ergonomics
Federal Institute for Occupational Safety and Health
Dortmund, Germany
funk.miriam@baua.bund.de

Patricia Helen Rosen
Human Factors, Ergonomics
Federal Institute for Occupational Safety and Health
Dortmund, Germany
rosen.patricia@baua.bund.de

Sascha Wischniewski
Human Factors, Ergonomics
Federal Institute for Occupational Safety and Health
Dortmund, Germany
wischniewski.sascha@baua.bund.de

ABSTRACT

In the emerging field of human-robot interaction (HRI) the adaptation to the situation and individual needs is increasingly recognized in research. In the context of work, individualization of robotic systems can be beneficial for a human-centered and safe workplace design. To realize the associated opportunities the continuous (real-time) collection and processing of personal data is necessary. These might not only be used to support employees but also for monitoring tasks or workplaces. Therefore, it is important to take into account the risks to which employees might be exposed when using these kind of systems. Furthermore awareness regarding potential risks should already be considered during the design process of robotic systems. This paper presents related work showing examples of beneficially individualized human-robot interaction. Nevertheless, the trade-off between an individualized human-robot interaction and the possible violation of data protection has to be considered carefully.

CCS CONCEPTS

•Computer systems organization~Embedded and cyber-physical systems~Robotics~Robotic autonomy •Human-centered computing~Interaction design~Interaction design process and methods~User centered design •Security and privacy~Human and societal aspects of security and privacy~Privacy protections

KEYWORDS

Human-robot interaction, data protection, occupational safety and health, individualization

Workshop on Behavioral Patterns and Interaction Modelling for Personalized Human-Robot Interaction 2020, March 23–26, 2020, Cambridge, United Kingdom

1 Introduction

Modern technologies have become indispensable in our daily lives. Currently, more and more robotic systems can be found both in the private sector as well as in the world of work. They can take over different kind of tasks and are mainly used to support or assist people at home or during work. The increasing prevalence of these systems lead to intensive research efforts to continuously improve human-robot interaction (HRI). Known from the field of human-computer interaction one important factor for successful interaction is the ability of the system to adapt to the individuals' needs and preferences. Previous findings also suggest that the adjustment of robotic behavior to the interaction situation has a positive impact on user experience regarding usability, acceptance [1] and trust [2]. In the context of work, individualization can also be beneficial for a human-centered and safe workplace design regarding physical [3] as well as psychological parameters [4]. To realize the associated advantages of individualized human-robot interaction, a large amount of personal data is required. Besides the potential improvements, the possibility of employees' data being misused might rise. It is therefore essential to consider whether the realised opportunities outweigh the associated risks with particular regard to the principles of data protection.

In the following, we present findings from our previous work that can be transferred when considering individualized human-robot interaction. Besides the positive implications we will also discuss the results in connection with workplace monitoring and data protection to give some first implications for the utilization of individualized occupational robotic systems.

2 Individualized Human-Robot Interaction: Preliminary Related Work

In order to ensure safe, healthy and productive working conditions, the individual and autonomic adaptation of technological systems to people's abilities is becoming more and more common. The process of individualization can refer to various factors: from the consideration of physical characteristics and situational behavior to individual personal characteristics, from manual input by the employee to automatic behavior adjustment by the robot itself. Different aspects of individualization were addressed in previous research. They will be presented in the following and provide first indications for further consideration of individualized human-robot interaction and their benefits in the context of work.

2.1 Physical Individualization

The project "INDIVA - Individualized socio-technical workplace assistance for industrial production" (funding number: 16SV6253) focused on the individual adaptation towards employee's physical demands. Within the project an industrial welding and handling robot with 150 kg payload and a lightweight robot with 14 kg payload were used. In the first industrial scenario the welding and handling robot was used to assemble a backplane module including a wire harness into the rear of a car body. In the second scenario the exact positioning of studs had to be completed. Within the project a strong focus was put on the recording and processing of worker-specific physical parameters, as they are the basis for an individualized workplace design when a semi-automatic robot collaborates with the human. Using digital human modeling, task specific movements were simulated to plan and implement an autonomous and coordinated acting of the robot. Humans were able to manually adjust the handling robot to their physical needs, e.g. the handling position of the robot in order to assemble parts at their preferred position. In order to investigate the usability of the procedure as well as technology acceptance parameters, a video-based evaluation study was conducted. The results showed an overall positive attitude of the 19 participants regarding the prototypical demonstrator. This result was underlined by positive impressions of potential users regarding functionality and usability [3].

2.2 Context Sensitive Individualization

The project "AIM - Work assistance system for the individualization of work organization and training methods" (funding number: 02L14A162) dealt with the context sensitive and person-specific supply of relevant information in current production environments. The aim of the project was to develop a concept using smart mobile devices as a context-sensitive assistance system based on trajectory recognition. As part of the project we conducted a laboratory study to investigate the effect of context-sensitive information provision on human work. In a multitask setting based on assembly tasks in the automotive industry, 45 participants carried out the task either under the context-sensitive or non-context-sensitive condition. This means that information where given at a convenient moment or at a more inconvenient point. The results showed that the perceived stress was statistically slightly higher under non-context sensitive conditions. Nevertheless, the average workload under this condition was still reasonable and for work performance no differences could be found. In order to act autonomously, this kind of assistance needs a permanent trajectory recognition and the integration of production data in real time is necessary [4].

2.3 Behavioral Individualization

Another study within the project "FRAME - Elevator use and room entry for robots involving human assistance" (funding number: 16SV7834) addressed the question of prosocial behavior regarding robots. In complex or dynamic environments, robots may not be able to adapt to new conditions and rely on human assistance to perform their task. Within the project determinants of an effective and human-centered request for help are investigated. From the interaction between humans it can be derived that trait models of personality can be used to predict prosocial behavior. To investigate if behavior patterns towards robots can be predicted by trait models of personality a laboratory experiment was designed likewise. For the experiment an autonomous platform with a mounted lightweight manipulator was used. In a real human-robot-interaction situation the robot repeatedly asked for help when opening a door and picking up a small product palette while the participant had to carry out its own task. The helping behavior of 50 participants was measured by observations of the participant's provision of help. Additionally personality traits were assessed using standardized questionnaires. The results showed that personality also affects human-robot interaction. Contrary to interactions between humans, the study results indicate that prosocial behavior towards robots is influenced by the personality dimensions of conscientiousness and openness. Whereas literature shows, that inter-human helping behavior is related to higher values on the dimensions agreeableness and honesty-humility. Nevertheless, in order to use findings from research on personality traits to improve human-robot interaction additional studies are needed. In such scenarios, the robot needs to know about all facets of the specific personality profile of the interacting individual [5]. Accordingly, using this knowledge to provide individual robot behavior customized to individual traits or characteristics raises strong ethical issues that must not be ignored and should be subject to future research activities.

2.4 Benefits of Individualized Human-Robot Interaction

The presented results share, that on one hand they give first implications on how individualized human-robot interaction can make a contribution in an occupational environment. To ensure a useful and user orientated system, the robot should enhance the employees' daily activities and not be associated with additional effort [6]. The perceived usefulness of the system is not only influencing technology acceptance, but also the intention to use or even the actual usage [7]. This is essential, because if the robotic system is not used in the intended way or even not at all, neither the beneficial assistive nor the associated efficiency and productivity opportunities can be realized. Thus, usability and acceptance are two important factors that shall be taken into account in respect of a humane workplace design. In this regard, the individual physical adaptation of work systems or assistant robots, first and foremost, address the common problems of awkward posture exposure in the short-term as well as long-term effects regarding musculoskeletal disorders. As shown before, an autonomous adjustment to the individuals' specific physical requirements can also have a positive impact regarding usability as well as technology acceptance. In addition, individualization can also effect mental workload, which is a decisive impact factor for work performance and an important parameter when it comes to occupational health and well-being: Mental underload may cause feelings of frustration or annoyance while mental overload can lead to confusion, decrease performance and increase the risk of failure [8]. When combining the reported results reveal that the adaptation to situational requirements may contribute to the optimization of mental workload. This can in turn positively effect the individual's mental state as well as performance parameters. To generally improve human-robot interaction processes and quality individual personality characteristics should be acknowledged. Personality traits not only affect the interaction between humans but also between robots and humans.

On the other hand, every individualization approach also involves the processing of a large number of different personal data. While the one-time collection of physical data in combination with access possibilities by the employee remains tangible, the continuous processing of highly sensitive personal data like individual assembly movement patterns or personality traits contains possibilities for exploitation.

3 Considering the Risk of Monitoring and the Principles of Data Protection

From the perspective of occupational safety and health, as shown before, individualization can contribute to optimize physical and psychological stress and strain. In the context of work, however, it is necessary to balance the benefits against the risks involved especially in turns of personal data. In the case of the INDIVA project described above, physiological personal data was used to avoid or counteract the possibility of physical illness. To do so, the necessary data is not collected continuously but only once. Besides, the user has the opportunity to adjust the system on their own. Even if the collected data must be adjusted in regular intervals to ensure optimal support there is no continuous data recording. From an employees' perspective the renouncement of real-time data gathering can be beneficial because the collected data cannot only be used to ensure the functionality of an adaptive system or to increase efficiency. It might also be misused to analyze the employees' performance and behavior to facilitate comprehensive workplace monitoring [9]. Especially when personal data is collected continuously and merged with data of the socio-technical work environment new and perhaps unwanted information can be generated. Under certain circumstances, it may open up or expand the possibilities for monitoring the individual employee in terms of habits or performance. Even if the accruing information is not collected for this purpose, work assistance systems can be perceived as monitoring tool. This perception can not only have a negative impact on motivation, satisfaction and organizational trust but can also be an additional stressor in the work situation [10]. Therefore, when considering individualization of robotic systems in the context of work, the principles of data protection following the General Data Protection Regulation (GDPR) have to be taken into account [11]. On one hand they entail the potential to reduce insecurity and fear of being monitored by the system. On the other hand they make a direct contribution to avoid arbitrary data use by the employer. In this regard, one important aspect is the principle of transparency [11]. It states that before collecting personal data, employees must be informed about the nature, purpose and duration of the storage. Furthermore, employees should have the opportunity to obtain information about their personal data being collected at any time. It should also be clearly defined, in advance, for what purpose the data will be collected, processed and stored. This purpose should directly be related to the employment relationship and not only be aimed at improving the organizations' effectiveness. If this purpose is fulfilled, there is no further need to collect and process data [11]. Besides the usage, access and exploitation rights regarding the collected data should also be clearly regulated. The limitation to a selected group of persons can help to reduce the probability of data being used in an inexpedient manner. Another way to restrict access of unauthorized persons and to avoid the risks of (perceived) monitoring is to store and process the necessary data only locally. However, from a technical perspective this is only possible if the system was designed that way. Therefore, along with choosing the system and developing the specific application, the GDPR principle of data minimization should be considered thoughtfully: "Personal data shall be limited to what is necessary in relation to the purpose for which they are processed" [11]. Yet preferable and necessary data might not always be in concordance. In this case, the principle of data minimization should be prioritized. Moreover, privacy should be built in the system by default, meaning that the setting should be chosen in such a way that ideally no personal data are captured. If a person does not change the setting proactively privacy protection should be ensured. These aspects become even more important if not only physical and behavioral aspects are taken into account. Depending on the scope and influence of data, it can be used enabling robots to influence people. On the basis of psychological knowledge in combination with information of individual personality aspects robotic systems can have an impact on human behaviour, attitude or even cognitive and decision making processes without coercion [12]. Besides positive effects on human-robot interaction this kind of technology can also be used to negatively affect the situation or even manipulate the interacting person [13]. In the context of work, this can entail considerable risks for the employees. An autonomous acting without intervention possibilities can exclusively focus on the interests of the employer at the expense of the employees. In this respect it is indispensable to carefully shape the introduction process of individualized robotic systems respecting the possible unintended consequences. These issues will be addressed within the EU-funded research project SOPHIA.

The project "SOPHIA – Socio-physical interaction skills for cooperative human-robot systems in agile production" aims at introducing human-robot teams to structured as well as unstructured production environments. Within the project a strong focus will be put on sensitive wearables and robotic systems that enable worker state monitoring as well as real-time feedback functions. For this purpose mobile robot platforms with mounted manipulators and interfaces as well as fixed lightweight robots will be used. In order to insure these functionalities the gathering of process data and personal data is vital. Within the development of each use-case the trade-off between data collection for an optimized interaction between employees and robots respectively wearables and collecting as little data as necessary is addressed actively. Especially in unstructured environments the ability of a robotic system to adapt increases its value. Nevertheless, neither flexibility nor real-time monitoring of worker states must be met on behalf of violating data protection issues.

4 Practical Implications and Conclusion

The benefits and risks associated with individualized robotic systems have to be considered carefully. Allowing robotics systems to get more and more individualized without any restrictions can be critical. More individualization possibilities do not necessarily lead to better work assistance. Individualization must not be harmful to health and safety as well as data privacy. Only if that fact is given, can individualization be a good means for human-centred workplace design. Therefore, not only the value of personal data and information but also the operational need to manage them responsibly have grown rapidly. When introducing new forms of human-robot-interaction into workplaces, especially ones that provide the option of individualization, a works council participation or agreement is highly recommended. Moreover, later steps regarding the system implementation require a definite involvement of the works council. For both, the organization and the employees it is beneficial to define clear guidelines. They should take into account the possibilities for efficient production and assistance improving but also address the importance of acceptance, human perception and minimizing the risk of data misuse. In order to increase acceptance and avoid adverse monitoring effects the employees should be involved and informed from the very beginning. This may also include the selection and evaluation of the used robotic system as well as the workstation. One possibility is to entrust a specific division to test the system with regards to actual requirements, before a comprehensive introduction. Another option to reduce risks and increase acceptance is to restrain an autonomous adaptation or only allowing it to a limited extent. It should be up to the employee to decide or determine to what extent the robot will adapt and thus also use personal data. This can not only have a positive effect on the perception and acceptance of the robot, but also counteract a potential feeling of being monitored and prevent permanent manipulation possibilities.

Overall, it can be concluded that employee involvement is not only mandatory but also a key element to achieve the intended improvements of individualized human-robot interaction whilst ensuring a safe and healthy workplace design.

ACKNOWLEDGMENTS

The research and development project “FRAME - Elevator use and room entry for robots involving human assistance - Human-centered interaction concepts in occupational settings (MIkA)” is funded by the German Federal Ministry of Education and Research (BMBF) and managed by the Project Management Agency VDI/VDE Innovation + Technik GmbH (funding number: 16SV7834).

SOPHIA project has received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement No 871237’. ‘The opinions expressed in this document reflect only the author’s view and reflects in no way the European Commission’s opinions. The European Commission is not responsible for any use that may be made of the information it contains.

The authors are responsible for the contents of this publication.

REFERENCES

[1] Elizabeth Broadbent, Rebecca Stafford and Bruce MacDonald. 2009. Acceptance of healthcare robots for the older population: Review and future directions. In *International Journal of Social Robotics* 1, 4 (2009), 319-330.
[2] Peter A. Hancock, Deborah R. Billings and Kristin E. Schaefer. 2011. A Meta-Analysis of Factors Affecting Trust in Human-Robot Interaction. In *Human factors* 53, 5 (2011), 517-527.
[3] Thomas Carsten and Bernd Kuhlenkötter. 2016. INDIVA - Individualisierbare sozio-technische Arbeitsassistenz für die Produktion: Abschlussbericht des LPS zum Verbundprojekt INDIVA. Report number: 16SV7354, 01150339. Available: www.tib.eu [Accessed Feb. 06th, 2020].
[4] Miriam Funk, Matthias Hartwig and Sascha Wischniewski. 2019. Evaluation of Assistance Systems for Manual Assembly Work. In *IEEE/SICE International Symposium on System Integration.* SII 2019, Paris, France, 794-798.
[5] Martin Westhoven and Patricia Tegtmeier. Submitted. Influence of Personality Traits on Helping Behavior in Human-Robot-Interaction. In *Transactions on Human-Robot Interaction*.
[6] Maartje M.A. de Graaf, and Somaya Ben Allouch. 2013. Exploring influencing variables for the acceptance of social robots. In *Robotics and Autonomous Systems* 61, 12 (2017), 1476-1486.
[7] Marcel Heerink, Ben Krose, Vanessa Evers and Bob Wielinga. 2010. Assessing the acceptance of assistive social agent technology by older users: The Almere model, In *International Journal of Social Robotics* 2, 4 (2010), 361-375.
[8] Luca Longo, 2017. Subjective usability, mental workload assessments and their impact on objective human performance. In *IFIP TC13 Conference on Human-Computer Interaction. INTERACT 2017, Mumbai, India,* 202-223.
[9] G. Stoney Alder and Maureen L. Ambrose. 2005. An examination of the effect of computerized performance monitoring feedback on monitoring fairness, performance, and satisfaction. In *Organizational Behavior and Human Decision Processes* 97, 2 (2005), 161-177.
[10] Nils Backhaus 2019. Context Sensitive Technologies and Electronic Employee Monitoring: a Meta-Analytic Review. In *IEEE/SICE International Symposium on System Integration.* SII 2019, Paris, France, 548-553.
[11] General Data Protection Regulation (GDPR) 2016. Retrieved February 6, 2020 from https://gdpr-info.eu/
[12] Nils Backhaus, Patricia H. Rosen, Andrea Scheidig, Horst-Michael Gross, Sascha Wischniewski. 2018. “Somebody help me, please?!” Interaction design framework for needy mobile service robots. In *IEEE Workshop on Advanced Robotics and its Social Impacts.* ARSO 2018, Genova, Italy, 8-15.
[13] Jason Borenstein and Ronald C. Arkin. 2017. Nudging for good: robots and the ethical appropriateness of nurturing empathy and charitable behavior. In *Ai & Society* 32, 4 (2017), 499-507.

Generating empathic responses from a social robot: An integrative multimodal communication framework using Sima Robot

Carmina Rodríguez-Hidalgo
School of Communications and Journalism
Universidad Adolfo Ibáñez
carmina.rodriguez@uai.cl

Felipe Araya
Sima robot
Fablab University of Chile
Santiago, Chile
felipearaya@simarobot.com

Virginia Dias
Sima robot
Fablab University of Chile
Santiago, Chile
virginia@simarobot.com

Alejandro Pantoja
Sima robot
Fablab University of Chile
Santiago, Chile
apantojas@gmail.com

Hugo Araya
Sima robot
Fablab University of Chile
Santiago, Chile
hugo@simarobot.com

ABSTRACT

Despite their rapid development, social robots still face difficulties for achieving smooth communication with humans, particularly when we consider interactions within an emotional context. To contribute in this search for meaningful human-robot interaction (HRI), we present an interaction framework focusing on children, that introduces a multimodal interaction model to parametrize the social robot's empathic responses according to type of emotional valence and intensity. This work in progress presents an interaction sequence of five steps between a child and a user-programmable social robot (SIMA), for five basic emotions and four mood/affective states. The goal of the model is that of parametrizing the robot's empathic emotional responses considering multimodal communication aspects in a simultaneous way, such as: communication chronemics (e.g. time of the day, speech reaction time) and kinesics (e.g. amplitude of micro arm movements to mirror emotions). Ultimately, through completion of the five stages, the model aims at boosting emotional self-awareness and learning in children.

CCS CONCEPTS

• Information systems • Human factors • User interfaces • Interaction styles

KEYWORDS

Social robots, empathy, multimodality, interaction, non-verbal language, emotion.

Workshop on Behavioral Patterns and Interaction Modelling for Personalized Human-Robot Interaction 2020, March 23–26, 2020, Cambridge, United Kingdom.

1 Introduction

Social robots have emerged as new communication partners [1] and as such, face the challenge of achieving new meaningful ways to communicate with humans. Although social robots pose new communicative affordances which allow a social interaction [2], there is still room for improvement in aspects such as synchronous communication between a robot and a human, and vice versa. For instance, social robots can still improve their simultaneous responses in a more empathic and social manner. Empathic communication involves vicariously responding to the perceived emotional state of others [3]. This work in progress posits that tailoring the verbal and non-verbal emotional responses of the social robot to users' self-reported emotional states, can result in increased perceptions of perceived empathy. Consequently, and through this enhanced empathy, a number of positive communicational and psychological outcomes can be reached, such as increased rapport, liking and persuasiveness. Also of importance, this empathic communication model is ultimately intended to reach prosocial objectives in the human user, such as enhanced emotional awareness.

2 Theoretical framework

The literature and theories informing the framework pertain to human machine communication and social psychology. First, the media equation theory [4] states that people tend to treat machines as they would treat people. This theory gives ground to assume that a meaningful interaction between a humanoid social robot and a child can be emotionally effective, since humans, and particularly children, tend to see the social robot as an interaction partner and may quickly adapt their behaviors to the social robots' actions [5].

Secondly, emotional mimicry in context theory [6] states that emotional mimicry occurs in highly contextual communicational situations and that this mimicry can act as a social regulator. This means that emotional mimicry in a one on one communicative situation can help people better deal with their emotions. Therefore, by having the social robot mimic the child's facial and bodily expression according to the emotion, we have grounds to assume that it can help children better recognize his/her emotions and thus offer a leeway to deal with them and become more emotional self-aware. Moreover, emotional mimicry according to the users' state may provide a sense of empathy. Although variously defined, for the purposes of this study we understand empathy as an emotional response to the perceived emotions of others [7]. In this case, the social robot would emotionally react to the perceived emotion of the child, which functions as a triggering input for the robot to display an emotionally tailored response.

Third, emotional self-awareness theory posits that higher attention to ourselves leads to judging our own behavior, fostering recognition of felt emotions [8]. Emotions are caused by a set of complex and synchronized component responses which often results from adjustments or disappointments in achieving one's goal [9]. Learning to identify emotions is an essential skillset to function well in life [9]. As children of young age are still learning to recognize and accept their emotions [10], we posit that social robots can help children reach higher emotional self-awareness to help the child better manage his or her negative emotions and also learn to savor the positive emotions even more.

This work in progress presents an integrated framework that aims to parametrize the social robot's verbal and non-verbal communication, according to type of emotion, with the goal of establishing a smooth interactive communication sequence with a child. Preliminary, our focus is on 8 to 11-year-old children, based on their developmental characteristics [11]. In sum, this research seeks to advance the field of Human Machine Communication since the proposed framework can aid in developing desirable traits in social robots, such as empathic responses. [12] has outlined a list of 'desiderata' or basic social functions that a social robot should have, such as: (a) performing multiple speech acts; (b) interact affectively; (c) respond to more than one command; (d) purposefully speak, among others. We believe that the proposed framework represents a step forward in this direction. Moreover, the framework could also help illuminate and improve the social and emotional responses of other social robots, particularly educational ones.

3 Explaining the framework

The framework is applicable in an anthropomorphic social robot, Sima robot (www.simarobot.com). The robot functions through a cell phone embedded in a 3D anthropomorphic body. It is able to perform a variety of emotional facial expressions and body movements. The robot's original purpose is to teach preschoolers, therefore comes with a number of educational games from which children can choose in a visual menu shown in the screen by tapping in different icons, e.g., vocabulary, simple additive and subtraction operations. The robot has a kawaii or 'cute' design (figure 1), specially tailored for communication with children to avoid the uncanny valley effect [13]. The robot comes with an array of facial emotional expressions and can react emotionally to whether the child has replied correctly or incorrectly to a question. For a small example of the facial emotional expressions of the social robot, see figure 2.

Figure 1. Sima robot (image source: simarobot.com)

Figure 2. Small example of SIMA's facial expressions.

The framework has two main components, the verbal and nonverbal. These have been programmed in a mobile phone application (SimaRobot). The interactions are powered by Watson IBM, a question-answering computer system, based on DeepQA software, which is capable of answering questions in natural language [14]. To establish a dialog, the robot uses the smartphone's Automatic Voice Recognition (AVR) feature through the app, similar to a system such as Google Assistant. When the child speaks, Watson activates its dialog features through a language algorithm - based on different keywords and sentence fragments - and is able to respond with the most 'sensible' solution. Therefore, the robot can recognize words via the app and when "hearing" certain keywords and intents in the child's speech to one of Sima's questions (e.g., how are you feeling?), it can trigger a spoken verbal response from an array (or combinations of) replies.

The nonverbal component consists of a series of bodily and facial expressions to respond contingently by using meaningful expressions and gestures in response to the child's self-reported emotion. The robot's body movements include for instance raising one or both arms, walk, or stand on one leg. In particular, the present framework utilizes the robots' arm movements as nonverbal means to represent the valence and intensity of the recognized emotion [15], similar as to how humans would physically react and express emotions. For instance, the emotion of surprise when e.g., hearing good news, implies that the robot energetically waves its arms towards the upper body (similar to a 'wave' motion). In turn, when triggered by a sad word, the robot would lower its arms in a slower motion in response, simultaneously to the verbal response. For an example of the robot's reactions to a trigger word or intent, see Figure 3.

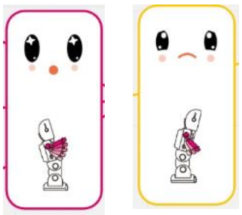

Figure 3. Examples of SIMA's emotional facial and bodily movements in reaction to the keywords (or related phrases), including the words "surprise" (left) and "sad" (right).

To what respects the robots' facial expressions, this nonverbal aspect uses the robot's "face," via the Simarobot app using the telephone screen. Although Sima already uses various facial expressions in the context of its educational games, the framework includes expressions specially developed for empathic communication, for instance those displayed on Figure 3. Because the telephone screen is able to show video images, Sima's "eyes" can rapidly change form (e.g., from happy to sad eyes) and wink. The robot's facial expressions (see figure 2), have been developed based on the universal expressions of the five basic affective states identified by several studies, particularly [16]. So far, the framework has been developed to support interactive sequences tailored to five basic emotions: happiness, surprise, fear, anger, sadness and four affective/dispositional states: excitation, relaxed, sleepy and bored. Both responses, verbal and nonverbal, occur simultaneously in response to a verbal input, as programmed via Watson and the phone app. Next, we detail the main interaction stages of the framework.

4 Five interaction stages

The framework considers five levels of an interaction sequence between the child and the robot. To make the talk run more smoothly, the idea is to make the child acquainted with the robot at class level, first, by introducing itself and starting a casual interaction (e.g., "Hi my name is Sima and I'm here to help you learn through play. What is your name?"). To increase personalization, two main features have been implemented. First, the robot shall talk to and respond to the child using the child's first name. Several studies account for positive interaction effects when using the person's name, such as increased ratings of a robot's friendliness [17]. Second, each of the five levels contains three main "conversation tracks" throughout, which are tailored to different times of the day (morning, afternoon, evening, e.g., "good morning! How are you feeling today? Or "it is late, how about we play a little before you go to bed?"). Each of these responses is given according to the real time, which is easy to implement as the smartphone is aware of the real time. Given that interpersonal communication can be influenced by contextual factors [18], we posit that inclusion of chronemic aspects such as the time of the day, can increase the degree of realism in the interaction.

We describe hereunder the five levels of an interaction sequence between the child and the robot. The five levels are as follows:

2.1. Communication initiation. After greeting the child according to time of the day (good morning/afternoon/evening) and having a brief "filler" conversation, SIMA robot initiates a conversation by launching probing questions to ask how the child is feeling.

2.2. Social interaction. When the child responds, the social robot speech recognition will be activated with the occurrence of a particular emotion word(s). This word will trigger the social robot to respond simultaneously both verbally and non-verbally, depending on the emotion, initiating a social interaction. In this phase, emotional 'mirroring' occurs because the social robot displays both verbal and non-verbal signs of the child's self-reported emotion.

2.3) Emotional self-awareness. After the child responds with a certain emotion, the child is prompted to reflect on their emotional state as the social robot would try to instigate emotional self-awaremess in the child. An example question would be: why do you think that you feel this way?

2.4) Transformation. In this stage the goal is to transform the emotion, either into a more positive emotional state (in the case of negative emotions), or to savour positive emotions even more (in case of positive emotions). What is key in this stage is that the social robot will attempt this transformation through what we term 'positive action', that is, proposing to play one Sima's educational games.

2.5) Action. In this stage, provided that the child agrees, the social robot and the child play a game or educational activity. In future versions, we study creating an educational or gaming activities tailored to the particular emotion.

In this way, the framework aims to: (1) identify the child's emotional state; (2) make the social robot respond in a socially and emotionally congruent (empathic) way, depending on the child's expressed emotion on the previous stage; (3) parametrize both the verbal and non-verbal emotional responses of the social robot towards the child; (4) initiate an action with the child (e.g. play an educational game) which would modify the child negative emotional state or increase his or her positive emotions. Ultimately, a broader goal is to achieve more meaningful communication between a social robot and a human, in this case, the child, in a situation which creates the best possible mood for learning, contributing to an effective learning environment [19].

5 Preliminary discussion

The present framework intends to make empathic responses feasible between a child and a social robot. Its goal is to parametrize the robot's responses by considering the child's emotional state.

The framework is integrative, because it comprises both verbal and nonverbal expressions of emotions concerning the robots' kinesics and non-verbal communication elements, such as facial expressions and body movements, which attempt to be congruent with the emotion expressed by the child, and thus has the potential of creating an illusion of emotional mirroring, a response which is congruent with the child's self-reported emotion, in this way enabling meaningful HRI.

Though the framework has been initially tested, our future research will determine whether (1) children can actually recognize the robots' mood; (2) to what extent the robot emotional expression affects the children's feelings. For (1), there are empirical grounds to assume that children would be able to recognize the emotions in the social robot, as an experimental study found that participants could distinguish between negative and positive robot mood from the robot physical movements and speech [18]. For (2), there are both theoretical and empirical evidence to presuppose that Sima's reactions will affect a child feelings and behaviors, based on mood transfer findings [19].

Although in initial pre-tests the phone Automatic Voice Recognition (AVR) behaves well for the first stage of communication initiation, we observe that the interaction can flow rather smoothly, particularly in chronemic aspects such as the robot's reaction time. Because it is powered through an app, the robot responds in a timely manner compared to other interactions in which robot takes a rather long time processing the verbal input. However, in moments that the robot does not recognize the children's speech, there is a communication breakdown. A possible solution is to continue working on the robot's array of recognized responses, a challenging task in itself. Further, a shortcoming in the current framework is that the emotion recognition occurs only as verbal input from the child, the robot is not yet able to identify the individuals' emotions considering aspects such as facial expressions or tone of voice. It must be noted that adding this possibility comes with an added set of complications (e.g. the child's privacy when activating facial recognition in the phone device).

Further, it is equally challenging to provide an effective response in every step because the robot should be able to capture and decode a vast array of phrases. Whereas an adult may be perfectly capable of stating "I feel sad," a child may be more spontaneous in his or her emotional expressions. Of relevance would be to improve the AVR to interpret non-verbal utterances such as "grrr" or howls, or other random noises as representative of emotional states.

Even though Sima is powered through a telephone app, and that the functionality of its voice recognition feature may be similar to Google Assistant, we believe that one main advantage of the present framework is that it provides situated, embodied interaction, allowing anthropomorphic nonverbal and verbal language, which already provides situatedness and embodiment to the interaction, features which have been shown to lead to improved learning outcomes [20]. According to [20] "interactions between a child and robot should be contingent and multimodal, and provide appropriate forms of feedback" to foster a learning environment. They further suggest that the robot could keep track of children's advancement "and perhaps their emotional states during the tutoring sessions and adapt to these". Although the framework considers interactions in only five levels, we believe it represents a first step towards making an interaction with a social robot more personable and empathic, in a manner that can contribute towards better learning outcomes and higher emotional awareness.

6 Conflict of interest statement

The first author declares that this study has been conducted in the absence of any commercial or financial relationship with Sima robot that could be construed as a potential conflict of interest.

REFERENCES

[1] Shanyang Zhao. Humanoid social robots as a medium of communication. New Media & Society (2006). 8(3). 401-419. doi: 10.1177/1461444806061951.
[2] Carmina Rodriguez-Hidalgo. Me and my robot smiled at one another: The process of socially enacted communicative affordance in human-machine communication. Human-Machine Communication (2020), 1(1), 4.
[3] Albert Mehrabian and Norman Epstein. A measure of emotional empathy. Journal of Personality. 40, 4. 525-543.
[4] Byron Reeves. Clifford Nass. 1996. The media equation: how people treat computers, television and new media like real people and places. Cambridge University press: UK.
[5] Marco Nalin, Ilaria Baroni, Ivana Kruijff-Korbayova, Lola Cañamero, Mathew Lewis, AryelBeck, Heriberto Cuayahuitl, Alberto Sanna. Children's adaptation in a multi-session interaction with a humanoid robot. In *21st International Symposium on robot and human interactive communication. Pages 351-357 IEEE , 2012.*
[6] Ursula Hess and Agneta Fischer. Emotional mimicry as social regulation. Personality and social psychology review. 17 (2013). 142-157. Doi: 10.1177/1088868312472607
[7] Geo Albert Mehrabian and Norman Epstein. A measure of emotional empathy. Journal of Personality (1972), 40, 4. 525 – 543.
[8] George Herbert Mead. 1934. Mind, self and society. University of Chicago Press: Chicago, IL.
[9] Klaus Scherer. What are emotions? And how can they be measured? Social science information. 44 (2005). 695-729. https://doi.org/10.1177/0539018405058216.
[10] Daniel Goleman. Emotional intelligence: Why it can matter more than IQ. Bloomsbury publishing. (1996).
[11] Herbert Ginsburg. Piaget's theory of intellectual development. Prentice-Hall, Inc: Washington.
[12] Nikolaos Mavridis. A review of verba land non-verbal human-robot interactive communication. Robotics and autonomous systems. (2015) 22-35. Doi: 10.1016/j.robot.2014.09.031.
[13] Masahiro Mori. The uncanny valley. Energy (1970). 33-35.
[14] IBM Corporation, DeepQA Project Frequently Asked Questions. Retrieved from: https://researcher. Watson.ibm.com/researcher/
[15] Junchao Xu, Joost Broekens, Koen Hindrinks and Mark Neerincx. Mood contagion of robot body language in human robot interaction. Autonomous agents and multi-agent systems, 29. (2015). 1216-1248.
[16] Scherer, K. R. Emotion. 2000. In M. Hewstone, & W. Stroebe (Eds.), Introduction to Social Psychology: A European Perspective (3rd ed.) (pp. 151-191). Oxford: Blackwell.
[17] Deborah Johanson et al. Smiling and use of first-name by a healthcare receptionist robot: Effects on user perceptions, attitudes and behaviours. Journal of behavioral robotics 11, 1. 40-51. doi: 10.1515/pjbr-2020-0008.
[18] Brant Burleson. 2010. The nature of interpersonal communication. The Handbook of communication science. Sage: London.
[19] Lev Vygotsky. 1978. Socio-cultural theory. Mind in society.
[20] Nikolaos Mavridis. A review of verbal and non-verbal human-robot interactive communication. Robotics and autonomous systems. (2015) 22-35. Doi: 10.1016/j.robot.2014.09.031.
[21] Roland Neumann and Fritz Strack. Mood contagion: the automatic transfer of mood between person7s. Journal of personality and social psychology.79. 211-223. (2000). doi: 10.1037//0022-3514.79.2.211.
[22] Paul Vogt et al. Child-robot interactions for second language tutoring to preschool children. Frontiers in human neuroscience, 2017, 11. doi: 10.3389/fnhum.2017.00073.

Personalized Close-Proximity pHRI

Katsu Yamane
kyamane@honda-ri.com
Honda Research Institute USA
San Jose, California

ABSTRACT

This position paper presents our vision on future robots that can perform close-proximity physical interactions with human partners in order to provide both physical and emotional support. After introducing the concept of *empathetic physical support*, we summarize the current challenges in realizing such robots in general, as well as personalizing their interaction models. As a case study, we briefly describe our ongoing work on modeling hug interactions by a learning-from-demonstration (LfD) approach using demonstrations obtained through teleoperation.

CCS CONCEPTS

• **Human-centered computing → User interface programming**; *Haptic devices.*

KEYWORDS

pHRI, interaction modeling, learning from demonstration

1 INTRODUCTION

Physical human-robot interaction (pHRI) is studied mostly in two domains: collaborative robots or *cobots* [8] and companion robots [11]. Research in the former domain usually focuses on task planning and intention recognition, and is evaluated based on the efficiency of task execution. Physical interactions between the robot and human take place through an object. The latter domain often involves close-proximity interactions such as hugging but the robot is typically composed of short limbs and low-power actuators to mitigate the safety issue. Evaluation of such interactions is more difficult because there is no quantifiable objective.

As seen in Fig. 1, these domains cover only a part of possible pHRI forms in the space of physical capability and proximity to human: on one hand, robots with near-human physical capability can provide physical support but they need some space between the human co-worker to ensure safety; on the other hand, companion robots provide emotional support by, e.g., letting humans hug them, but the robot's body remains mostly passive due to lack of physical capability. In order to take full advantage of physical embodiment, robots should be able to perform close-proximity pHRI with physically capable hardware. In fact, humans are able to perform both heavy-duty physical tasks and delicate, intimate physical interactions using the same body. We believe that the commercial difficulty many social robots are facing can be attributed to the fact that they cannot perform physically meaningful tasks.

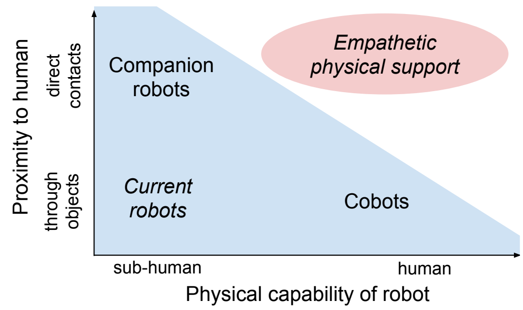

Figure 1: Relationship between the robot's physical capability and proximity to humans in pHRI research

We envision future robots that can provide *empathetic physical support*, in which the robot physically supports humans through direct contacts while constantly monitoring and adapting to the users' comfort, preference, and even emotion. Examples of this type of tasks include assisting a person standing up and carrying a person. Such interactions can be placed in the upper-right corner of Fig. 1.

Close-proximity interaction with physically capable robots will give rise to new challenges in terms of not only safety but also personalization because comfort, preference and emotion are highly subjective compared to, for example, efficiency of task execution. Large-scale user study has been the only way to evaluate the effect of this type of interactions, where statistical results may be used for evaluating the perception of average users but not of individuals.

Hence our position statement is: **Quantitative evaluation through human emotional state estimation is critical for personalizing close-proximity pHRI**.

Workshop on Behavioral Patterns and Interaction Modelling for Personalized Human-Robot Interaction 2020, March 23–26, 2020, Cambridge, United Kingdom

In the remainder of this paper, we first discuss a few challenges toward personalized close-proximity pHRI. We then describe the current status and future plans of our research in this direction using hug interaction [4, 9] as a case study.

2 PERSONALIZED CLOSE-PROXIMITY PHRI

Hardware

At least two major factors should be considered in hardware design: appearance and physical capability. Appearance has a significant impact on the human partner's expectation of the robot's physical and intellectual capabilities, or in other words, the robot's "personality." The higher the expectations are, the larger the disappointment is when they are not met. It is therefore recommended to choose the robot's appearance that correctly reflects its capability. Robots also have to be equipped with enough physical capability to support humans both emotionally and physically. The joints should be able to not only generate enough torques to actively apply forces to the human partner but also control the contact forces in order to convey subtle information such as intention and emotion through contacts.

Other aspects specific to close-proximity pHRI are 1) feel of touch, which is also critical for both safety and comfort, and 2) whole-body tactile sensing. The robot surface should be at least soft and perhaps even warm [2], and ideally include embedded tactile sensors. Unlike fingertips, however, the sensor placement does not have to be particularly dense.

Recently, soft robotics is gaining attention thanks to its potential to achieve safe pHRI through mechanical compliance, although most of the current technologies are not mature enough to be used in practical robot hardware. An alternative is to retrofit traditional robot arms with soft skin and tactile sensors. Alspach et al. [1] proposed to use 3D-printed air cavity with pressure sensors as soft skin as well as force sensors, while Block et al. [3] presented a similar idea with additional sound sensors to distinguish different contact types. Other researchers aim at developing high-density tactile sensors that can cover the whole body [10].

Interaction Modeling

Interaction models for close-proximity pHRI are most likely multimodal because contact force and robot motion equally affect the perception of an interaction. Obtaining such models will probably have to rely on learning-from-demonstration (LfD), a.k.a. imitation learning, since it would be very difficult to model interactions analytically. However, using tactile data can be challenging for the learning process due to the high-dimensionality and sparsity.

For close-proximity pHRI, collecting training samples for LfD can be another challenge. Although in principle it is possible to use human-human interactions as samples, contact force and motion measurement as well as human-to-robot motion mapping are not straightforward. Since we need robot hardware with the properties mentioned in the previous subsection anyway, it makes more sense to directly control the robot by teleoperation. However, building a teleoperation interface for close-proximity physical interaction itself poses an interesting research question of how to provide the operator with soft, wide-area contact forces.

Evaluation

Evaluating the effect of close-proximity pHRI is also challenging, primarily because it is difficult to quantitatively define and measure its magnitude. Researchers therefore tend to rely on subjective feedback from users obtained by questionnaires. This fact also implies that any evaluation requires a large number of users to draw statistically significant conclusions.

Because one of the goals of close-proximity pHRI is to improve the emotional state of the human partner, it would be useful to quantitatively estimate the emotional state through sensor information. Many studies in psychology [12] have investigated the relationship between a person's emotional state and physiological data such as heart rate, skin resistance, and facial muscle activity. If these techniques can reliably detect subtle changes in the emotional state, it may allow us to evaluate interaction models quantitatively using a relatively small number of subjects.

Personalization

Current approaches for interaction modeling and evaluation discussed so far are also problematic for personalization due to the large number of required demonstrations and experiments: the resulting models tend to be generic and can only be statistically proven to be effective. Two key technical challenges need to be addressed in order to realize personalized close-proximity pHRI:

- Quantitative evaluation: A more quantitative method for evaluation, such as estimating the emotional state from physiological data, could accelerate the evaluation process compared to using subjective evaluation. This approach will enable personalization of an interaction model if the evaluation can be done for a specific person.
- Online learning: A possible approach for personalization is to start from a generic interaction model and then adapt it to personal preference. For this approach to work, however, an active learning algorithm combined with a model that can be updated online would be necessary.

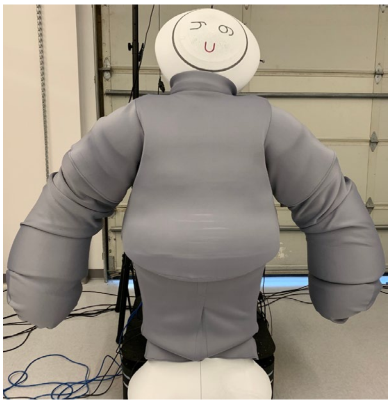

Figure 2: Robot hardware developed for hug interaction

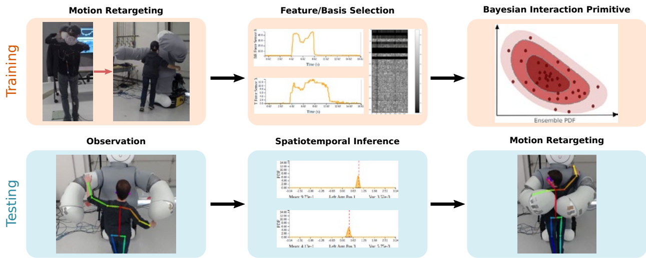

Figure 3: Overview of the learning framework

3 CASE STUDY: MODELING HUG INTERACTION

As a case study of close-proximity pHRI, we develop an LfD framework based on multimodal Bayesian inference for autonomous human-robot hug interaction [4]. The robot (Fig. 2) has two 6-degrees-of-freedom torque-controlled arms, and its torso and arms are covered by soft foam with a total of 61 embedded contact sensors. The arms have enough motion and torque ranges to enclose the human torso and exert decent contact force.

Learning from Demonstration

We apply multimodal Bayesian inference [5, 6] to model the relationship between the actions of the robot and human partner (Fig. 3). The model is capable of inferring the robot's action (joint angle commands and the arm contact forces) from the human partner's motion detected by the OpenPose library [7] and the contact force applied to the robot's body.

A major extension required for modeling hug interactions is feature selection to address high-dimensional but sparse tactile data. Feature selection utilizes two types of sparsity: *temporal sparsity* i.e. tactile sensors are activated only for a limited time during interactions, and *spatial sparsity* i.e. only a few tactile sensors are activated during a specific interaction. In [4], we demonstrate that dimensionality reduction

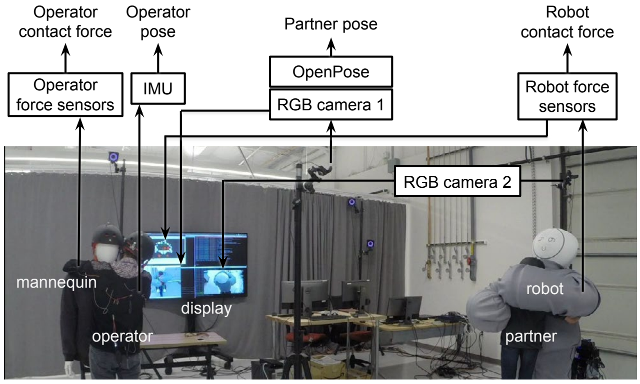

Figure 4: Teleoperation setup for collecting demonstrations

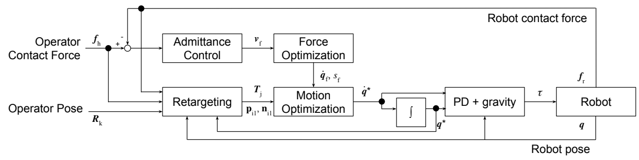

Figure 5: Controller overview

by feature selection does not cause statistically significant degradation of inference quality.

Demonstrations are generated by the teleoperation setup shown in Fig. 4. The operator has access to the human partner's pose and contact force applied to the robot's body through a display. In response to this information, s/he generates the motion and arm contact force commands by directly hugging a mannequin. The demonstrations are expected to provide the interaction model with basic "social norms" associated with hug: reciprocate the timing and intensity of the human partner's action. Because there is no commercially available haptic device that can display soft, wide-area contact forces, we opt for visualizing the contact force information as the size and color of the blobs displayed at the corresponding sensor locations.

Control

Controlling the robot given the motion and contact force commands has two main challenges. First, the problem is overconstrained because motion and force depend on each other even though both are critical for human perception. Our controller employs hierarchical optimization with contact force having the higher priority, while allowing some deviation from the commanded contact force to prevent the robot motion from deviating too much from the operator's.

The second issue is that the human partner's body size and shape may vary and therefore the contact states on the operator and robot sides may be different. If the human partner is large, for example, the robot may touch the human before

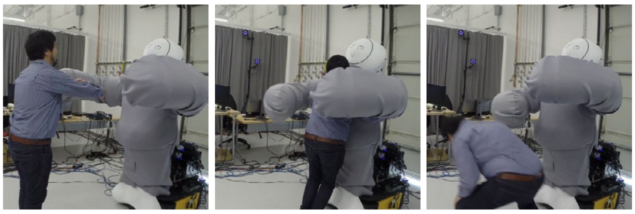

Figure 6: Edge cases. Left: hugging the air; the hug did not complete because the phase did not proceed further. Center: delay before hugging; hug was successful because the model correctly recognized the beginning of a hug. Right: hugging without making contact; the robot failed to release the partner because the model received conflicting information.

the operator touches the mannequin. To solve this issue, we use a simple motion retargeting algorithm that adjusts the robot motion based on arm contact force information only. Figure 5 depicts the overview of the controller [9].

Results

121 demonstrations were collected from 6 participants interacting with the robot teleoperated by the same operator. After training, the interaction model was tested on one of the participants from data collection as well as 2 new participants. Overall, approximately 82% of test cases resulted in successful hugs, which we define as the robot hugging the human participant and responding to their cues.

As shown in Fig. 6, the learned model was able to generate reasonable reactions to many of the edge cases where human partner's motion was completely different from the demonstrations. On the other hand, the interaction did not complete when the model received conflicting motion and contact force information.

4 CONCLUSION AND FUTURE WORK

In this paper, we first introduced the concept of *empathetic physical support*, in which physically capable robots provide both emotional and physical support through close-proximity pHRI including direct contacts. We then reviewed the technical challenges toward personalizing close-proximity pHRI such as quantitative evaluation and online learning.

In the second half of this paper, we presented a case study of close-proximity pHRI: modeling hug interactions through an LfD framework using demonstrations collected by a teleoperated robot. The robot is covered with soft skin and tactile sensors for comfort and contact force sensing. The learned model demonstrated the ability to generalize to a wide variety of human partners and hug styles.

Although we are able to obtain a generic hug interaction model, personalization would be required to address differences in culture, context, and emotional state where simply reacting to the human partner's action is not enough. If a person is depressed, for example, his/her hug may be weak and short but it does not necessarily mean that robot should respond with similar intensity. Rather, strong and reassuring hug may help overcome the depression.

In order to detect and model subtle differences caused by these aspects, we need more sophisticated methods for evaluating and teaching individual interactions. First, we will explore the possibility of quantitatively estimating the human partner's emotional state using the data from wearable sensor suit equipped with a number of physiological sensors. If successful, this technology will enable personalized evaluation of the effect of interactions. Secondly, we will develop a haptic device that can display soft, wide-area contact forces such as those experienced in a hug. This device would be an essential component of the system because subtle personal differences in the contact force pattern are difficult for the operator to recognize with the current system that simply visualizes the contact force information.

REFERENCES

[1] A. Alspach, J. Kim, and K. Yamane. 2015. Design of a Soft Upper Body Robot for Physical Human-Robot Interaction. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. 290–296.

[2] A.E. Block and K.J. Kuchenbecker. 2018. Emotionally Supporting Humans Through Robot Hugs. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. 293–294.

[3] A.E. Block and K.J. Kuchenbecker. 2019. Inflatable Haptic Sensor for the Torso of a Hugging Robot. Work-in-progress paper (2 pages) presented at the IEEE World Haptics Conference (WHC).

[4] J. Cambell and K. Yamane. 2020 (accepted). Learning Whole-Body Human-Robot Haptic Interaction in Social Contexts. In *IEEE International Conference on Robotics and Automation*.

[5] J. Campbell and H. Ben Amor. 2017. Bayesian Interaction Primitives: A Slam Approach to Human-Robot Interaction. In *Conference on Robot Learning*. 379–387.

[6] J. Campbell, S. Stepputtis, and H. Ben Amor. 2019. Probabilistic Multimodal Modeling for Human-Robot Interaction Tasks. In *Robotics: Science and Systems*.

[7] Z. Cao, G. Hidalgo, T. Simon, S.E. Wei, and Y. Sheikh. 2018. OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields. In *arXiv preprint arXiv:1812.08008*.

[8] J.E. Colgate, J. Edward, M.A. Peshkin, and W. Wannasuphoprasit. 1996. Cobots: Robots for Collaboration with Human Operators. In *ASME Intl. Mechanical Engineering Congress and Exposition*. 433–439.

[9] A. Kaplish and K. Yamane. 2019 (in press). Motion Retargeting and Control for Teleoperated Physical Human-Robot Interaction. In *IEEE-RAS International Conference on Humanoids Robots*. Toronto, Canada.

[10] P. Mittendorfer and G. Cheng. 2011. Humanoid Multimodal Tactile-Sensing Modules. *IEEE Transactions on Robotics* 27, 3 (2011), 401–410.

[11] M. Shiomi, A. Nakata, M. Kanbara, and N. Hagita. 2017. A Hug from a Robot Encourages Prosocial Behavior. In *IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 418–423.

[12] J. Wagner, J. Kim, and E. André. 2005. From Physiological Signals to Emotions: Implementing and Comparing Selected Methods for Feature Extraction and Classification. In *IEEE International Conference on Multimedia and Expo*. 940–943.

Motivating Incremental, Personalized Models of Human Behavior for Structured Environments

Christopher K. Fourie
ckfourie@csail.mit.edu
Massachusetts Institute of Technology
Cambridge, Massachusetts

Przemyslaw A. Lasota
plasota@csail.mit.edu
Massachusetts Institute of Technology
Massachusetts

Julie A. Shah
julie_a_shah@csail.mit.edu
Massachusetts Institute of Technology
Cambridge, Massachusetts

ABSTRACT

Models of human behavior are critical to the fluent interaction of a human and a robot. Typical approaches to action and intent recognition, however, rely on explicit action labels, negating the ability to adapt to new data without human intervention. Further, we provide preliminary evidence that timing models of human behavior transfer poorly between individuals, supporting an argument for individual modeling. We argue for the use of an event-based framework for understanding a human's interaction with their environment, to support and enable incremental profiling, the automatic segmentation of data, and dynamic robot behavior in Human-Robot Interaction. We introduce events as a semantic representation when modeling a human's behavior, and recast the traditional problems of action recognition, segmentation and prediction using events.

CCS CONCEPTS

• **Human-centered computing** → **Interaction design theory, concepts and paradigms**; • **Computer systems organization** → **Robotics**.

KEYWORDS

human modeling, intent recognition, activity segmentation, motion prediction, human-robot interaction

1 INTRODUCTION

Humans have the capacity to work extremely effectively with other humans, easily adapting to and predicting a collaborator's behavior. This process has been widely studied - a field of Cognitive Psychology, the theory of Joint Action [11], studies the coordinated activity of people to achieve joint goals. A human dyad's capacity to anticipate, react and coordinate is largely enabled through the complementary prediction of the spatial, temporal, and semantic aspects of a collaborator's behavior [12].

Workshop on Behavioral Patterns and Interaction Modelling for Personalized Human-Robot Interaction 2020, March 23–26, 2020, Cambridge, United Kingdom

The prediction of human behavior has been widely studied in Human-Robot Interaction (HRI). Numerous systems, particularly those developing a mechanism for *anticipation* [1, 4, 8] utilize predictions of human behavior to reason on the action that a robot should take. These capabilities typically express and operate on some form of semantic prediction, but they are also occasionally combined with either spatial or temporal aspects of predictions to further improve the robot's behavior. Understanding where a person will be in the future can help a robot plan its actions [7, 10], while understanding when a human will complete a current action can help a robot schedule its own action [2, 9].

Our previous work has focused on developing prediction capabilities within structured environments, such as those typically found within manufacturing [5, 13]. Yet, even within these domains, prediction remains a challenge and traditional approaches to modeling humans pose several limitations. Traditional action or intent recognition approaches treat the recognition of an action or plan as a machine learning problem, relying on explicit labels and supervised learning techniques. At present, actions tend to be broadly defined, and are typically chosen by the algorithm designer.

Consider a simple assembly process, in which five parts must be collected. Assuming one action per part (the act of retrieving it), and without constraining the order in which parts are collected, there are 5! = 120 sequences of actions a person could exhibit. An individual is unlikely to exhibit every plausible sequence of actions, and a typical approach to ensure maximal coverage would be to collect data from multiple people. While such an approach could be exhaustive, it may overgeneralize. The use case focuses on an individual, and much of the training data is likely to be irrelevant - potentially confounding results.

We will demonstrate that models of the temporal behavior of individuals transfer poorly to other individuals, and even to variations on the task for the same individual (i.e. when swapping the order of similar reaching motions). However, the same individuals demonstrate consistency within their own action timings. This has implications for prediction methodologies that rely on time-series to inform their models, but do not contain information tailored to the individual in question. In these cases, the data can fail to be

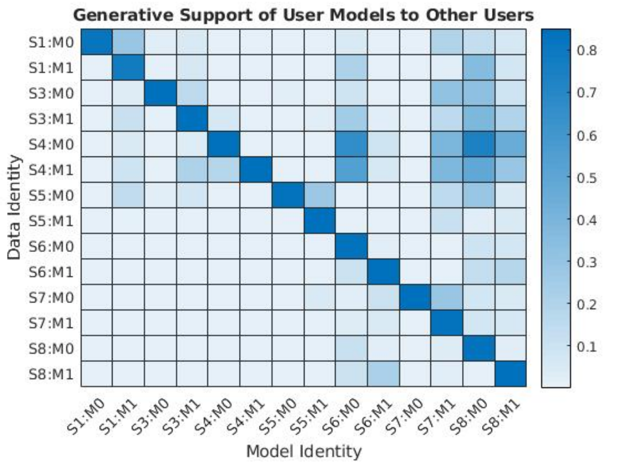

Figure 1: An illustration of the "support" a model gives to another dataset (see text). Note that user profiles did not transfer well across participants.

representative, and algorithms that work well when trained on some people may fail when tested on others. Worse, these models are unable to adapt to a person: when people are consistent, the algorithm may perform consistently poorly.

Traditional approaches have limited capacity for incrementally improving and updating models of human behavior. Without an oracle or extremely robust recognition and/or labeling function, such capabilities lie beyond the purview of most algorithms. Scalability is also problematic - reconsidering the earlier example, at 8 parts and a single, non-repeating action per part, the number of unique sequences of actions rises to $8! = 40320$, limiting the practicality of an approach centered on data collection and offline annotation.

We motivate a unified event-based modeling approach that is appropriate for accommodating the demands and requirements of incremental human behavioral modeling, and particularly for enabling responsive, intelligent robot behavior that adapts to a user's variations in behavior. Events are dense semantic representations indicating key points in an activity or task, and occur at discrete time instances of the task. We use these events for rigorous definitions of the prediction problems and discuss the merits of using an event-based design for enabling incremental human profiling. We demonstrate that temporal behavioral patterns appear to be unique to individuals, and argue for a shift toward incremental, personalized behavioral models in HRI.

2 VARIATIONS IN ACTION DURATION

Our own work in manufacturing settings has provided us with observations of people performing structured, repetitive tasks in routine jobs. We have observed that people develop consistent behavioral patterns (e.g. ordering and timing) but that these patterns varied widely between people.

These observations suggested a need for personalization and inspired a short data collection of several participants performing a simple reaching task. The task involved placing 8 bolts into 8 linearly spaced holes in the assembly in front of them. The bolt holes were within reaching distance, and the bolts were placed directly in front of the participant. The participants were free to select the order in which the bolts were placed, but were requested to maintain the initially chosen order for 20 demonstrations. This was then repeated for a second, different order of bolt placements. Participants were encouraged to perform the task at a natural pace, with no time constraint. The data was collected with a PhaseSpace motion capture system and separated into actions by detecting when the hand entered or exited a prespecified radius of the bolthole. The prespecified radius was identical across participants.

To evaluate the temporal consistency of the participants, we constructed a model based on a multivariate Gaussian. For each action sequence, for each participant, we construct a multivariate Gaussian distribution, approximated using the sample mean and covariance calculated from the vectors of action durations (Δt_i) of the participant performing the task. Considering the standard deviation of an action for each participant, we saw that all participants varied, on average, by 15% of the mean time for their individual actions. The action durations between participants, however, differed greatly.

Our primary consideration is whether timing models transferred well between individuals. To establish this, we calculated a transfer statistic as follows:

$$\gamma_{i,j} = \exp\left(\frac{\alpha}{N_i}\sum_{k=1}^{N_i}(\Delta t_k - \mu_j)^T\Sigma_j^{-1}(\Delta t_k - \mu_j)\right)$$

α is an arbitrary scaling coefficient intended for numerical conditioning. The resulting matrix of values represents the transformed likelihood of a *model* built for a participant performing an action sequence (j) to the *raw data* collected for a participant performing a particular action sequence (i). We term this score the "support" of a profile built on one dataset to another dataset. This is shown graphically in Figure 1. The profiles lend "support" to the data from which they were generated, but not to data from which they were not. The implication is that these profiles do not transfer well between subjects (or even across action sequences per a particular subject) and that it may be necessary to learn a timing profile for each person that interacts with such a system. Note that data relating to Subject S2 has been omitted as the subject failed to follow study protocol.

3 EVENTS AS A UNIFIED MODELING REPRESENTATION

Events are a natural representation for embedding semantic information within a time-series. To accommodate personalization, individual preferences, and support automatic data segmentation, we advocate for the use of events as a *unified modeling representation*. We consider a unified modeling representation as one that enables incremental behavioral modeling and the prediction of spatial, semantic, and temporal aspects of human behavior.

We consider events to be categorical variables that are the output of a function $e_i = f_i(\lambda)$, where λ is the set of variables relevant to the detection of an event. Events are a function of the behavior of a person in a particular space. An event occurs at a particular time, t_{e_i}, is understood to be *causal* (in that it is dependent on the behavior of a particular person), and is assumed to be *observable*. An event could be the act of an associate triggering a light curtain or it could be a pose that a person transitions through in a motion (such as in the context of the work in Xia et al. [14] or Hayes and Shah [3]).

Each action is defined by two events - one at the beginning and one at the end of the action: $a_{ij} : e_i \longrightarrow e_j$. The events act to *segment* the actions, while also defining them. The number of explicitly detectable actions $\mathbf{A} = \{a_{11}, a_{12}, \ldots\}$ within this framework is the square product of the number of events, $\mathbf{E} = \{e_1, e_2, \ldots\}$ (i.e. $|\mathbf{A}| = |\mathbf{E}|^2$), while the number of possible event *sequences* is the factorial of the number of events.

The labeling of actions happens automatically based on observed events, negating the need for the manual annotation of actions and instead requiring the detection of an event. Events, in a similar fashion to the traditional action approach, require explicit definitions from the system designer for implementation. However, events are far easier to detect than actions: for instance, consider a person entering a particular space compared to predicting a person entering the same space in the future. Detecting the event could be accomplished through a standard industrial sensor, while detecting a human walking towards the space poses significant machine learning challenges.

Events are sources of semantic and temporal information within the framework. They exist at a level of abstraction that allows fluents or logical propositions to be effectively formed, allowing for representation within Linear Temporal Logic (LTL) formulations or Probabilistic Domain Definition Language (PDDL) plans.

However, the introduction of events leads to a transformation of the traditional problems of action recognition, segmentation and prediction:

- **Action Segmentation:** this is recast as event detection, in which a set of events, $\mathbf{e}$, is considered latent within the data λ:
$$\hat{\mathbf{e}} = \arg\max_{\mathbf{e}} p(\mathbf{e}|\lambda)$$
- **Action Recognition:** this is recast as the prediction of a future event, e_j, given the past event e_i and the data λ:
$$\hat{a}_{ij} = \hat{e}_j|e_i = \arg\max_{e_j} p(e_j|e_i, \lambda)$$
- **Action Prediction:** this is recast as predicting the *set* of future events $\mathbf{e}_p$, given the set of transpired events, $\mathbf{e}$, and the data λ:
$$\hat{\mathbf{e}}_p = \arg\max_{\mathbf{e}_p} p(\mathbf{e}_p|\mathbf{e}, \lambda)$$

Action segmentation implies the detection of events latent within the data available to the system. While we have so far considered explicit forms of events, the problem could be regarded to extend to *implicit* events - the problem of *discovering* events. Discovering new events and inferring functions for their detection would increase the granularity of a robot's model of a human's interaction with the environment, posing a challenging but rewarding inference problem.

Similar to the manner in which an action within a PDDL plan results in a change in the fluents forming the state, a correctly identified action within our framework implies the occurrence of a future event. As a result, the recognition of an ongoing action is the prediction of an upcoming event. Standard approaches remain applicable, and the event based modeling approach allows for the automatic segmentation of training data, allowing methods for action recognition to be retrained after a person interacts with the system. The framing also allows for a level of robustness and recovery: in the case that an action is incorrectly recognized, the eventual detection of the event signaling the end of the action allows the correct label for the action to be inferred.

Action prediction, or the prediction of the set of upcoming actions (i.e. intent/plan), can be understood to be the prediction of a set of upcoming events. This remains remarkably close to existing techniques, particularly those based on PDDL-style representations of goals and states (e.g. Freedman and Zilberstein [1] or Jain and Argall [6]). As previously noted, the semantic information latent in events allows for the construction of fluents that can be directly integrated into such approaches.

However, the event based framing also allows for the introduction of problems related to *timing*. These problems are of particular relevance to a robot operating in close proximity to a human, where the human's action is contingent on that of the robot's. The knowledge of when a human will finish a particular action or require a robot's action to be completed (e.g. for handover), allows a robot to optimize or schedule its actions to maximize fluency (the synchronous meshing of human and robot actions [4]). We further define two problems related to timing:

- **Short-term Temporal Prediction:** detect the time, $\hat{t}_{e_j}$, at which the predicted next event, $\hat{e}_j$, will occur, given the previously observed event, e_i, the time at which it occurred t_{e_i}, and the system data λ :

$$\hat{t}_{e_j} = \arg\max_{t_{e_j}} p(t_{e_j} | t_{e_i}, \hat{e}_j, e_i, \lambda)$$

- **Long-term Temporal Prediction:** detect the set of times, $\hat{\mathbf{t}}_{\mathbf{e}_p}$, at which the set of predicted future events $\hat{\mathbf{e}}_p$, will occur, given the previously observed events, $\mathbf{e}$, the times at which they occurred, $\mathbf{t}_\mathbf{e}$, and the system data λ:

$$\hat{\mathbf{t}}_{\mathbf{e}_p} = \arg\max_{\mathbf{t}_{\mathbf{e}_p}} p(\mathbf{t}_{\mathbf{e}_p} | \mathbf{t}_\mathbf{e}, \hat{\mathbf{e}}_p, \mathbf{e}, \lambda)$$

The first temporal prediction problem focuses on detecting the time that the next event will occur (i.e. when the current action will conclude). The knowledge allows a robot to appropriately plan its short term motions, ensuring that it concludes its own action at an appropriate time. The second temporal prediction problem focuses on detecting the time at which the sequence of upcoming events will occur. To maintain fluency, the robot must have a reasonable estimate of when a person will perform certain actions, to enable the robust scheduling of its own.

Finally, we note the role of spatial-temporal predictions in robot motion planning. Prior work in spatial prediction has focused on the prediction of the human's trajectory, $\hat{x}_p$, using the observed trajectory up to that point, x_q. We define the spatial prediction problem using the event-based framing:

- **Spatial Prediction:** predict trajectory, $\hat{x}_p$, using knowledge of the previous event e_i, the predicted next event, $\hat{e}_p$, and the system data λ:

$$\hat{x}_p = \arg\max_{x_p} p(x_p | \hat{e}_j, e_i, \lambda)$$

The spatial prediction leverages the previously observed event and predicted future event to only consider spatial data of relevance to the prediction. This helps to ensure computational tractability in both the robot's planning and in the generation of a dense spatial-temporal prediction. Maintaining a multimodal hypothesis space can be accomplished with a distribution over the prediction activities, while predicting beyond the horizon of the predicted future event could be accomplished using short-term motion propagation techniques at the boundary between events.

4 CONCLUSION

We advocate for the use of event-based modeling to support incremental human profiling and the automatic labeling and segmentation of data in structured environments. We describe limitations in traditional approaches, provide preliminary evidence of the suboptimality of general temporal models when predicting individuals, and recast the traditional prediction problems to use events. We believe that this can enable robot behavioral adaptation to user preferences.

REFERENCES

[1] Richard G Freedman and Shlomo Zilberstein. 2017. Integration of Planning with Recognition for Responsive Interaction Using Classical Planners.. In *AAAI*. 4581–4588.

[2] Kelsey P Hawkins, Shray Bansal, Nam N Vo, and Aaron F Bobick. 2014. Anticipating human actions for collaboration in the presence of task and sensor uncertainty. In *International Conference on Robotics and automation (ICRA)*. IEEE, 2215–2222.

[3] Bradley Hayes and Julie A Shah. 2017. Interpretable models for fast activity recognition and anomaly explanation during collaborative robotics tasks. In *International Conference on Robotics and Automation (ICRA)*. IEEE, 6586–6593.

[4] Guy Hoffman and Cynthia Breazeal. 2007. Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team. In *International Conference on Human-Robot Interaction*. ACM.

[5] Tariq Iqbal, Shen Li, Christopher Fourie, Bradley Hayes, and Julie A Shah. 2019. Fast online segmentation of activities from partial trajectories. In *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 5019–5025.

[6] Siddarth Jain and Brenna Argall. 2019. Probabilistic Human Intent Recognition for Shared Autonomy in Assistive Robotics. *ACM Transactions on Human-Robot Interaction (THRI)* 9, 1 (2019), 1–23.

[7] Yun Jiang and Ashutosh Saxena. 2014. Modeling High-Dimensional Humans for Activity Anticipation using Gaussian Process Latent CRFs.. In *Robotics: Science and systems*. 1–8.

[8] Przemyslaw A Lasota and Julie A Shah. 2015. Analyzing the effects of human-aware motion planning on close-proximity human–robot collaboration. *Human factors* 57, 1 (2015), 21–33.

[9] Steven James Levine and Brian Charles Williams. 2018. Watching and Acting Together: Concurrent Plan Recognition and Adaptation for Human-Robot Teams. *Journal of Artificial Intelligence Research* 63 (2018), 281–359.

[10] Jim Mainprice and Dmitry Berenson. 2013. Human-robot collaborative manipulation planning using early prediction of human motion. In *International Conference on Intelligent Robots and Systems (IROS)*. IEEE.

[11] Natalie Sebanz, Harold Bekkering, and Günther Knoblich. 2006. Joint action: bodies and minds moving together. *Trends in cognitive sciences* 10, 2 (2006), 70–76.

[12] Natalie Sebanz and Guenther Knoblich. 2009. Prediction in joint action: What, when, and where. *Topics in Cognitive Science* 1, 2 (2009).

[13] Vaibhav V Unhelkar, Przemyslaw A Lasota, Quirin Tyroller, Rares-Darius Buhai, Laurie Marceau, Barbara Deml, and Julie A Shah. 2018. Human-aware robotic assistant for collaborative assembly: Integrating human motion prediction with planning in time. *IEEE Robotics and Automation Letters* 3, 3 (2018), 2394–2401.

[14] Lu Xia, Chia-Chih Chen, and Jake K Aggarwal. 2012. View invariant human action recognition using histograms of 3d joints. In *Computer vision and pattern recognition workshops (CVPRW), 2012 IEEE computer society conference on*. IEEE, 20–27.

Adaptivity as a Service (AaaS): Enabling Deep Personalisation for a Heterogeneous Ambient Assisted Living Landscape

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

ABSTRACT

The need for personalisation of *Ambient Assisted Living* (AAL) solutions is widely recognised in current research activity, with a variety of approaches to embedding user preferences and needs, and to modelling various aspects of the user. However, a strong focus on lab-based development and evaluation, as opposed to in situ deployment, has made it easy to overlook the reality of future AAL: a future where a variety of heterogeneous platforms, devices and robots will need to share their experiences, within and outside the home. Without scalability, AAL will struggle to succeed. *Adaptivity as a Service* (AaaS) poses that deep individual personalisation is better managed by a highly specialised standalone service, rather than by individual smart home platforms and devices.

CCS CONCEPTS

• **Applied computing** → *Health informatics*; • **Computer systems organization** → *Distributed architectures*; • **Human-centered computing** → **Ambient intelligence**.

KEYWORDS

Deep Personalisation, Adaptivity, Ambient Assisted Living (AAL), Digital Twin, Human-in-the-Loop (HITL)

1 INTRODUCTION

This position paper argues that current approaches to personalisation in Ambient Assisted Living systems, in terms of user needs and preferences, fail to consider real world scalability implications of their solutions in the context of a heterogeneous AAL landscape.

AAL is an umbrella label for technology and processes implemented within home and care facilities to support the elderly with the intent to create a more natural environment to sustain physical and mental health. Existing examples rely on wearables and/or sensors embedded throughout the environment to track movement, detect changes in the individual's health status (through historical data analysis), alert carers of falls, and trigger interventions when needed. Autonomous, interactive, and robotic technologies are increasingly proposed as part of AAL provision, for their potential to bring active monitoring as well as social and cognitive support and responsive physical assistance.

Adaptivity is a term covering virtually all aspects of customisation in an AAL context. AAL systems must adapt to the changing habits, situations, individual preferences and evolving needs of their users: from the way they interpret user activities, from sensor data, to the way they provide them with personalised social and physical interaction.

Consider that traditional approaches view personalisation and customisation as things that have to be baked in and custom-engineered (often as an afterthought) for any given AAL system, the following **position statement** is presented: *personalisation in Ambient Assisted Living (AAL) should be delivered primarily by a dedicated personalisation service that uses a Digital Twin (DT) to model the user throughout their life, consolidating user data to enable immediate adaption of compatible systems to a user's wants and needs.*

To accomplish this, AaaS sandboxes adaptivity into its own research domain in which personalisation is the number one priority.

2 CURRENT APPROACHES IN AAL

A variety of bespoke approaches to personalisation and user modelling exist, in the domain of AAL. The concept of AaaS draws from many ideas in previous work that, although sound in principle, have been constrained by the larger system within which they exist.

Personalisation

Evidence suggests that the preferred mode of interaction between humans and robotic smart home environments is

Workshop on Behavioral Patterns and Interaction Modelling for Personalized Human-Robot Interaction 2020, March 23–26, 2020, Cambridge, United Kingdom

context-dependent [7]. Smart homes, upon which AAL solutions are built, provide unique opportunities to harness distributed perception and multimodal interaction, which means there is no need to commit to a single mode of interaction when it comes to enabling personalisation.

‘Hybrid’ approaches are well established in *Human Activity Recognition* (e.g. [2]): hybrid models fuse knowledge- and data-driven sources to enable adaption to individual users and improve baseline performance and scalability. This ultimately boils down to starting with initial ‘seeds’ (templates) of what a system can recognise or do (from knowledge engineering at design time), while data collection in situ enables a semi-supervised learning process to grow capability over time. Such approaches have improved personalisation by reducing user burden in having to specify activities manually. This is especially useful in ‘if this, then that’ approaches, where an activity is commonly the ‘if this’.

Users themselves are best placed to evaluate their own needs, and so a number of recent approaches propose the use of *Graphical User Interface* (GUI) for both control and interaction. Examples exist in projects such as AAL platform ‘RADIO’ [1] and in the user-driven customisable companion/care robot featured in ‘Teach Me-Show Me’ [6]. Although a step forward, configuration options in these approaches are limited to those pre-programmed at design time.

Existing approaches encapsulating multi-modal personalisation include the *Global Public Inclusive Infrastructure* (GPII), which seeks to enable automatic interface personalisation on a large scale [10]. However, GPII is not geared towards robotics, where an opportunity exists for personalisation to be handled primarily in the planning domain.

Another facet of personalising AAL is co-design (aka participatory design), wherein end users and key stakeholders (e.g. primary/secondary caregivers) are heavily involved in the design phase of an assistive product or service [3]. While the resulting solution may itself include methods of personalisation, like those seen thus far, co-design effectively ‘bakes in’ customisations for a specific group of individuals.

User Modelling

There are now two common themes in user modelling in AAL: sets of pre-defined user models, and ontology-based profiles. The former typically use several predefined templates of potential users (e.g. "dependent, assisted, at risk, and active" in the approach described in [9]) to inform the system’s behaviour. Model-based approaches struggle to fully represent elderly individuals, who have a wide range of care needs. The latter, ontology-based, offer greater flexibility and extensibility. They may, for example, capture personal, health, and preference data, such as in [8]. A downside of this approach is that to utilise the ontology, third party developers require intricate knowledge of its structure.

In the ‘GrowMeUp’ project, a profiling mechanism is included to capitalise on the use of multiple ‘GrowMu’ social robots. Using a pre-defined profile schema, robots gathered and collated data to expand profile information [4]. User routines were also detected in a similar fashion [5]. This is, in effect, a hybrid method insofar as it starts with an initial minimal defined profile and fleshes it out with learned data.

3 RESEARCH GAPS

A significant research gap exists relating to evaluating the value of a highly specialised personalisation and adaptivity service that can operate with a range of smart home platforms, devices, and robots. This service will need to provide a high degree of personalisation in the face of heterogeneity.

One such characteristic to consider is the range of interaction modalities on offer in AAL, which should leveraged for personalisation. Personalisation will need to account not only for user preferences within interactions, but also preferences regarding the mode of interaction itself.

It is unreasonable to expect a user to contend with various personalisation processes from different platforms and services at the same time, and so personalisation for the user should ultimately become a ‘create once, publish everywhere’ experience. Decoupling personalisation from individual platforms and devices will require the establishment of a high-level formalism for personalisation and adaptivity, effectively a common language to allow one-to-many translations of service functionality.

A novel approach seeking to fulfil this role should cater for a range of essential adaptions and incorporate the user model, in order to maximise knowledge reuse.

A Role for AaaS

In light of the given research gaps, consider some example scenarios to demonstrate the role that could be filled by AaaS. First, some background information: Jack is 72 and completely deaf in his left ear due to an industrial accident earlier in life, with age related hearing loss in his other ear. He also has generally poor mobility and struggles to walk unaided. Jack’s smart home is comprised of low-cost *off-the-shelf* (OTS) equipment that is not specialised for his needs.

(1) **Scenario:** Jack begins standing to head for the bathroom. The AAL system has no plan for this scenario. **Adaption:** based on Jack’s profile information, AaaS has automatically selected an assistive policy for this scenario. Having learned Jack’s routine, a robot is already nearby to guide Jack to the bathroom. AaaS has learned that Jack does not like robots to step into the bathroom. It steps aside at the bathroom door, waiting to help him return to his chair.

(2) **Scenario:** Jack is sitting in his chair when someone rings his doorbell. The smart home recognises the visitor as his friend Victor and plans to dispatch a mobile robot with the intention to announce Victor's arrival. **Adaption:** AaaS knows of Jack's disabilities and modifies the original plan, so that the robot knows to: stand on Jack's right side, make the announcement at high volume, and display an image of Victor on its screen.

4 ADAPTIVITY AS A SERVICE

AaaS addresses three types of long-term adaptation in AAL systems, namely:

(1) Adapting context-awareness itself to account for predicted physical and mental decline, based on individual's known conditions and principles of ageing.
(2) Adapting assistive functionality, including interaction modalities, to fit an individual's exact needs/wants.
(3) Adapting quickly to new users, based on experience.

AaaS slots into AAL platforms as a high-level intermediary between the local control/decision making, and sensing/actuating devices. Elements are divided between the local level, the point of service delivery, and the global level, which provides data aggregation and centralised learning.Within the home, AaaS can intercept communication between these components. Most user data, including user models, exist primarily in the global level, where the user is considered as one of many in pool of individuals requiring personalised and adapted AAL. The collocation of this data, where an individual is part of a comparable population, accelerates learning potential and the amount of useful knowledge generated. Consequently, the ins and outs of explicit social interactions are not of concern to AaaS, but of the specific social and/or robotic agents operating in the home.

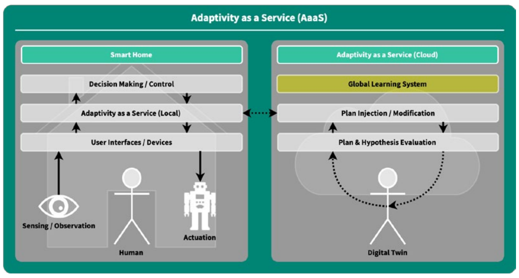

Figure 1: High-level architecture of AaaS. Note the division of labour across the local (left) and global (right) layers.

AaaS addresses the three types of adaption with a two-pronged approach to *Human-in-the-Loop* operation: (1) implicit consideration of the user in all decision making, and (2) continuous learning of preferences and desired behaviours to update locally (user) and globally (system) held beliefs. Figure 1 shows a (simplified) conceptual architecture of AaaS.

5 A HYBRID MODEL

Hybrid models should be employed in AaaS for both user modelling and personalisation, with their implementation intrinsically linked. AaaS relies on a *Digital Twin* (DT) and a *Human-in-the-Loop* (HITL) approach to, respectively, represent and modify initial templates for both aspects.

Digital Twin

Fundamental to enabling 'deep' personalisation in AaaS is the Digital Twin: a model of the user generated based on profile information gathered from the user and/or caregivers. When a new user is created, a set of initial assistive policies are generated and assigned to the DT, which can then be adapted to suit.

Assistive policies are nested high-level plans that are originally created by humans and added to a global policy bank. These plans describe: (1) assistive services that can be carried out in a robotic AAL environment, as well as the tunable parameters in those policies, and (2) modifications that can be made to external plans to accommodate specific user wants/needs, again including tunable parameters. These plans are the product of co-design: a combination of expert input with input from the target audience on *what types* of assistance/modifications they would find useful, and *which parameters* they would want to personalise. Importantly, they do not rely on specific hardware, but on translating plans into commands for execution with commonly available APIs.

Over time, through an adaptivity process explained herein, these plans are adapted and the DT becomes increasingly attuned to the wants and needs of an individual. This has short- and long-term forecasting implications. Short-term, AaaS is able to assess whether actions planned by the local AAL system (which may do little to *actively* adapt to the user) 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.

A key benefit of AaaS therefore lies in how the DT enables knowledge reuse within and across users. Since user preferences are recorded in high-level terms, it is possible to re-apply them in virtually any environment (e.g. when moving home or purchasing a new robot). Similarly, when a new user begins using AaaS, it is possible to identify existing users with similar traits and generate initial policies based on prior experience with those users, so that each user starts with policies already evolved from the baseline. Transferred knowledge between users therefore includes both types of assistive policy, since policies modified and rules applied by other users can be transferred verbatim, thanks to their high-level implementation. This ultimately produces a benefit of reduced knowledge engineering as a single policy specified at design time is evolved into many variations for later use.

Policy Personalisation

It is envisaged that assistive policies be updated using *Human-in-the-Loop* (HITL) *Reinforcement Learning* (RL) approaches. In essence, modification of policies initially assigned to a user (at enrolment) will be feedback-driven. Globally held beliefs about the impact of psychological traits (e.g. relating to *Mild Cognitive Impairment* [MCI]) can be updated via data-driven aggregation of real-world experiences.

Over time, policies for specific scenarios will settle on what the user *wants* (note the emphasis over *needs*), which raises its own challenges. This is a key differentiator between AaaS and some prior approaches such as GPII: personalisation goes beyond meeting accessibility requirements. Once evolved, these policies are (continuously) fed back to the global system both to enhance the user's Digital Twin, and to reuse learning to benefit other users. For instance, this feedback enables automatic policy selection for new users, drawing from experiences with similar existing users.

6 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 a variety of 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?

7 CONCLUSION

The fundamental principles of AaaS have been outlined here in relation to addressing existing challenges in AAL personalisation. There is novelty in proposed HITL operation across a distributed architecture with local instances (e.g. individual homes) and the cloud. Within AaaS, an individual home can benefit from and contribute to a wider network of adaptivity specialisation. Significant future research will focus on the best approaches to meet key goals of AaaS. Ultimately, the Digital Twin will serve as a rich source of data that accompanies a user for life, which systems that deal with personalisation independently may fail to replicate.

ACKNOWLEDGMENTS

This work was supported by the Engineering and Physical Sciences Research Council (grant EP/L016834/1), EPSRC Centre for Doctoral Training in Robotics and Autonomous Systems at Heriot-Watt University and the University of Edinburgh.

REFERENCES

[1] Christos Antonopoulos et al. 2015. Robots in assisted living environments as an unobtrusive, efficient, reliable and modular solution for independent ageing: The RADIO perspective. In *International Symposium on Applied Reconfigurable Computing* (Bochum). Springer, 519–530.

[2] Liming Chen, Chris Nugent, and George Okeyo. 2013. An ontology-based hybrid approach to activity modeling for smart homes. *IEEE Transactions on human-machine systems* 44, 1 (2013), 92–105.

[3] Ezequiel Duque, Guilherme Fonseca, Heitor Vieira, Gustavo Gontijo, and Lucila Ishitani. 2019. A Systematic Literature Review on User Centered Design and Participatory Design with Older People. In *IHC* (Vitória, Espírito Santo, Brazil) *(IHC '19)*. Association for Computing Machinery, New York, NY, USA, Article 9, 11 pages. https://doi.org/10.1145/3357155.3358471

[4] Goncalo S. Martins, Luis Santos, and Jorge Dias. 2017. BUM: Bayesian user model for distributed social robots. In *RO-MAN*. IEEE, Lisbon, 1279–1284. https://doi.org/10.1109/ROMAN.2017.8172469

[5] Luis Santos and Jorge Dias. 2016. User Routine Model Using a Cloud-Connected Social Robot. In *2016 5th IEEE International Conference on Cloud Networking (Cloudnet)*. IEEE, Pisa, Italy, 182–187. https://doi.org/10.1109/CloudNet.2016.34

[6] Joe Saunders, Dag Sverre Syrdal, Kheng Lee Koay, Nathan Burke, and Kerstin Dautenhahn. 2016. "Teach Me–Show Me"—End-User Personalization of a Smart Home and Companion Robot. *IEEE Transactions on Human-Machine Systems* 46, 1 (Feb. 2016), 27–40. https://doi.org/10.1109/THMS.2015.2445105

[7] Birte Schiffhauer, Jasmin Bernotat, Friederike Eyssel, Rebecca Bröhl, and Jule Adriaans. 2016. Let the User Decide! User Preferences Regarding Functions, Apps, and Interfaces of a Smart Home and a Service Robot. In *ICSR*. Springer, 971–981.

[8] Kerry-Louise Skillen et al. 2014. Ontological user modelling and semantic rule-based reasoning for personalisation of Help-On-Demand services in pervasive environments. *Future Generation Computer Systems* 34 (May 2014), 97–109. https://doi.org/10.1016/j.future.2013.10.027

[9] Stefanos Stavrotheodoros, Nikolaos Kaklanis, and Dimitrios Tzovaras. 2018. A personalized cloud-based platform for AAL support to cognitively impaired elderly people. *IFMBE Proceedings* 66, 2 (2018), 87–91. https://doi.org/10.1007/978-981-10-7419-6_15

[10] Gregg C. Vanderheiden, Jutta Treviranus, Manuel Ortega-Moral, Matthias Peissner, and Eva de Lera. 2014. Creating a Global Public Inclusive Infrastructure (GPII). In *Universal Access in Human-Computer Interaction. Design for All and Accessibility Practice*. Springer International Publishing, Cham, 506–515.

IMPRINT

Editors
Kathrin Pollmann | kathrin.pollmann@iao.fraunhofer.de
Daniel Ziegler | daniel.ziegler@iao.fraunhofer.de

Graphic design
Milena Velić

Contact
Fraunhofer Institute for Industrial Engineering IAO
Human-Technology Interaction
Nobelstraße 12
70569 Stuttgart
www.hci.iao.fraunhofer.de