

RADBOUD UNIVERSITY NIJMEGEN



FACULTY OF SOCIAL SCIENCES

Personalization of a Social Health Care Companion

THESIS MSc ARTIFICIAL INTELLIGENCE

Author:
Britt DECKERS
s1007564

Supervisor:
dr. Martijn VASTENBURG

Second reader:
dr. Serge THILL

July 2023

Abstract

The objective of this thesis is to gain insights into the benefits, limitations, and future directions of personalization techniques in social robotics. The study focuses specifically on Lizz, a digital healthcare companion, within the context of exercise recommendation during rehabilitation. The aim is to contribute to the understanding of this field by exploring the various aspects of personalization in social robotics. Three different prototypes are tested to showcase different mechanisms of personalization: a rule-based system without user feedback, a combination of rule-based and reinforcement learning with user feedback, and a combination of rule-based and reinforcement learning with user feedback and other dynamic factors. The prototypes aim to demonstrate the advantages and disadvantages of different personalization approaches and contribute to improved well-being and recovery outcomes. The findings indicate that greater personalization leads to exercise recommendations that better matches individual preferences and needs. However, it is worth noting that higher levels of personalization require additional effort for the model to reach optimal performance. Future studies could explore the generalizability of personalization approaches in different application domains. Additionally, conducting user studies in real-world settings is necessary to gain deeper insights into practical implementation challenges.

Keywords: *Socially Assistive Robotics; Personalization; User Profiling; Rule-based systems; Reinforcement Learning.*

1 Introduction

Over the past decade, the robotics industry has experienced rapid and remarkable expansion [1]. Socially Assistive Robotics (SAR) [2] is a field of research and development focused on designing robots that assist and support individuals through social interaction rather than physical interaction. These robots employ non-invasive interaction methods to provide emotional, cognitive, and social engagement to users, aiming to enhance their well-being and quality of life [3]. From healthcare to education, and from customer service to socialization, these advancements in technology have the potential to address societal challenges while providing innovative solutions [4].

In the healthcare sector, one pressing challenge is the shortage of healthcare professionals. According to the World Health Organization [5], the global population aged 60 and above is projected to increase from 12% to 22% between 2015 and 2050 [6]. This demographic shift will place a substantial burden on our healthcare and social systems, necessitating alternative solutions to meet the growing demand for healthcare services [5]. SAR can assist in healthcare delivery by providing personalized care, monitoring vital signs, reminding patients to take medications, and supporting rehabilitation exercises [7]. These robots can also engage in social interactions with patients, offering emotional support and companionship, which can be particularly valuable for individuals facing isolation or cognitive impairments [8]. Furthermore, SAR can facilitate telemedicine and remote healthcare monitoring, enabling healthcare professionals to remotely assess and provide guidance to patients, especially in underserved or remote areas [9]. By augmenting the capabilities of healthcare professionals

and enhancing patient experiences, SAR has the ability to improve healthcare outcomes and transform the way healthcare services are delivered. [7].

SAR is increasingly recognized as a valuable asset in the field of education [10]. Robots can serve as interactive learning companions, offering personalized instruction and support to students [11]. By adapting their teaching strategies to individual learning styles, providing real-time feedback, and engaging students in interactive discussions, SAR is able to create a dynamic and engaging learning environment that promotes knowledge acquisition and retention [12]. Additionally, SAR addresses the challenges of limited resources and growing student-to-teacher ratios by providing additional support and guidance, thus optimizing the learning process and fostering educational inclusivity [13].

In the field of customer service, SAR has emerged as a valuable solution as well, revolutionizing how businesses interact with their clientele [14]. SAR is able to enhance customer experiences by offering personalized recommendations, assisting in product or service location, and providing prompt responses to inquiries [15]. This advanced technology enables businesses to deliver exceptional customer service, fostering higher satisfaction and loyalty [16]. Moreover, SAR helps organizations meet the increasing demand for personalized assistance while optimizing operational efficiency [14]. By incorporating SAR, businesses can elevate customer engagement, deliver top-notch service, and gain a competitive edge in the market [15].

While the potential of socially assistive robots extends to various domains such as healthcare, education, and customer service, addressing certain limitations is crucial to fully leverage their benefits in meeting user demands [17]. One critical limitation lies in the challenge of achieving effective personalization, particularly for vulnerable individuals or those with cognitive impairments [18]. For instance, in education, socially assistive robots can offer tailored reminders to students for timely completion of assignments, considering their preferred learning styles and workload. Similarly, in healthcare these robots can be programmed to understand a user's dietary restrictions or preferences, providing meal suggestions that align with their specific needs and enhancing their overall experience. Overcoming these limitations is vital to ensure that socially assistive robots cater to the unique requirements of users across diverse domains, ultimately enhancing outcomes and quality of life.

Personalization is an essential aspect of socially assistive robots, ensuring that individual users' needs and preferences are considered during the design and implementation of the robot's features and functions [19]. This approach enhances user engagement, satisfaction, and adherence to various programs in different domains [20]. Overcoming the limitations associated with personalization is crucial to enable these robots to provide necessary support and assistance to a wide range of users, ultimately leading to improved outcomes and quality of life [21]. Therefore, the objective of this project is to identify the various forms of personalization that can be implemented and provide insights into the optimal utilization of each form.

2 Related Work

SAR, or socially assistive robotics, is an emerging field that has garnered significant attention due to its potential to provide innovative solutions to meet the increasing demands across various domains, including healthcare. The goal of this research field is to provide personalized support and assistance through non-invasive interaction, aiming to enhance users' well-being and quality of life. Since the 1990s, many researchers have investigated the use of SAR in numerous applications, including physical therapy, emotional support, and cognitive rehabilitation. In this section, we examine various pertinent research projects and SAR-related technological advancements.

2.1 Perceptions and Attitudes towards SAR

Positive perceptions and acceptance of socially assistive robotics (SAR) among patients and clinicians have been observed in several studies, indicating the potential benefits of SAR in various healthcare applications [22]. Raigoso et al. conducted a survey focusing on the use of social robots in physical rehabilitation (PR) treatments, and the results demonstrated a favorable perception of SAR among participants. Although some negative impressions were noted, indicating complexity in patients' views on social robot usage, the findings highlighted the importance of clarifying the robot's goals and capabilities from the beginning to enhance acceptance. The study also suggested that incorporating speech recognition technology could improve natural interaction with the robot and boost user engagement.

Furthermore, exploring older individuals' attitudes towards SARs in healthcare, Pino et al. found that most participants had positive sentiments regarding the use of SARs and recognized their potential in supporting daily tasks [23]. However, concerns were raised regarding the need for individualized care and the potential loss of human interaction. The study emphasized the importance of designing SARs with consideration for the requirement of human interaction and personalized treatment. Addressing user concerns, such as the suitability of the robot for individual needs, ease of operation, and familiarity with technology, are crucial factors for successful SAR adoption in healthcare settings. Understanding the perspectives of older individuals regarding SARs can provide valuable insights for the development and implementation of SARs in healthcare.

Studies investigating perceptions and attitudes towards SAR highlight the overall positive reception among patients, clinicians, and older individuals. These findings emphasize the potential of SAR to enhance healthcare delivery and rehabilitation while emphasizing the need for clear communication of goals, consideration for individualized care, and addressing user concerns to ensure successful integration and acceptance of SAR in various healthcare settings.

2.2 User Profiling for Personalized and Adaptable SAR

Personalization plays a crucial role in the development of socially assistive robotics (SAR), enabling adaptive interaction and tailored assistance to users based on their specific needs and contexts [24]–[27]. By leveraging user profiling techniques, SAR systems can customize their behaviors and provide effective support in various domains.

One approach proposed by Benedictis et al. [24] focused on adaptive interaction, aiming to personalize assistance based on users' health-related conditions and adapt behaviors to their changing needs. By incorporating a user model grounded on the International Classification of Functioning, Disability, and Health (ICF) framework, the authors designed a control architecture inspired by the dual-process theory. The study showcased the application of this approach in the context of cognitive stimulation for older adults, where the social robot customizes verbal behavior according to user characteristics and dynamic reactions during interaction. This general approach demonstrated the potential of personalized and adaptable SAR in various scenarios.

Cesta et al. [25] proposed a cognitive loop architecture that connected reasoning on sensed data with the process of acting, enabling SAR to engage in high-level planning and low-level reactive behaviors. The user profile, aligned with the ICF framework, was integrated into the reasoning part of the cognitive loop, facilitating the generation of customized plans and goals that cater to specific user requirements and preferences. The study emphasized the importance of considering human interaction and individualized treatment in the design and deployment of SARs in healthcare settings. By understanding user needs and preferences, SAR has been able to overcome challenges, such as restricted mobility, and generate appropriate plans to improve daily tasks and overall well-being.

Umbrico [26] adopted a similar approach, utilizing an architecture inspired by the theory of dual processes to create intelligent and adaptive behavior in social robots. The proposed architecture incorporated modules for long-term assistive planning and dynamic plan execution, adapting to contingencies during interaction. Although exemplified in cognitive rehabilitation, the architecture's flexibility allowed its application in various domains, including human-robot collaboration.

The UPA4SAR project [27] presented an example of user modeling that incorporates both static and dynamic profiling. The project focused on developing a personalized robotic application for home care of elderly patients with cognitive impairments. By utilizing a service-oriented approach and the ICF model, the application enabled personalization and adaptation based on the patient's needs, preferences, and available technology. The user model encompassed static profiles, such as personality traits, and dynamic profiles, including daily observations like current mood. The autonomous robotic application performed daily assistive tasks while considering individualized care and support.

User profiling techniques in SAR research have offered opportunities for personalization, adaptation, and improved user experiences. These approaches, grounded in the ICF framework, have contributed to the advancement of SAR by addressing individual needs, enhancing engagement, and tailoring assistance across various domains

2.3 Addressing Needs and Challenges through End-User-Centered Design

End-user-centered design is crucial for the successful adoption of socially assistive robots (SAR), especially in healthcare settings. Cooper et al. [28] developed ARI, a SAR designed by PAL Robotics specifically to address the needs of older people, individuals with physical constraints, and those in isolation due to infectious diseases. However, barriers such as robot

acceptability and constraints in speech interaction and empathy still hinder the widespread adoption of SAR robots in healthcare. The paper has highlighted the importance of end-user-centered robot design to establish SAR as true partners in healthcare.

In the context of providing companionship and independence to people with dementia, Khosla et al. [29] emphasized the two main abilities required for social robots to achieve a person-centered approach. Firstly, social robots should perform functions that are perceived as useful and easy by older people with dementia. Secondly, they should engage older people in a unique experience of human-robot interaction. The study emphasized the preference for personalized robot services over standardized ones and highlighted ease of use as a crucial factor when designing intervention services for home-based use. Designing for emotion engagement, visual engagement, and behavioral engagement, as well as creating a conducive interactional environment and effective communication modalities, played significant roles in facilitating a long-term and meaningful reciprocal relationship between social robots and their human partners.

End-user-centered design is essential for addressing the needs and challenges of socially assistive robots in healthcare. Personalization, ease of use, and creating engaging human-robot interactions are crucial for establishing SAR as valuable partners in healthcare settings.

2.4 Application of Machine Learning Techniques in SAR

Machine learning techniques have gained significant interest in the field of socially assistive robotics (SAR) due to their potential for personalization and adaptation. Several studies have explored the application of machine learning algorithms in SAR, demonstrating their effectiveness in enhancing the interaction and learning capabilities of robotic systems.

One area of application is personalizing the interaction between robot companions and older adults with cognitive impairments. Moro et al. [18] developed a novel robot learning architecture that combined learning from demonstration (LfD) and reinforcement learning (RL) algorithms. The experiments conducted with the socially assistive robot Casper showcased the ability to teach personalized behaviors for activities of daily living, such as tea-making. The combination of LfD and RL algorithms significantly reduced the learning time required for the robot to acquire new activities.

Integrating robotics with artificial intelligence (AI) is another avenue for leveraging machine learning techniques in SAR. Umbrico et al. [30] proposed a cognitive architecture that integrates knowledge reasoning, planning, and acting using the KOaLa ontology. By incorporating the International Classification of Functioning, Disability, and Health (ICF) framework, the system allowed for adaptability and personalization in providing assistive services based on the individual's profile. The feasibility of the proposed system was demonstrated in a simulated environment, with future work focusing on customizing the system for specific human ecology goals and incorporating learning capabilities to capture changes in the individual's daily living.

In the domain of education, machine learning techniques have been employed to create personalized learning experiences. Park et al. [31] developed a social robot learning companion system for early literacy education. The system utilized machine learning algorithms to train a personalized policy for each student based on their verbal and nonverbal affective cues. The personalized policy optimized engagement and linguistic skill progression

by selecting stories tailored to each child’s needs. The deployment of this system in schools showed improved engagement and learning outcomes compared to other groups, highlighting the efficacy of personalized policies in maximizing learning gains.

In summary, the application of machine learning techniques in SAR enables personalization and adaptation in various contexts. These techniques enhance the interaction capabilities of robot companions, facilitate the provision of individualized assistive services, and optimize learning experiences. By leveraging machine learning algorithms, SAR systems can better meet the specific needs and preferences of users, ultimately improving their overall well-being and quality of life.

In conclusion, while there is ample research on personalizing socially assistive robots, there is currently no standard approach for doing so. This lack of standardization makes it challenging to compare results across studies, as different research groups employ varying methods and metrics for evaluating personalization’s effectiveness [17]. Additionally, researchers must first understand the unique preferences and needs of each user, which is an area where there is still limited research on how to gather and use this information effectively [17]. Overall, the related work highlights the crucial role that personalization plays in the acceptance and effectiveness of SAR. Although some research has explored these areas, there is still a need for further investigation on how to effectively implement personalization and establish and maintain trust between the robot and its user. This study aims to fill these gaps in the literature.

3 Theoretical Framework

In the preceding chapter, the significance of personalization in socially assistive robotics (SAR) and its impact on the acceptance and effectiveness of SAR systems were explored. The need for a standardized approach to personalization and the challenges associated with understanding individual user preferences and needs were discussed. Building upon this groundwork, the focus of this chapter is to establish a theoretical framework that precisely defines various personalization parameters along with their associated challenges. This framework will serve as the bedrock for making informed decisions when selecting among the different personalization mechanisms, which will be elaborated on in the forthcoming chapter.

Personalization within the context of a social companion pertains to tailoring the agent’s behavior, interaction and recommendations to the unique preferences, needs, and context of each individual user [32]. The level of personalization can vary based on various factors, such as the situation the user is currently in. For instance, if the user is engaged in a conversation, the social companion might adjust its responses and attentiveness accordingly. Similarly, the level of personalization can be influenced by the user’s needs. For example, if the user requires assistance with medication reminders, the companion can provide tailored reminders and prompts. Furthermore, personalization can also take into account the user’s preferences. For instance, if the user prefers a particular tone of voice or communication style, the companion can adapt its interactions to align with those preferences.

In certain use cases, users may have limitations or reservations in providing extensive

data and feedback to the robot, which calls for lower levels of personalization. For example, individuals with limited technological skills or those facing a health crisis may lack the ability or inclination to spend time configuring and customizing the robot. In such cases, it becomes crucial to design a social companion with a simple and straightforward interface that requires minimal setup and maintenance.

Conversely, there are situations where personalization plays a vital role in ensuring a positive user experience. For example, consider individuals with diverse cultural backgrounds or communication styles. In order to feel understood and supported, they may require personalized recommendations and interactions that align with their cultural values and communication preferences. Additionally, personalization can be crucial in tailoring fitness or health advice to meet the unique needs and preferences of each user. For instance, some users might benefit from personalized exercise routines based on their specific fitness goals or physical limitations, while others might prefer dietary recommendations tailored to their nutritional preferences or restrictions.

The required level of personalization is contingent upon various factors, including the user’s preferences, needs, technological skills, and the specific context of the healthcare companion’s use [33]. Users with specialized demands or cognitive impairments may require a higher level of personalization, while those familiar with the technology and possessing fewer specialized needs may require a lower level of personalization.

In this chapter, we established a theoretical framework in order to effectively capture the diverse factors involved in the personalization process.

3.1 The Framework

In order to effectively capture the diverse factors involved in personalization, a comprehensive framework can be established. This framework is depicted in figure 1, where the user profile serves as a repository of information that encompasses the various personalization parameters [34], enabling a holistic understanding of the individual user in the context of a social healthcare companion. These parameters can be categorized into four main dimensions: context, user characteristics, user preferences, and user state. The context dimension incorporates situational factors that influence personalization, such as the user’s current environment, social setting, and activity level. User characteristics can encompass a wide range of factors, including personality traits, cultural background, age, and cognitive abilities, which shape the individual’s unique needs and responses. User preferences capture the individual’s specific likes, dislikes, communication styles, and preferences for interaction. The last dimension, user state, refers to the current condition or state of the user at a given moment. It encompasses aspects such as the user’s emotional state, physical health, level of fatigue, and overall well-being. By systematically organizing these personalization parameters within a user profile, a more nuanced and tailored approach can be adopted to enhance the social companion’s adaptability, responsiveness, and overall user experience [35].

In order to understand how the user profile interacts with the personalization of the social healthcare companion, Figure 1 was constructed. This figure illustrates the dynamic interaction between the companion and the user, highlighting two key aspects that can be personalized in this context: the treatment and the social interaction. The figure also captures how the user profile is leveraged to personalize the different modules of the companion.

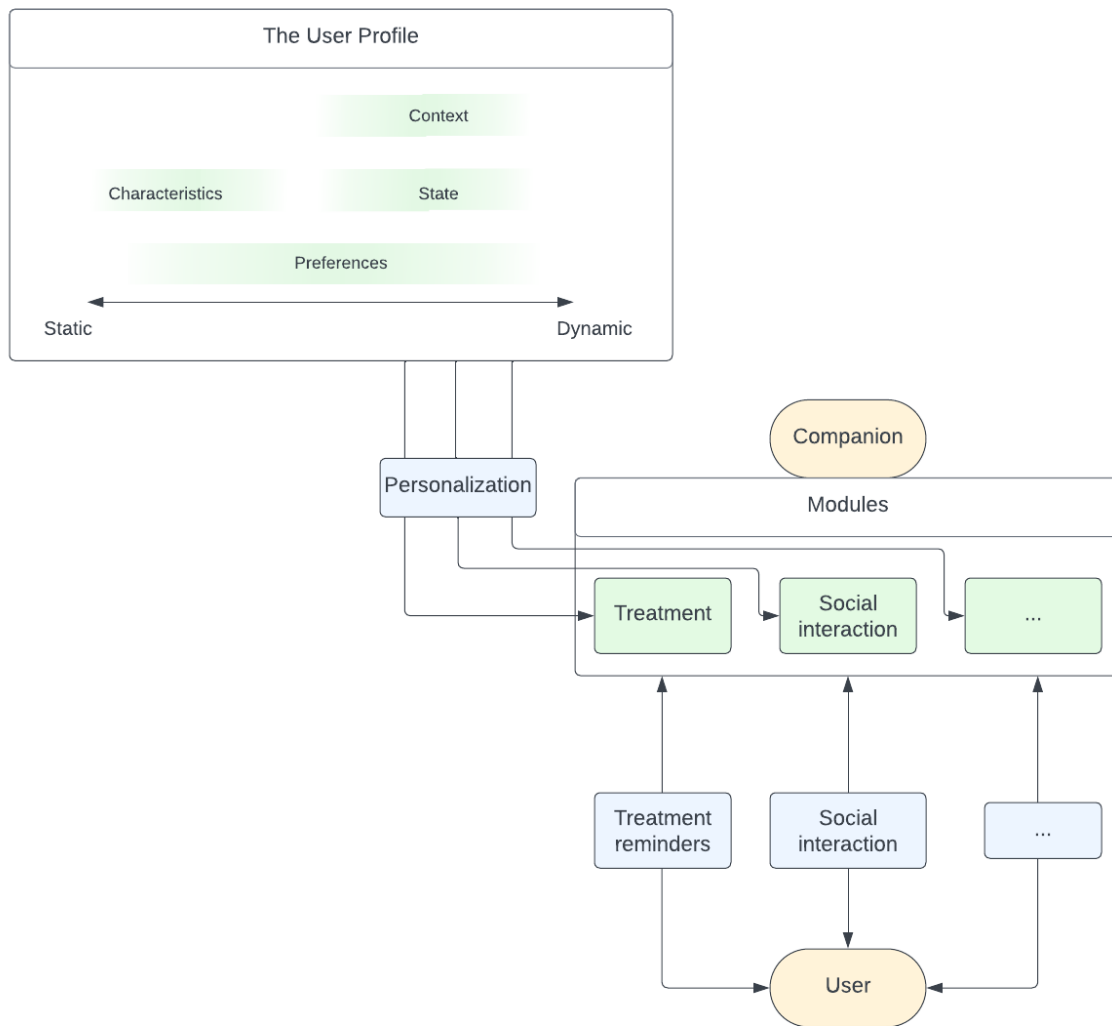


Figure 1: The Personalization Framework for the Social Healthcare Companion - The figure illustrates the interaction between the user profile and the personalization of the social healthcare companion. The framework highlights two modules of the companion that can be personalized: the treatment and the social interaction, the third module is placed in the figure to emphasize that there are more modules possible. Within the user profile, a comprehensive collection of personalization parameters is organized into categories, including context, user characteristics, user preferences, and user state. These categories are strategically positioned along an axis that represents their degree of dynamism.

The treatment module of personalization involves tailoring the healthcare interventions and recommendations provided by the companion to align with the user's specific healthcare goals, treatment protocols, and individualized needs. The user profile, with its categorization of personalization parameters, plays a pivotal role in this process. By considering factors such as the user's medical history, dietary restrictions, physical capabilities, and preferences, the companion can customize its treatment and reminders to ensure they are relevant, effective, and aligned with the user's unique requirements.

The social interaction aspect of personalization focuses on adapting the companion’s behavior, communication style, and interactions to match the user’s preferences, characteristics, and context. The user profile provides valuable insights into the user’s communication preferences, cultural background, personality traits, and situational factors that influence personalization. By leveraging this information, the companion can adjust its tone of voice, language style, and social cues to create a more comfortable and engaging social interaction with the user.

The third module in the figure acts as a placeholder to emphasize that treatment and social interaction are not the only modules available for personalization in the context of a social healthcare companions, and there are other possible options, depending on the user’s goals.

3.2 Challenges

The process of collecting data for the user profile comes with its own set of challenges, particularly when considering the nature of the data involved. An important consideration regarding the user profile data is its classification as either static or dynamic. Static data encompasses information that remains relatively stable over time, such as the user’s cultural background, personality traits, and medical history. On the other hand, dynamic data captures the user’s ever-changing context and preferences. The dynamic profile contains parameters that are updated more frequently during use, such as the user’s current mood or health status, which are represented by the context and preferences dimensions in Figure 1. By incorporating both static and dynamic parameters within the user profile, a more comprehensive understanding of the user can be achieved.

Contextual data are dynamic because they encompass information that can change over time and in different situations. Factors such as time, location, environment, and presence of the user contribute to the dynamic nature of contextual data.

User characteristics, such as personality traits, cultural background, age, and cognitive abilities, remain relatively stable over time and help us understand the user’s fundamental attributes and needs. User preferences lie between the static nature of user characteristics and the dynamic nature of user state. They can evolve or change over time but generally remain more consistent than the user’s immediate state. Preferences include communication style, likes, dislikes, and personal choices.

The user state dimension captures the dynamic aspects of the user’s condition, such as their emotional state, physical health, level of fatigue, and overall well-being. This dynamic aspect adds complexity to personalization, as the social healthcare companion needs to adapt its treatment and social interaction strategies in real-time based on the user’s changing state, providing effective support and assistance.

One specific approach to user profiling in socially assistive robotics that aligns with the person-centered approach is the International Classification of Functioning, Disability and Health (ICF) model [36]. This model facilitates the tailoring of the user profile to individual needs and preferences by incorporating a static and a dynamic user profile, which was introduced in figure 1. To illustrate the components of the user profile and their role in behavior personalization and adaptation, a graphical representation is provided in Figure 2. This representation demonstrates how the user profile serves as a foundation for understanding

the individual user, enabling the social healthcare companion to adapt its behavior, communication style, and interventions to align with the user’s preferences, characteristics, and context.

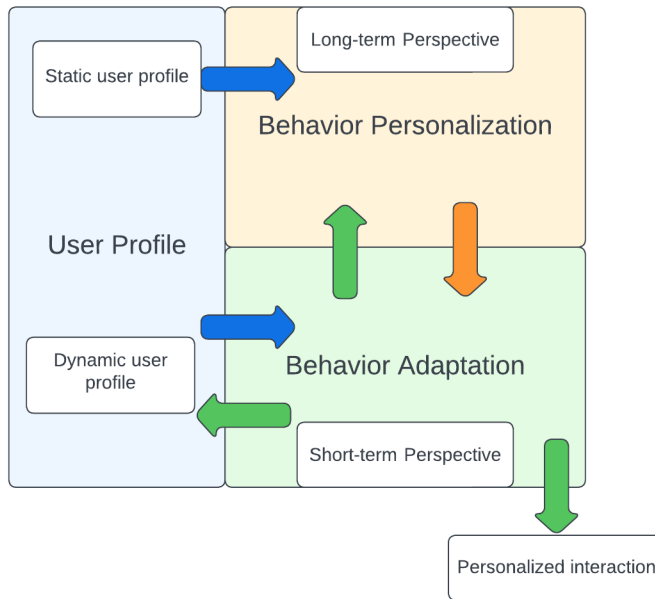


Figure 2: User Profile Based on the ICF Model [36] - A graphical depiction showcasing the user profile in socially assistive robotics, incorporating static and dynamic components that interact synergistically to enable personalized interactions and tailor the agent’s support and recommendations to individual user needs and circumstances.

4 Personalization Mechanisms

In the previous chapter, we presented a comprehensive theoretical framework for personalization in the context of a social healthcare companion. The framework established a user profile as a repository of information, encompassing various personalization parameters that enable a holistic understanding of the individual user. It highlighted the dynamic interaction between the user profile and the personalization modules of the companion. Building upon this foundation, the current chapter introduces the different mechanisms for personalization that will be considered in this research. Specifically, we will explore pre-programmed rules, learning while using, and a hybrid approach. By examining and comparing these mechanisms, we can gain insights into their strengths, limitations, and potential for delivering effective personalized support in social healthcare robotics.

Learning while using and pre-programmed rules are two approaches for personalizing the companion based on the information that is in the user profile. These specific two are chosen, because a rule-based system is the most fundamental approach; the pre-programmed rules provide guidelines to personalize the interactions and ensure consistency [37]. Learning while using, also referred to as reinforcement learning, allows the agent to improve its interactions

with the user by collecting data on their behaviors and adapting accordingly, without the need of pre-training on large data sets, as other machine learning techniques would [37]. Afterwards, a hybrid approach is suggested, which aims to combine pre-programmed rules and reinforcement learning in order to overcome a lot of the inefficiencies.

4.1 Pre-programmed rules

The first mechanism that will be explored is a system using pre-programmed rules. Pre-programmed rules are guidelines or instructions that are created in advance to dictate how an agent should interact with users in different scenarios [38]. These rules provide a static basis for personalization; it does not allow for adaptability, but it does allow for the companion to interact with the user profile. These rules can be built on the information in the user profile, described in chapter 3. The rules are often based on static information from the user profile, such as user characteristics and preferences, but in some cases rules can also be based on more dynamic parameters, such as state and context (figures 1,2). Arguably, this rule-based approach is not a direct mechanism for personalization, but it does allow the system to act using the information of the user profile. The rules are pre-defined and do not change over time.

One potential advantage of pre-programmed rules is that they can be designed to respond to users' emotional states in a more sensitive and appropriate way. For instance, if a user is feeling sad, the agent can be programmed to provide emotional support and comfort through predetermined responses. Additionally, pre-programmed rules can help ensure that the agent consistently adheres to certain ethical and moral principles in its interactions with the user [38].

However, there are also potential disadvantages of pre-programmed rules. The biggest concern in terms of personalization is that while the guidelines may be based on research or expert opinions, they might not always align with the individual user's personal preferences or unique situation. The pre-programmed rules can overlook the nuances of individual differences and may not fully account for the complexity of the user's needs. Therefore, relying solely on pre-programmed rules without considering the user's specific situation or preferences can limit the effectiveness of the social companion in providing personalized support. It is essential to continuously monitor and adjust the pre-programmed rules based on the user's feedback and engagement to ensure that the system is adaptable and responsive to the individual user's needs. Therefore, as mentioned in the previous section, a combination of reinforcement learning and pre-programmed rules is often optimal [39].

Overall, pre-programmed rules can be a valuable tool for personalizing the interactions between a social companion and users. However, they should be designed with care to balance the benefits of personalization with the potential drawbacks of rigidity and artificiality.

4.2 Learning while using

A second approach to personalizing a social companion is through learning while in use, also referred to as reinforcement learning. Reinforcement learning (RL) involves the agent gathering feedback data on the user's interactions and behaviors to improve its future interactions with the user. Reinforcement learning is a sub-field of machine learning that focuses

on training agents to make decisions in a dynamic environment by using feedback signals in the form of rewards or penalties [40].

Reinforcement learning involves an agent interacting with an environment by taking actions based on its current state. The agent receives feedback in the form of rewards from the environment, which depend on the actions it chooses. Figure 3 provides a schematic depiction of the reinforcement learning cycle. For instance, a social companion can adjust the frequency and content of medicine reminders based on the user’s responses. The agent can adapt the timing and language of the reminders to better match the user’s preferences and needs if the user frequently ignores or becomes irritated by them.

Through repeated interactions with the user, the agent can learn to make improved decisions that maximize the user’s satisfaction and well-being. This process is more suited for dynamic information from the user profile, depicted in figure 1, as it can learn about user preferences through feedback, while taking dynamic user states and context into account. By continuously learning from the user’s behavior and feedback, the companion can become an even more effective tool for supporting the user’s health and well-being.




Figure 3: The reinforcement learning cycle - A schematic representation illustrating the iterative process of reinforcement learning, showcasing the interaction between the agent, environment, and rewards to optimize the agent’s decision-making and behavior.

However, in some circumstances, using RL for healthcare companions can also be ineffective due to a variety of factors, such as the requirement for ongoing user feedback to gather sufficient data for interaction personalization. The user may become agitated by having to give the robot feedback or input all the time, especially if they are already coping with a health issue or other stressors. This requirement for feedback might be troublesome.

It may take a while before the agent can effectively tailor its interactions with the user because obtaining sufficient data through RL can be a time-consuming process [41]. The effectiveness of the intervention may suffer if the user gets frustrated or loses interest in the agent during this time.

Another potential drawback of RL for social companions is that it might not be appropriate for all kind of healthcare interventions. For instance, some interventions might work better using an expert system or rule-based approach rather than a machine learning method.

In order to personalize healthcare companions, it might be required to combine techniques like rule-based systems and machine learning to overcome these inefficiencies [39]. Additionally, it might be necessary to carefully strike a balance between the demand for

data and the desire for a good user experience, as well as to create interventions that are both successful and user-acceptable, as mentioned in section 3.2.

4.3 Hybrid Approach

A hybrid approach that combines pre-programmed rules and reinforcement learning offers a comprehensive solution for personalizing social companions. Pre-programmed rules provide a foundational framework based on expert knowledge and the user profile, ensuring ethical adherence and addressing user preferences. Reinforcement learning allows the companion to learn and adapt based on real-time user feedback, enhancing personalization over time.

This hybrid approach provides flexibility, adaptability, and efficient learning. It balances personalization with user acceptance by starting with familiar experiences guided by pre-programmed rules and gradually adapting based on specific user behaviors and preferences. Ongoing refinement and evaluation are necessary to optimize the balance between rule-based guidelines and the learning component.

By combining pre-programmed rules and reinforcement learning, the hybrid approach enables personalized interactions and tailored assistance while improving adaptability. The hybrid approach is perfectly suited for using both the dynamic and static user profile parameters, as it can use the rules to leverage the static user characteristics as a starting point, and update during use by learning how to interact based on more dynamic parameters. Figure 4 presents a diagram depicting the hybrid approach employed by social companions in tailoring their responses to user needs and preferences.

4.4 Evaluation of Personalization Approaches

In evaluating personalization methods in socially assistive robots we will need to take into account the trade-off between the cost to put each technique into practice and the benefits that are obtained. We will do this by looking at the effort required from the user and the degree of personalization that can be reached.

In reinforcement learning, the agent interacts with humans and learns from their feedback to improve its behavior and adapt to different situations. The required effort involves the load on the user, the effort required to provide feedback or guidance to the robot during the learning process [42]. Ideally, the goal is to minimize the required effort for the user while achieving effective and efficient learning, which we already touched upon in section 3.2.

The degree of personalization refers to how well a system or service is customized to individual users. It involves tailoring the experience to suit each user’s specific needs [43]. Evaluating personalization involves assessing the effectiveness of delivering personalized recommendations by utilizing user-specific information.

User profiling serves as a foundational element in all approaches, albeit in varying forms. Pre-programmed rules often employ the static user profile, while learning while using mostly relies on the dynamic user profile. The hybrid approach integrates both the static and dynamic user profiles. In this section, these three approaches will be compared. A summary of the comparisons is displayed in table 1.

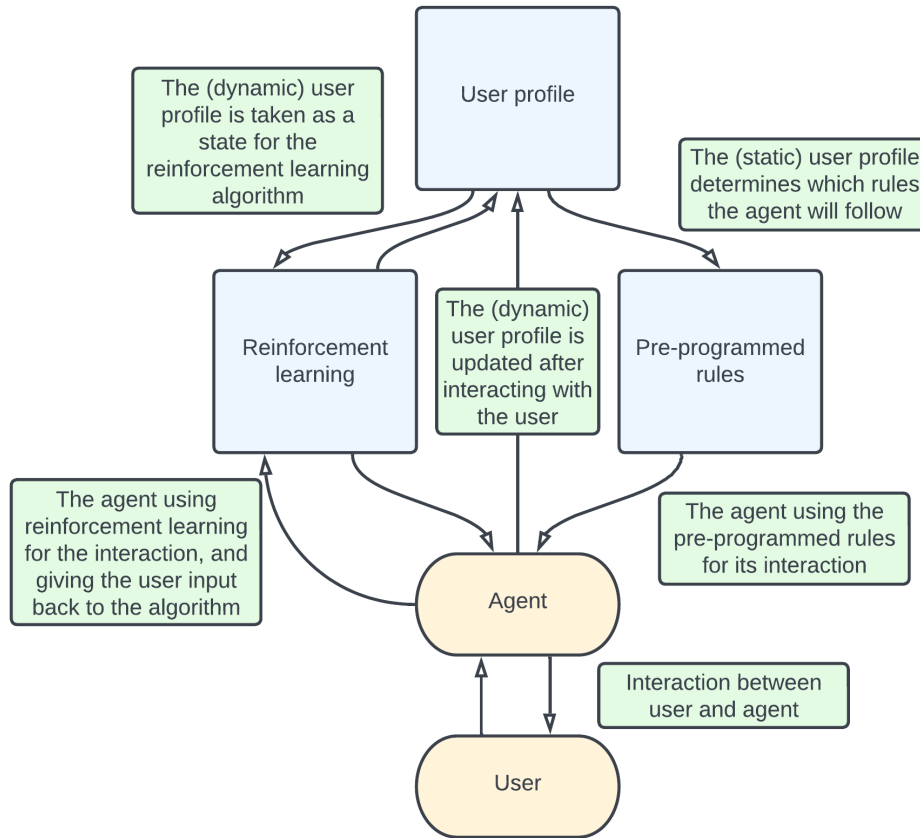


Figure 4: Hybrid Approach for Personalization in Social Companions

Approach	Effort required	Personalization level
Pre-programmed rules	Small amount of effort is required, only at initialization	Minimal level of personalization; rules do not change over time.
Learning while using	Large amount of effort is required, extensive user feedback required over long period of time	High level personalization is possible; model learns user's preferences over time
Hybrid approach	Large amount of effort due to RL, but combining RL and pre-programmed rules gives the model an initial advantage	High level of personalization is possible because of the flexibility of the model

Table 1: Comparison of Personalization Approaches from a Technical Perspective: Effort Required and Personalization Levels

In the pre-programmed rules approach, only a small amount of effort is necessary, primarily during initialization. This approach operates based on predefined rules that do not change over time. As a result, the level of personalization achieved with this approach is minimal since the rules remain static and cannot adapt to individual user preferences.

Contrastingly, the learning while using approach demands a significant amount of effort from the user. To obtain the desired output, users must provide extensive feedback to the system over an extended period of time. This continuous feedback allows the model to learn and adapt to the user's preferences gradually. As a result, the personalization level achievable with this approach is high, as the model dynamically evolves and incorporates user feedback to improve its performance.

The hybrid approach combines reinforcement learning (RL) with pre-programmed rules, resulting in a more complex and versatile system. This approach requires a substantial amount of effort due to the involvement of RL. However, by leveraging the combination of RL and pre-programmed rules, the model gains an initial advantage. It starts with pre-existing rules that provide a head-start in the learning process. As the system interacts with users and employs RL, it can further personalize its responses and adapt to individual preferences. The hybrid approach potentially offers a high level of personalization, benefiting from the flexibility of the model to adjust and improve over time.

In summary, while the pre-programmed rules approach requires minimal effort from the user and can suffice in many applications, the learning while using approach demands significant user effort but often results in a higher level of personalization through continuous learning. The hybrid approach strikes a balance by utilizing RL and pre-programmed rules, requiring a substantial effort but offering a high level of personalization due to its adaptability and flexibility. The choice of approach depends on the desired trade-off between user effort and the degree of personalization sought, which is in turn dependent on the application of the companion, of which we will try to gain a better understanding in the following section.

5 Prototypes



Figure 5: Lizz, the digital healthcare assistant developed by ConnectedCare [44].

This project focused on a specific use case: a digital healthcare companion named Lizz (figure 5) will be used as a carrier for this research. Lizz is an embodied digital healthcare assistant, developed by ConnectedCare [44], with the aim to provide remote patient support, including rehabilitation, exercise guidance, attention and structure support, and dietary advice. Given its current implementation in 4 clinics as a supportive tool during rehabilitation, Lizz serves as an ideal platform for our research on personalization.

While Lizz can be personalized through a web dashboard in its current version, it currently lacks dynamic personalization mechanisms. To address this limitation, our project aimed to implement various means of personalization in Lizz by leveraging user profiling and a combination of pre-programmed rules and learning while in use through reinforcement.

5.1 Prototype Scenario

To explore the advantages and disadvantages of different personalization mechanisms within Lizz, we will start by describing a generic scenario and the underlying challenges and goals. This scenario served as a framework for exploring various ways to implement personalization, as outlined in Chapter 4. By taking this approach, we demonstrated the diverse possibilities of personalization and determined the most suitable form for each specific context.

In this scenario, Lizz was used for exercise suggestion during rehabilitation. The scenario revolved around a predefined set of exercises: short walks, long walks, yoga, cycling, stretching, and swimming. This selection prioritized diversity and manageability, accommodating different user preferences and needs. It is important to note that future iterations of the algorithm may require adjustments to this exercise set based on treatment provider recommendations.

In addition, Lizz incorporated specific timings for exercise recommendations: 8 AM, 12 PM, 4 PM, and 8 PM. The inclusion of these time slots was crucial as it allowed Lizz to seamlessly integrate with users' daily routines. By suggesting exercises during these specific intervals, Lizz facilitated the integration of physical activity into users' schedules. This temporal personalization promoted the development of consistent exercise habits and maximizes the effectiveness of rehabilitation efforts.

In summary, the careful selection of the exercise set and the incorporation of specific exercise timings within Lizz contributed to delivering a personalized and effective rehabilitation experience. Rather than presenting multiple scenarios, we used a single scenario as a basis and explored different approaches to its implementation. This approach ensured a stronger narrative and aligns with our commitment to diversity, accommodating user preferences, and seamless integration into users' daily routines.

The initial user profile

In all prototypes, a static user profile as described in section 3 was created by asking the user a specific set of questions when they started using Lizz. This profile comprised information about the user that generally remains unchanged over a short period of time, making it more static in nature. The selection of these questions was carefully made to gather information that is relevant for exercise recommendation.

For the implementation of these prototypes, we chose to have Lizz ask the questions and allow the user to answer through the interaction with Lizz itself. This choice was influenced by the current limitations, where establishing a direct connection between the dashboard and WebSockets is not feasible. However, it is worth noting that for future development of Lizz, establishing this connection would be beneficial, as it would enable healthcare providers to also contribute to filling the static user profile from the web dashboard.

The specific set of questions chosen for this project was designed to capture information that is valuable for exercise recommendation. The questions used in this example were as follows:

- What is your age? (options: younger than 40, between 40 and 60, older than 60)
- How would you rate your activity level? (options: lightly active, moderately active, highly active)
- At what time during the day are you most active? (options: 8 AM, 12 PM, 4 PM, 8 PM)



Figure 6: Images where Lizz asks the user the initial questions for setting up the (static) user profile.

Figure 6 shows the interface of how these questions were asked using Lizz. By gathering the user’s responses to these questions, we populated the static user profile. This profile served as the foundation for the rule-based part of the algorithm in all three scenarios. The information gathered from the user’s answers helped to tailor the exercise recommendations provided by Lizz, ensuring they are better aligned with the user’s age, activity level, and preferred time of activity. Overall, the selection of these specific questions was driven by the intention to construct a static user profile that contains information directly relevant to exercise recommendation.

Prototype 1: Rule-based system without user feedback

In this prototype, Lizz used a rule-based system to recommend exercises to the user based on their initial, static user profile. The system did not require any user feedback beyond initial setup. The user was asked to provide basic information as described in the previous section. The system relied on pre-programmed rules to make recommendations and did not incorporate any machine learning. The level of personalization in this prototype was thus very low.

Prototype 2: Combination of rule-based and reinforcement learning with user preferences

This prototype was an extension of prototype 1, where Lizz used reinforcement learning to improve the exercise recommendation engine over time. To burden the user as little as possible, Lizz provided the user with two different options of recommended exercise at different times during the day. The user could choose which exercise they preferred from the two options presented, and based on their selection, the algorithm learned from the user’s behavior and improved the recommendations over time. The user could also choose to not perform any of the two recommended exercises, which also helped the system learn which recommendations are not appropriate at certain times. While the system still took into account the user’s profile, the incorporation of machine learning through reinforcement

learning helped tailoring the recommendations to the individual user’s needs and preferences. For example, if the user consistently chose certain exercises at a certain time of day, Lizz prioritized recommending those exercises during that time. There is a moderate level of personalization in this prototype, as the recommendations were based on information learned from the user, but there was no state monitoring involved.

Prototype 3: Combination of rule-based and reinforcement learning with user preferences and user state

This prototype built upon prototype 2, where Lizz used a combination of rule-based and reinforcement learning to suggest exercises to the user. However, in this prototype, Lizz considered a wider range of factors to enhance the exercise recommendations. The system kept track of the user’s dynamic profile through state monitoring. Furthermore, reinforcement learning techniques were employed to refine the recommendation algorithms, taking into account the user’s feedback. The exercise recommendations took into account various factors, which were the user’s emotional state, fatigue level, initial information from the static user profiling, the different time points during the day, and the exercises previously selected by the user. For example, if the user reported feeling particularly tired during a certain time of day, Lizz avoided recommending intensive exercises during that period. The level of personalization in this prototype is a lot higher, which might improve the user experience in terms of personalization, but providing Lizz with information on their current state, might also exhaust the user.

5.2 Implementation

The implementation of the prototypes was done in JavaScript, a programming language primarily used for creating interactive and dynamic web pages. It runs on the user’s web browser and enables developers to add functionality, respond to user interactions, and communicate with servers. JavaScript was chosen because this allows us to communicate with Lizz using WebSockets. WebSockets are a communication protocol that provides a persistent and bidirectional connection between a web browser and a server. Unlike traditional HTTP requests, which are stateless and require the client to initiate a new connection for each request, WebSockets allow for real-time, two-way communication between the client and server. This enables the server to push data to the client instantly, facilitating dynamic and interactive web applications such as chat applications, live updates, and multiplayer games.

For the implementation of the three prototypes, as mentioned in section 5.1, a set of exercises was determined (short walk, long walk, yoga, cycling, stretching, swimming). We also used a set of times during the day (8 AM, 12 PM, 4 PM, 8PM), at which Lizz suggested exercises to the user.

5.2.1 Prototype 1: Rule-based system without user feedback

The implementation of the first prototype was fairly self-explanatory; Lizz asked the initial static user profile questions, and based on the user’s answers, Lizz set up an exercise recommendation schedule. These recommendations did not change over time, as this imple-

mentation was solely based on pre-programmed rules. Figure 7 shows the interface of Lizz doing the suggestions.



Figure 7: Images where Lizz suggests exercises in prototype 1.

5.2.2 Prototype 2: Combination of rule-Based and reinforcement learning with user preferences

This second prototype built upon the implementation of the first. However, this prototype used reinforcement learning in the form of a Q-learning algorithm. This algorithm was implemented in the following structure:

- A **Q-table** was initialized, which is a table of each state-action pair, which in our case were time and exercise, so the size of the **Q-table** was 4x6. All time-exercise pairs were initialized based on the rules that were defined for the first prototype, so every exercise and time had an initial score based on the static user profile.
- An **exploration rate** was defined, which was initialized at 0.7 and descended with 0.01 at every iteration (with a minimum value of 0.2). The **exploration rate** is used to balance exploration and exploitation. It determines the probability with which an agent chooses to explore new actions rather than exploiting the actions it already believes to be optimal based on its current knowledge.
- A **learning rate** was defined and set to 0.4. The **learning rate** controls the extent to which newly acquired information overrides the existing **Q-value** estimates in the **Q-table** during the update process. It determines the balance between the weight given to new information and the weight given to existing knowledge. A higher **learning rate** leads to more rapid updates but can be sensitive to noise, while a lower learning rate provides more stability but slower convergence.

- A **discount factor** was defined and set to 0.6. The **discount factor** determines the importance of future rewards compared to immediate rewards. A higher **discount factor** values long-term rewards more, while a lower discount factor prioritizes immediate rewards. It affects the agent’s decision-making process by considering the cumulative reward over time.
- A function **suggestExercise** was called, which first checks whether the algorithm is going to exploit or explore in the current iteration by taking a random value between 0 and 1, and checking if it is smaller than the **exploration rate**, if it is, it return two random exercises. If the algorithm is going to exploit in this iteration, it takes the **Q-values** at the current time, sorts the exercises based on their **Q-values** in descending order, and it returns the top two exercises with the highest **Q-values**.
- A **reward** was calculated every time a user chose one of the suggested exercises (or chose none of the suggested exercises). When the user chose an exercise, the **reward** for this exercise was set to 1.0, and the **Q-value** for the other suggested exercise remained what it was before. When the user has chosen to do none of the suggested exercises, this meant that the suggestions were not fitting at the current time, so the **reward** for the two suggested exercises was set to -1.0.
- Then a function **updateQTable** was called which retrieves the current **Q-value** for the given time and exercise. It calculates the maximum **Q-value** among all exercises for the current time. It applies the **Q-learning** update rule to calculate a new **Q-value** based on the current **Q-value**, the **reward**, the **learning rate**, and the **discount factor**. This formula is as follows:

$$Q(s, a) = Q(s, a) + \alpha \cdot (r + \gamma \cdot \max'_a Q(s', a') - Q(s, a))$$

Where:

- $Q(s, a)$ represents the Q-value of taking action (exercise) a in state (time point) s .
- α is the learning rate.
- r is the immediate reward received after taking action a in state s .
- γ is the discount factor.
- $\max'_a Q(s', a')$ represents the maximum Q-value among all possible actions a' in the resulting state s' after taking action a in state s .

The function updates the **Q-table** with the new **Q-value** for the current time and exercise.

This algorithm iterates over the pre-defined time points during the day, and thus updates the values of the **Q-table** after every iteration. Figure 8 shows the interface of Lizz suggesting the different exercises.



Figure 8: Images where Lizz suggests exercises in prototype 2 and 3.

5.2.3 Prototype 3: Combination of rule-based and reinforcement learning with user preferences and user state

The implementation of the third prototype extended the implementation of the second. However, for this prototype the states in the **Q-table** had more variables. In this prototype, Lizz asked the user at every first time point of the day (which was 8 AM in our implementation) questions to update the dynamic user profile. We chose to ask the following questions:

- How was your sleep? (options: good, not good)
- How are you feeling today? (options: well, unwell)

We chose two simple questions with only two options per question to demonstrate how a small extension of the user profile affects the time it takes for the reinforcement learning algorithm to converge, which will become apparent in sections 6.2 and 6.3. The **Q-learning** algorithm remained the same as in prototype 2, however, the states consisted of not only time, but (time, sleep, mood) which increased the size of the **Q-table** from 4×6 to $4 \times 2 \times 2 \times 6 = 24 \times 6$. Figure 9 shows the interface of the questions that Lizz asks the user every day to keep up the dynamic user profile. The interface for the exercise suggestions did not differ from that of prototype 2, see figure 8.




Figure 9: Images where Lizz asks the user questions for the dynamic user profile.

6 Evaluation

To assess the effectiveness of the models, we will examine two perspectives: required effort for the user, considering the level of burden imposed, and the degree of personalization, evaluating how well the model aligns with the user’s preferences. In this section we will first describe how we will evaluate the prototypes in terms of required effort and degree of personalization, after which we will present our results, and finally we will interpret the results.

6.1 Evaluation Measures

6.1.1 Required effort

As discussed in section 4.4, we aim to minimize the user’s effort in enabling the algorithm to effectively personalize the suggestions. To assess required effort of high-level personalization using reinforcement learning, the following measurement approach was employed. prototype 1 was not considered for this calculation, as no reinforcement learning was used for that prototype, we already know that the required effort only entails that of answering the initial questions.

A script was developed to execute the algorithms of prototypes 2 and 3, simulate user input, and determine the convergence of the Q-table. Convergence implies that the algorithm has learned the optimal policy. This convergence was evaluated by calculating the average difference in Q-values across all exercise-time combinations and comparing it to a convergence

threshold. The following formula was used:

$$c(q_t, q_{t-1}, \varepsilon) = \begin{cases} 1, & \text{if } |q_t(s, a) - q_{t-1}(s, a)| < \varepsilon, \forall s \in S, a \in A \\ 0, & \text{otherwise} \end{cases}$$

Where

- c is the "convergence" function (termination condition) that returns 1 (true) if the RL algorithm has converged to some small enough neighbourhood of value functions (where those value functions are "indistinguishable"), and 0 otherwise
- q_t is the value function at iteration t
- ε is a threshold (precision or tolerance) value, which is a hyper-parameter that you can set depending on your "tolerance".

In this particular test, the convergence threshold ε was set at 0.25. This specific value was carefully chosen through a process of experimentation and fine-tuning, aiming to find a balance between attaining satisfactory convergence and minimizing unnecessary computational burden.

In addition to the average difference in Q-values, we considered the visit count of state-action pairs. The visit count represents the number of times each state-action pair has been visited during the learning process. This information provides insights into the exploration and exploitation balance of the algorithm. The convergence calculation only considered the state-action pairs that have been visited a sufficient number of times, ensuring that the evaluation focused on the most explored areas of the Q-table.

To evaluate the required effort, we looked at convergence of the Q-tables of the implemented models for 2 and 3. We measured the duration in days it took for the Q-values to reach convergence. To expedite the learning process and enhance the early recommendations, a static user profile was implemented. This user profile provided the algorithm with initial insights and helped it make more informed suggestions right from the start. Although these initial values gave the algorithm a head start, it's worth mentioning that they did not impact the convergence time of the Q-table.

In Q-learning, the algorithm improves exercise recommendations by continuously updating its estimates based on observed rewards and future Q-values. While non-zero initial values can speed up initial learning, their impact diminishes as the algorithm explores the environment and relies more on actual rewards. Consequently, initializing the Q-table with non-zero values does not significantly affect the overall convergence time compared to initializing with zeros. The algorithm's interactions with the environment and reinforcement learning updates are the key drivers of convergence and the discovery of optimal exercise choices.

6.1.2 Degree of personalization

As mentioned in section 4.4, the level of personalization pertains to the extent to which a system or service is tailored to suit the unique preferences of individual users. This encompasses customizing the companion's modules to meet their specific requirements. The evaluation of

personalization entails gauging the success of providing personalized recommendations based on user-specific data.

To assess the level of personalization in the three prototypes, we examined the following factors to provide qualitative insights into the level of personalization. The first measure is the extent of user profile information as a measure of personalization. The availability of user data allows the system to better understand individual needs and tailor its recommendations accordingly. Justifying this measure, we recognize that a rich user profile contributes to a higher potential for personalization, as it enables the system to make more informed and targeted recommendations. We will assess the user profile size by the amount of parameters that are in the user profile for each prototype.

Additionally, we assess the model’s adaptability as a measure of personalization. The ability of the system to learn from user feedback and adapt its recommendations over time is crucial for effective personalization. By evaluating the model’s adaptability, we can determine whether the system dynamically adjusts its recommendations based on user interactions and changing preferences. The adaptability is assessed by the amount of user feedback the model receives for each prototype. This qualitative assessment enables us to capture the dynamic nature of adaptability and understand how well the system responds to the evolving needs and preferences of individual users.

6.2 Results

In this section, we present the results of the evaluation of the three prototypes, each exploring different approaches of personalization in Lizz, the digital healthcare companion. The required effort was evaluated as well as the degree of personalization, which was evaluated using user profile size and adaptability.

Prototype	Effort required	Personalization level	
		User profile size	Adaptability
Prototype 1	None	3 static variables	Not adaptive
Prototype 2	Convergence after 30 days	3 static variables; 1 dynamic variable	Moderately adaptive
Prototype 3	Convergence after 100 days	3 static variables; 3 dynamic variables	Moderately adaptive

Table 2: Results of the Evaluation of the Three Prototypes - It shows the effort required which is expressed in amount of days it takes for the model to converge, and the personalization level, which entails the user profile size in the amount of variables, and the adaptability of the system, based on the amount of user feedback.

Table 2 contains the results of the evaluation of the three prototypes. We can see that the first prototype, the rule-based system without user feedback required no effort, had a user profile size of three static variables, namely the initial setup questions, and the system was not adaptive at all.

The second prototype, the combination of rule-Based and reinforcement learning with moderate user feedback, showed convergence after an average of 30 days, the user profile

contained the same three static variables and one dynamic variable. As this prototype uses feedback in the form of which exercise the user chooses at a certain time, the adaptability was categorized as moderate.

The third prototype, the combination of rule-Based and reinforcement learning with user feedback and user state monitoring, showed convergence after an average of 100 days, the user profile again contained the initial three static variables and three dynamic ones, and it was equally adaptive as the second prototype, as the same amount of feedback is used.

6.3 Discussion

Three prototypes were developed to demonstrate the advantages and disadvantages of different personalization tools within Lizz. Prototype 1 used a rule-based system without user feedback, prototype 2 introduced reinforcement learning with user preferences, and prototype 3 extended reinforcement learning with consideration of additional factors on the user state.

Prototype 1: Rule-based system without user feedback

The first prototype, which relied solely on pre-programmed rules, shows no required effort and a very low level of personalization. The exercise recommendations were static and did not adapt to user feedback or individual preferences. The evaluation of the required effort in prototype 1 reveals that the Q-table convergence is not applicable in this case, as the recommendations were not based on reinforcement learning or user feedback. The burden on the user in terms of effort is minimal since the recommendations were determined by the initial static user profile.

In this prototype, only the initial static user profile was populated and no other data were collected during use. This makes the user profile very small, as it only contained three variables, namely age, activity level and most active time during the day. As for adaptability, the exercise recommendations were merely based on the initial profile, thus remained static and did not adapt to individual user needs or preferences over time. We can conclude that the model is not adaptive at all.

Prototype 2: Combination of rule-Based and reinforcement learning with user preferences

This second prototype, with reinforcement learning and moderate user feedback, requires more effort from the user and achieves a higher level of personalization than the first prototype. The evaluation of the required effort shows that the Q-table reached convergence after an average of 30 days, indicating that the algorithm successfully learned the optimal exercise recommendations based on user feedback. Notably, despite the relatively small size of the Q-table in this case (4x6), it is crucial to acknowledge that the algorithm still demanded a substantial amount of time to discover the optimal policy, even with this limited number of states. This

This prototype also used the initial user profile information, and during use it collected data about the user's preferred exercises at certain times during the day. In comparison to prototype 1, the user profile in this case becomes slightly more comprehensive, encompassing four variables: age, activity level, the most active time of the day, and preferred exercises.

The incorporation of reinforcement learning in this prototype allowed Lizz to improve exercise recommendations over time based on user feedback. The algorithm learned from user behavior and adapted the recommendations accordingly. By offering choices and learning from user selections, the algorithm gradually tailored the exercise plan to individual preferences. The model is thus more adaptive than that of prototype 1. The moderate level of personalization in this prototype demonstrates the potential of reinforcement learning techniques to recommendations based on user behavior.

The 30-day convergence time observed in this prototype might be considered reasonable for achieving the level of personalization attained. It allows the algorithm to learn user preferences thoroughly while still being able to adapt to gradual changes over time. This balance between effort required, measured by the convergence time, and the achieved level of personalization ensures that the algorithm can provide tailored exercise recommendations that align with the evolving preferences of the user.

Prototype 3: Combination of rule-based and reinforcement learning with user preferences and user state

In the third prototype, a combination of rule-based and reinforcement learning techniques was employed to improve the exercise recommendation system. This prototype requires a lot more effort than the other two prototypes, while it is also more personalized. The increased complexity of considering a wider range of factors, namely the emotional state and fatigue level, contributed to the substantially longer convergence time of 100 days. In this prototype, the Q-table expanded to a size of 26x6, highlighting the significant increase in the number of states due to the addition of merely two binary factors. This indicates that incorporating higher levels of personalization and considering more factors can impose a substantial computational load on the algorithm.

The third prototype again uses the initial user profile and learns exercise preferences as in prototype 2, but also expands the user profile by asking the user about their sleep and emotional state on a daily basis. This expands the user profile to six variables, namely age, activity level, the most active time of the day, preferred exercises, sleep, and emotional state, resulting in a more refined exercise recommendation. The algorithm, similar to prototype 2, was able to gradually learn to provide more suitable recommendations, making it an adaptive model.

However, the 100-day convergence time in this prototype raises concerns about potential misalignment between user preferences and the algorithm's learning capabilities. User preferences can change faster than the algorithm can adapt, diminishing the practicality of higher levels of personalization achieved. The significant effort and computational load required for expanded factors may outweigh the benefits. It is crucial to strike a balance between personalization and practicality, considering the frequency and magnitude of user preference changes to determine the optimal level of personalization that provides meaningful recommendations without overwhelming the algorithm.

In conclusion, the three prototypes of the exercise recommendation system showcase the advantages and disadvantages of different personalization tools within Lizz. Prototype 1 demonstrates the limitations of a rule-based system without user feedback, offering minimal personalization and no adaptability. Prototype 2, combining rule-based and reinforcement

learning with moderate user feedback, achieves a higher level of personalization by gradually learning from user behavior. It strikes a reasonable balance between effort required and the achieved level of personalization, ensuring the algorithm’s adaptability to gradual changes in user preferences. Prototype 3, incorporating user feedback and considering additional factors, demonstrates the potential for further personalization but at the cost of increased effort and longer convergence time. It highlights the importance of striking a balance between personalization and practicality, as rapid changes in user preferences could challenge the algorithm’s ability to adapt effectively. Overall, these prototypes provide insights into the trade-offs and considerations involved in developing personalized recommendation systems, allowing for informed decision-making in tailoring the system to individual user needs.

7 Discussion

Although the personalization of Socially Assistive Robotics has gained increasing attention in recent years, the absence of a standardized approach remains a significant challenge. This lack of standardization hinders the comparison of findings across studies, as different research groups employ diverse methods and evaluation metrics to gauge the effectiveness of personalization [17]. Moreover, there is still a limited understanding of how to effectively gather and leverage user preferences and needs, which significantly impacts the development of personalized SAR systems [17].

To address these critical gaps in the literature and contribute to the advancement of personalized SAR, this study aimed to explore various forms of personalization within the context of a social companion. The primary objective was to provide insights into the optimal utilization of each form. In this section, we delve into the potential implications and limitations of the findings, and outline potential avenues for future research.

7.1 Implications of the Results

The study utilized Lizz, a digital healthcare companion developed by ConnectedCare [44], as a platform for exploring personalization techniques in remote patient support. Three prototypes were developed to demonstrate the advantages and disadvantages of different personalization tools within Lizz. Prototype 1 used a rule-based system without user feedback, prototype 2 introduced reinforcement learning with moderate user feedback, and prototype 3 extended reinforcement learning with consideration of additional factors on the user state. The prototypes involved JavaScript programming and WebSockets for communication, showcasing Lizz’s interaction with users and personalized exercise recommendations.

The three prototypes of the exercise recommendation system in Lizz demonstrate the advantages and disadvantages of different personalization tools. Prototype 1, a rule-based system without user feedback, offers minimal personalization and no adaptability. Prototype 2, combining rule-based and reinforcement learning with moderate user feedback, achieves a higher level of personalization by gradually learning from user behavior. It strikes a reasonable balance between effort required and the achieved level of personalization. Prototype 3, with consideration of additional factors, demonstrates the potential for further personalization but at the cost of increased effort and longer convergence time. Rapid changes in user

preferences may challenge the algorithm’s ability to adapt effectively in this case.

The distinction between dynamic and static data within the user profile has important implications for the choice of personalization mechanisms. Static data, encompassing information that remains relatively stable over time, such as user characteristics, can be effectively utilized by rule-based systems without user feedback. These mechanisms provide initial personalized recommendations based on static data but may lack adaptability to individual needs and preferences as they do not consider dynamic changes. On the other hand, incorporating dynamic data, such as the user’s current context and state, into the personalization process requires more advanced mechanisms like reinforcement learning and user feedback. These mechanisms allow for continuous learning and adaptation based on the evolving user state. However, it is crucial to consider the trade-off between personalization and practicality, as rapidly changing user preferences may pose challenges to the algorithm’s ability to adapt effectively. Therefore, achieving a balanced approach that combines static and dynamic data, along with appropriate personalization mechanisms, is necessary to strike the right balance between personalized recommendations and system manageability.

In terms of optimal implementation, the choice of personalization mechanism depends on user preferences, the desired level of personalization, and the willingness to invest effort. For users who prefer simplicity and minimal effort, a rule-based system without user feedback can be suitable. It provides basic exercise recommendations based on initial user profiles but lacks adaptability to individual needs and preferences over time. For users seeking a moderate level of personalization and who are willing to provide feedback, a combination of rule-based and reinforcement learning mechanisms can be beneficial. This approach gradually learns from user behavior and offers suitable recommendations while maintaining a reasonable convergence time. However, for users who highly value personalization and are willing to invest more effort, incorporating user feedback and user state into the personalization process can provide refined exercise recommendations. It is important to consider the trade-off of longer convergence time and potential misalignment with rapidly changing user preferences in this case.

Overall, achieving a balance between personalization and practicality is crucial, considering the frequency and magnitude of user preference changes. Tailoring the system to individual user needs while maintaining a manageable convergence time ensures meaningful recommendations without overwhelming the algorithm. These findings have significant implications for the development of personalized recommendation systems, enabling informed decision-making in exercise recommendation platforms.

7.2 Limitations

However, it is important to acknowledge the limitations of this research. Firstly, it should be noted that this study specifically uses the example of Lizz being for exercise recommendations as a carrier. While exercise recommendations are critical for healthcare assistance, it is important to understand that personalization strategies may have distinct implications and effectiveness across different domains. As a result, while evaluating the findings, it is important to remember the specific context of exercise recommendations and exercise caution when extrapolating the results to other domains. Future studies could investigate how personalization approaches might be utilized in a broader range of social circumstances to

acquire a better understanding of their benefits and limitations. This will provide significant insight into their generalizability and effectiveness beyond exercise guidelines.

Second, the evaluation of the personalization strategies relied on a simulated scenario conducted within a controlled environment. While this approach enabled controlled research and comparisons, it may not fully capture the intricacies and complexities of genuine user interactions. Actual user interactions possess subtleties and nuances that may be absent in a simulated scenario. In real-world situations, personalization methods can face challenges and be less effective due to factors such as varying user motivation, preferences, and external influences. To gain a deeper understanding of the effects of personalization strategies on user engagement, adherence, and overall outcomes, future research should encompass user studies and longitudinal evaluations conducted in real-world settings. These studies would provide valuable insights into the practical implications and challenges associated with the implementation of personalized healthcare companions. Regrettably, the scope of this thesis did not permit the inclusion of user studies and evaluations, thus limiting the availability of insights into the practical aspects of implementing personalized healthcare companions in real-world scenarios.

Additionally, the evaluation of the required effort focused primarily on the convergence of the Q-tables as a measure of learning effectiveness. While the convergence of the Q-table indicates the algorithm’s ability to learn optimal exercise recommendations, it may not fully capture all aspects of user effort or system usability. User effort can include various factors such as the time and cognitive load required to interact with the companion, willingness to provide feedback, and perceived usefulness of personalized recommendations. Future research may include more comprehensive assessments of user effort and user satisfaction through user research, feedback collection, and subjective measurements of user experience.

7.3 Future work

In terms of future work, there are several promising avenues for further exploration in the field of personalization in social companions like Lizz, which have been briefly touched upon in the previous section. First, using more data sources could significantly improve personalization possibilities. Real-time data on the user’s health and well-being could be gathered by incorporating data from wearable technology or physiological assessments. With the help of this new information, exercise recommendations might be made that are even more individualized and accurate and do not involve the user exerting a lot of effort. Considerations like heart rate, sleep quality, or stress levels could be made collect data for personalization. The next stage would be to investigate the integration of such data sources and assess how they affect the personalization procedure.

Secondly, it is also important to examine how different tools for personalization can influence user motivation and engagement. Exploring the effects of various personalization techniques on factors such as user motivation, retention, and overall engagement can provide valuable insights that go beyond the scope of this study. By discovering strategies that successfully inspire and captivate users, we can create personalized systems that maximize long-term commitment and produce positive results.

Additionally, extending the research beyond exercise recommendations to other aspects of the digital healthcare companion could contribute to a more holistic understanding of

personalization’s role in optimizing overall effectiveness. For instance, exploring personalization techniques in the context of attention and structure support, such as reminders or goal-setting strategies, could further enhance the user experience and outcomes. Furthermore, investigating personalization in dietary advice or nutritional support could provide a comprehensive approach to improving overall health and well-being. Understanding how personalization can be effectively applied to different aspects of a digital healthcare companion would enable the development of comprehensive and tailored systems that address multiple dimensions of user health.

8 Conclusion

This thesis aimed to gain valuable insights into the benefits, limitations, and future directions of personalization techniques in social robotics. By exploring various personalization techniques within the context of Lizz, a digital healthcare companion, this research contributed to the understanding of this field.

Our study demonstrated that increasing levels of personalization resulted in more adaptive and tailored exercise recommendations. We implemented three examples, ranging from rule-based systems to reinforcement learning with extensive user feedback. The results indicated that increasing personalization led to exercise recommendations that were better aligned with individual preferences and needs. However, it is important to note that higher levels of personalization required additional effort and time before the model reached optimization and this might not weigh up to the benefits of the increased personalization.

It is important to consider the limitations of this research. The findings are specific to exercise recommendations in healthcare and may not directly apply to other domains. Future studies should explore the generalizability of personalization approaches across different social contexts.

Additionally, our evaluation relied on a simulated scenario, which may not fully capture real-world complexities. Future research should include user studies in authentic settings to gain deeper insights into practical implementation challenges.

Moving forward, there are several promising avenues for future work. Incorporating real-time data from wearable technology and assessing the effects of personalization on user motivation and engagement during longitudinal user studies are areas worth exploring. Extending research beyond exercise recommendations to other aspects of a digital healthcare companion, such as attention and structure support or dietary advice, would contribute to a more comprehensive understanding of personalization’s role in optimizing effectiveness.

In conclusion, this thesis highlights the significance of personalization in social companions and its potential to enhance user support. By effectively designing and implementing personalized systems, we can address individual needs, preferences, and health goals. Personalized social companions have the potential to revolutionize healthcare, providing tailored support, improving engagement, and ultimately leading to better outcomes. We hope this research serves as a foundation for future studies and inspires further exploration in the field of personalization, ultimately benefiting individuals in their daily lives and overall well-being.

Acknowledgements

I would like to express my sincere gratitude to dr. Martijn Vastenburg for his guidance and support as the supervisor of this thesis. His dedication is evident through our weekly meetings and the constructive feedback he provided throughout the process, which greatly contributed to the development of this research.

I would also like to extend my appreciation to dr. Serge Thill for his role as the second reader of this thesis. His expertise and willingness to assist, particularly in the more technical aspects, proved to be immensely beneficial throughout the research process.

Furthermore, I am grateful to Luca Leisten, Krunoslav Domić, and Janna Albers, members of the ConnectedCare team, for their valuable contributions. Their input during brainstorming sessions and assistance in setting up Lizz have been very helpful to this research.

Lastly, I would like to thank all the individuals who have supported me during this journey, including my friends, roommates, partner, and family members. Their encouragement and understanding have been a constant source of motivation.

References

- [1] A. Robaczewski, J. Bouchard, K. Bouchard, and S. Gaboury, “Socially assistive robots: The specific case of the nao,” *International Journal of Social Robotics*, vol. 13, pp. 795–831, 2021.
- [2] Z. Zhao, Y. Ma, A. Mushtaq, *et al.*, “Applications of robotics, artificial intelligence, and digital technologies during covid-19: A review,” *Disaster Medicine and Public Health Preparedness*, vol. 16, no. 4, pp. 1634–1644, 2022.
- [3] D. Feil-Seifer and M. J. Mataric, “Defining socially assistive robotics,” in *9th International Conference on Rehabilitation Robotics, 2005. ICORR 2005.*, IEEE, 2005, pp. 465–468.
- [4] M. Čaić, D. Mahr, and G. Oderkerken-Schröder, “Value of social robots in services: Social cognition perspective,” *Journal of Services Marketing*, vol. 33, no. 4, pp. 463–478, 2019.
- [5] “World health organization. ageing and health.” Accessed: 2022-02-11. (2022).
- [6] G. P. Soriano, Y. Yasuhara, H. Ito, *et al.*, “Robots and robotics in nursing,” in *Health-care*, MDPI, vol. 10, 2022, p. 1571.
- [7] M. Mataric, A. Tapus, and D. Feil-Seifer, “Personalized socially assistive robotics,” in *Workshop on Intelligent Systems for Assisted Cognition Rochester, NY*, Citeseer, 2007.
- [8] F. Dino, R. Zandie, H. Abdollahi, S. Schoeder, and M. H. Mahoor, “Delivering cognitive behavioral therapy using a conversational social robot,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2019, pp. 2089–2095.
- [9] B. Zhou, K. Wu, P. Lv, *et al.*, “A new remote health-care system based on moving robot intended for the elderly at home,” *Journal of healthcare engineering*, vol. 2018, 2018.
- [10] S. Serholt and W. Barendregt, “Students’ attitudes towards the possible future of social robots in education,” Aug. 2014.
- [11] M. Smakman, P. Vogt, and E. A. Konijn, “Moral considerations on social robots in education: A multi-stakeholder perspective,” *Computers & education*, vol. 174, p. 104317, 2021.
- [12] W. Johal, “Research trends in social robots for learning,” *Current Robotics Reports*, vol. 1, pp. 75–83, 2020.
- [13] R. Van den Berghe, J. Verhagen, O. Oudgenoeg-Paz, S. Van der Ven, and P. Lese-man, “Social robots for language learning: A review,” *Review of Educational Research*, vol. 89, no. 2, pp. 259–295, 2019.
- [14] L. Xiao and V. Kumar, “Robotics for customer service: A useful complement or an ultimate substitute?” *Journal of Service Research*, vol. 24, no. 1, pp. 9–29, 2021.

- [15] A. Følstad and M. Skjuve, “Chatbots for customer service: User experience and motivation,” in *Proceedings of the 1st international conference on conversational user interfaces*, 2019, pp. 1–9.
- [16] C. V. Mischia, F. Poetze, and C. Strauss, “Chatbots in customer service: Their relevance and impact on service quality,” *Procedia Computer Science*, vol. 201, pp. 421–428, 2022.
- [17] A. Kubota and L. D. Riek, “Methods for robot behavior adaptation for cognitive neurorehabilitation,” *Annual review of control, robotics, and autonomous systems*, vol. 5, pp. 109–135, 2022.
- [18] C. Moro, G. Nejat, and A. Mihailidis, “Learning and personalizing socially assistive robot behaviors to aid with activities of daily living,” *ACM Transactions on Human-Robot Interaction (THRI)*, vol. 7, no. 2, 2018.
- [19] E. Wiese, G. Metta, and A. Wykowska, “Robots as intentional agents: Using neuroscientific methods to make robots appear more social,” *Frontiers in psychology*, vol. 8, p. 1663, 2017.
- [20] G. D’Onofrio, L. Fiorini, H. Hoshino, *et al.*, “Assistive robots for socialization in elderly people: Results pertaining to the needs of the users,” *Aging clinical and experimental research*, vol. 31, pp. 1313–1329, 2019.
- [21] K. Winkle, P. Caleb-Solly, A. Turton, and P. Bremner, “Social robots for engagement in rehabilitative therapies: Design implications from a study with therapists,” in *Proceedings of the 2018 acm/ieee international conference on human-robot interaction*, 2018, pp. 289–297.
- [22] D. Raigoso, N. Céspedes, C. A. Cifuentes, A. J. Del-Ama, and M. Múnica, “A survey on socially assistive robotics: Clinicians’ and patients’ perception of a social robot within gait rehabilitation therapies,” *Brain sciences*, vol. 11, no. 6, p. 738, 2021.
- [23] M. Pino, M. Boulay, F. Jouen, and A.-S. Rigaud, ““are we ready for robots that care for us?” attitudes and opinions of older adults toward socially assistive robots,” *Frontiers in aging neuroscience*, vol. 7, p. 141, 2015.
- [24] R. D. Benedictis, A. Umbrico, F. Fracasso, G. Cortellessa, A. Orlandini, and A. Cesta, “A dichotomic approach to adaptive interaction for socially assistive robots,” *User Modeling and User-Adapted Interaction*, pp. 1–39, 2022.
- [25] A. Cesta, G. Cortellessa, A. Orlandini, and A. Umbrico, “A cognitive loop for assistive robots-connecting reasoning on sensed data to acting,” in *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, IEEE, 2018, pp. 826–831.
- [26] A. Umbrico, R. De Benedictis, F. Fracasso, A. Cesta, A. Orlandini, and G. Cortellessa, “A mind-inspired architecture for adaptive hri,” *International Journal of Social Robotics*, pp. 1–21, 2022.
- [27] C. Di Napoli, G. Ercolano, and S. Rossi, “Personalized home-care support for the elderly: A field experience with a social robot at home,” *User Modeling and User-Adapted Interaction*, pp. 1–36, 2022.

- [28] S. Cooper, A. Di Fava, C. Vivas, L. Marchionni, and F. Ferro, “Ari: The social assistive robot and companion,” in *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, IEEE, 2020, pp. 745–751.
- [29] R. Khosla, M.-T. Chu, S. M. S. Khaksar, K. Nguyen, and T. Nishida, “Engagement and experience of older people with socially assistive robots in home care,” *Assistive Technology*, vol. 33, no. 2, pp. 57–71, 2021.
- [30] A. Umbrico, A. Cesta, G. Cortellessa, and A. Orlandini, “A holistic approach to behavior adaptation for socially assistive robots,” *International Journal of Social Robotics*, vol. 12, no. 3, pp. 617–637, 2020.
- [31] H. W. Park, I. Grover, S. Spaulding, L. Gomez, and C. Breazeal, “A model-free affective reinforcement learning approach to personalization of an autonomous social robot companion for early literacy education,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 687–694.
- [32] T. W. Bickmore, L. M. Pfeifer, D. Byron, *et al.*, “Usability of conversational agents by patients with inadequate health literacy: Evidence from two clinical trials,” *Journal of health communication*, vol. 15, no. S2, pp. 197–210, 2010.
- [33] H. L. Bradwell, K. J. Edwards, R. Winnington, S. Thill, and R. B. Jones, “Companion robots for older people: Importance of user-centred design demonstrated through observations and focus groups comparing preferences of older people and roboticists in south west england,” *BMJ open*, vol. 9, no. 9, e032468, 2019.
- [34] S. Rossi, F. Ferland, and A. Tapus, “User profiling and behavioral adaptation for hri: A survey,” *Pattern Recognition Letters*, vol. 99, pp. 3–12, 2017.
- [35] S. Schneider and F. Kummert, “Comparing robot and human guided personalization: Adaptive exercise robots are perceived as more competent and trustworthy,” *International Journal of Social Robotics*, vol. 13, no. 2, pp. 169–185, 2021.
- [36] W. H. Organization *et al.*, “Icf: International classification of functioning, disability and health,” 2001.
- [37] M. I. Ahmad, Y. Gao, F. Alnajjar, S. Shahid, and O. Mubin, “Emotion and memory model for social robots: A reinforcement learning based behaviour selection,” *Behaviour & Information Technology*, vol. 41, no. 15, pp. 3210–3236, 2022.
- [38] H. Abdollahi, M. Mahoor, R. Zandie, J. Sewierski, and S. Qualls, “Artificial emotional intelligence in socially assistive robots for older adults: A pilot study,” *IEEE Transactions on Affective Computing*, 2022.
- [39] M. H. Lee, D. P. Siewiorek, A. Smailagic, A. Bernardino, and S. B. Badia, “Towards personalized interaction and corrective feedback of a socially assistive robot for post-stroke rehabilitation therapy,” in *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, IEEE, 2020, pp. 1366–1373.
- [40] R. S. Sutton, A. G. Barto, *et al.*, “Reinforcement learning,” *Journal of Cognitive Neuroscience*, vol. 11, no. 1, pp. 126–134, 1999.

- [41] J. Hemminghaus and S. Kopp, “Towards adaptive social behavior generation for assistive robots using reinforcement learning,” in *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, 2017, pp. 332–340.
- [42] C. Di Napoli and S. Rossi, “A layered architecture for socially assistive robotics as a service,” in *2019 IEEE international conference on systems, man and cybernetics (SMC)*, IEEE, 2019, pp. 352–357.
- [43] I. Papadopoulos, R. Lazzarino, S. Miah, T. Weaver, B. Thomas, and C. Koulouglioti, “A systematic review of the literature regarding socially assistive robots in pre-tertiary education,” *Computers & Education*, vol. 155, p. 103 924, 2020.
- [44] “Lizz – the digital health care assistant.” Accessed: 2022-02-15. (2022).