Comparing Robot and Human guided Personalization: Adaptive Exercise Robots are Perceived as more Competent and Trustworthy

Sebastian Schneider^{1,2} · Franz Kummert¹

Accepted: 22 January 2020 / Published online: 8 February 2020
© The Author(s) 2020

Abstract

Learning and matching a user's preference is an essential aspect of achieving a productive collaboration in long-term Human–Robot Interaction (HRI). However, there are different techniques on how to match the behavior of a robot to a user's preference. The robot can be *adaptable* so that a user can change the robot's behavior to one's need, or the robot can be *adaptive* and autonomously tries to match its behavior to the user's preference. Both types might decrease the gap between a user's preference and the actual system behavior. However, the Level of Automation (LoA) of the robot is different between both methods. Either the user controls the interaction, or the robot is in control. We present a study on the effects of different LoAs of a Socially Assistive Robot (SAR) on a user's evaluation of the system in an exercising scenario. We implemented an online preference learning system and a user-adaptable system. We conducted a between-subject design study (*adaptable* robot vs. *adaptive* robot) with 40 subjects and report our quantitative and qualitative results. The results show that users evaluate the *adaptive* robots as more competent, warm, and report a higher alliance. Moreover, this increased alliance is significantly mediated by the perceived competence of the system. This result provides empirical evidence for the relation between the LoA of a system, the user's perceived competence of the system, and the perceived alliance with it. Additionally, we provide evidence for a proof-of-concept that the chosen preference learning method (i.e., Double Thompson Sampling (DTS)) is suitable for online HRI.

Keywords Human–robot interaction · Preference learning · Adaptive robots · User experience

1 Introduction

Future scenarios of social robots envision a personable system that is flexible and adapts itself to the user's preferences [20]. Typical applications of social robots include, for example, social assistance for physical or cognitive exercising [12,26]. Different users or target groups can have different preferences or requirements towards the system interaction. However, anticipating all user types and pre-programming the system for their needs will be an obstacle to deploy robots in

domestic settings and engage users beyond an exploration phase. In these situations, robots would need capabilities to adjust to different user profiles (e.g., match a user's personality [1]). In web-based applications, this requirement is already widely accomplished (e.g., recommender systems on Amazon, Netflix). However, it remains a challenge for social robots that have no access to an extensive user database and thus cannot utilize techniques like collaborative filtering. Thus, social robots face the *cold start problem*, which requires the system to gather initial data to personalize the interaction experience. Nevertheless, deploying an *adaptive* system still comes with some known difficulties:

First, querying the user for information in real-time HRI might be more cost-intensive than in web-based applications. Cakmak et al. [6] showed that a constant stream of questions in a Learning by Demonstration (LdB) task annoys users.

Second, autonomous personalization decisions by robots could result in diametral effects for the HRI experience satisfaction. When a robot controls the personalization process, it could lead to technology disuse due to wrongly learned

✉ Sebastian Schneider
sebschne@techfak.uni-bielefeld.de;
sebastian@robot.soc.i.kyoto-u.ac.jp

Franz Kummert
franz@techfak.uni-bielefeld.de

¹ Applied Informatics, Faculty of Technology, Bielefeld University, Bielefeld, Germany

² HRI Lab, Department of Social Informatics, Kyoto University, Kyoto, Japan

user profiles or users prefer to be in control [5]. Thus, it is essential to consider whether people prefer an interface to adjust the system behavior or prefer to let the system control the adaptation process.

These different personalization strategies would influence the system's autonomy and affect the interaction experience. Based on the theory by Epley et al. [10] on anthropomorphization, an autonomous *adaptive* system could create unexpected user experiences [10]. This unanticipated experiences can increase a user's perceived human-likeness of the robot and enhance the credibility of, and trust in the system. In contrast, a user-controlled system would increase the match between the user's expectation and the robot's behavior and therefore reduce anthropomorphic effects. The investigation of these two aspects is the core of this work. We try to find an answer to the question:

What effects have different types of personalization methods on the user's perception of the system, trust in the alliance, and motivation to interact with the system?

To answer this question, we investigate the effects of different personalization behaviors of the system and present a study in the area of robotic exercising companions that compares the impact of an *adaptive* robot versus an *adaptable* robot as an exercising partner for physical activities.

Previous work on robots for exercising and coaching have investigated the motivational effects of using such coaching systems [12,16,44]. However, most studies used only one type of exercise (e.g., arm, or plank exercises). In this work, we present a system that offers a range of activities to the user, which we use to investigate a suitable preference matching framework to provide personalized interaction. Therefore, in the *adaptive* condition, the robot proposes different activities for the user and tries to learn a user's exercising ranking based on comparative user preference feedback. In the *adaptable* condition, the robot is directly controlled by the user, and the user can decide which exercise she/he wants to do with the robot.

Our work contributes to the community by showing that the k-armed Dueling Bandit Problem is a suitable approach for online preference learning in HRI scenarios [55]. Moreover, we provide evidence that the alliance to an *adaptive* robotic exercising partner is perceived as more trustworthy and that this is mediated by a perceived higher competence of the system by the user.

The manuscript is organized as follows. The difference between *adaptive* and *adaptable* robots will be explained in Sect. 2 along with the concepts of automation and alliance, which might be important variables when looking at the adaptivity of a system. Section 3 introduces the system design and Sect. 4 explains the study design to test the effects of a robot's different personalization mechanisms. Section 5 presents the results of the study, which are discussed in Sect. 6. Finally, Sect. 7 gives a conclusion of this work.

2 Adaptation, Automation and Alliance

This section gives a brief introduction into the concepts of adaptation, automation and alliance. Discussing these topics is a challenging task because they are used differently across disciplines (e.g., philosophy, psychology, economics, biology). Therefore, the following explanations can not be exhaustive and will focus mainly on a computer science and psychology perspective.

2.1 Adaptation: Adaptivity Versus Adaptability

In computer science adaptation refers to the informative-based process of adjusting the behavior of an interactive system to meet the need of individual users [47]. Even though computer software or robots are running through many software design cycles, it is hard to anticipate the requirements for every possible user. The goal of *adaptive* processes is to minimize the discrepancy between the user's needs and system behavior after the deployment. This *adaptive* process (see Fig. 1) can either be automatically initiated by the system, in this case, the system is *adaptive* (e.g., the system chooses exercises for the users by itself), or users can adjust the system by themselves, in this case the system is *adaptable* (e.g., the users can choose the exercises by themselves). Adaptation can include different user profiles (e.g., personality), various times (e.g., morning/evening, days of the week, summer/winter), or other user characteristics (e.g., mood, experience).

Previous work in HRI investigated the implementation of *adaptive* processes to match a user's personality, generate empathic behavior, adapt therapy sessions, interaction distance, linguistic style or puzzle skills [18,25,26,31,40,49,50]. The results of these works show an improvement in task performance based on personalized lessons or user personality matching [18,26,49]. Additionally, providing *adaptive* empathic feedback also improved user engagement [25]. Other works present evidence for the feasibility of certain adaptation algorithms (e.g., [31]). Overall, there is evidence that *adaptive* personalization is a crucial capability for robots. Though, there are still many open issues that future works need to target. For example, to which objectives should the robot to adapt? How can the system adapt when the objectives are not apparent? Should the robot communicate the adaptation process and thus make it transparent? Finally, who should be in control of the adaptation process?

Although some have compared *adaptive* robots with experimental baseline control conditions (e.g., [25,26]), to the best of our knowledge, no investigation looked at the effects of robot-initiated personalization versus user-initiated personalization. It is reasonable to argue that the users could control and adjust the robot behavior to their preferences. Leyzberg et al. [26], for example, investigated the

Fig. 1 The adaptable system (dashed lines) comes with a set of behaviors it can offer to the user. The user selects her/his preferred behavior, and the system performs the action. The adaptive system (solid lines) explores which behavior is preferred by the user by querying preference feedback. The system updates its user model and can exploit the obtained information over time. Alternatively, the system can also make stereotypical predictions based on the user personality, as well as the valence or experience of the user. Used images in this figure were not modified. <https://pngimg.com/download/45265> (Robot) by <https://pngimg.com> is licensed under <https://creativecommons.org/publicdomain/zero/1.0/> CC BYNC4.0, no changes made. <https://freessvg.org/tai-chiwoman-silhouette> Silhouette of a woman practicing Tai Chi by <https://FreeSVG.org> is licensed under <https://creativecommons.org/licenses/by-nc/4.0/> CC0 1.0, no changes made

effects of a robot that gives personalized lessons to the user. The robot selects these tutorials based on a decision algorithm. However, also the user could have requested for a specific experience.

Both strategies might match a user's preference and increase interaction satisfaction, but the underlying difference in decision making is fundamental. One can interpret the approaches as either more transparent or as more competent. Generally, the question of whether to build an *adaptive* or *adaptable* system raises the concern of who is in control and how does it affect the interaction experience. The issue of who is in control is, in general, associated with the LoA of the system.

2.2 Level of Automation

An autonomous agent acts based on the information it receives from its sensors, knows in which state it is, and makes a decision accordingly (see [42, ch. 1]). The LoA of a system changes an agent's capability to act and react based on information on its own without any other external control instances. Thus, the agent's LoA alters depending on the task, the agent's environment, and whether a human can interfere in the agent's control loop. This distinction becomes essential where robots are carrying out delicate tasks (e.g., lethal autonomous weapons). There are various frameworks that can be used to identify the LoA of a system (e.g., [9,48]). However, most recently, Beer et al. [2] have

proposed a taxonomy to classify the level of robot autonomy for HRI.

Regardless of the exact different LoA, systems can be categorized as human-in-the-loop systems where the human has to approve a control decision by the autonomous agents, human-on-the-loop where the human is informed about the decision but the agent would carry out a decision if the human operator is not interfering or human-off-the-loop, where a human cannot interfere with the agent's decisions.¹

The relevance to consider the different LoA is apparent in sensible domains such as military operations or medical applications (e.g., surgery or medicine dispenser), but (yet) less apparent in socially assistive fields (e.g., rehabilitation or teaching). Nevertheless, social situations will require to understand whether a social robot should act autonomously, semi-self controlled, or is in full human control also. For the interaction experience, it will be crucial to understand the effects of different LoA. In the course of this work, we are interested in the impact of whether a robot exercising companion is in control to choose the exercises or whether the users can decide which tasks they want to do. The question of which LoA is appropriate and the effects it will have on the interaction experience will be related to the alliance and trust between the users and the SAR [2].

2.3 Alliance

The impact of trust in the alliance between a user and a robot has recently been investigated in use cases in which a robot shows a faulty behavior, gives explanations for actions, or varies the degree of expressivity and vulnerability [29,41,43,51]. Besides these aspects, human's trust in a robot's capabilities also depends on, for example, a high LoA and whether the system makes autonomous decisions [13].

Trust is defined in Human–Computer Interaction (HCI) as “the extent to which a user is confident in and willing to act by, the recommendations, actions, and decisions of an artificially intelligent decision aid” [30, p. 25]. As Madsen and Gregor [28] state, this definition “encompasses both the user's confidence in the system and their willingness to act on the system's decision and advice” [28, p. 1]. Thus, it already incorporates a notion of user trust regarding the willingness to take a system's recommendations into account.

To understand how trust influences HRI, Hancock et al. [17] reviewed different applications where confidence is an essential factor when robots and humans are working together in a team. They state that it is a crucial aspect of industry, space, or warfare applications. However, due to the rise of SAR for rehabilitative, therapeutic or educational tasks, understanding trust for social tasks is also an

¹ Earliest examples of hands-off-the-loop agents are land and naval mines.

essential research topic [12,22]. Hancock et al. [17] found several factors influencing trust in HRI, which are related to the human, the robot, and the environment. Though, the robot-related factors were the most important ones in their meta-review. They found that essential factors influencing the associated trust are the human's perception of the system's behavior, adaptability, competence, and performance [17]. Considering how different types of personalization change the LoA and how this might alter the perceived trust, we question how the manipulation of the LoA (for example how the system adapts or can be adapted) influences the associated competence and the perceived confidence in the system.

Rau et al. [37] investigated the influence of a social robot's LoA on the user's trust in the HRI based on the robot's decision making. They manipulated the robot's LoA by either giving the human the possibility to make a team decision and the robot could suggest a different decision (low autonomy) or the robot makes the team decision and the human can either reject or accept this decision (high autonomy). They hypothesized that a highly autonomous robot would increase the associated trust. Their results show the influence of an autonomous robot on human decision making, but in contrast to the hypothesis, people rated that they trust the low autonomous robot more. However, their result is only marginal significant ($p = .084$). And further investigation on this is needed.

Other works investigated how perceived anthropomorphization influenced perceived trust in autonomous vehicles [52]. Waytz et al. [52] found that the degree of anthropomorphization is associated with higher confidence in its competence. This indicates that the perceived level of skill might also influence the related trust. However, there are, to the best of our knowledge, no other works that investigated the influence of a social robot's LoA, based on its decision-making capabilities, on the perceived trust in the HRI and competence besides the work of Rau et al. [37].

2.4 Objectives and Contribution

This paper has two major objectives. The first objective is to test whether a preference learning approach is suitable for online interaction which is novel for the HRI community and test the feasibility in a realistic use case study in comparison to a user-controlled adjustment method.

The second objective is to investigate how the different personalization methods influence the user's evaluation of the system in terms of its social attributes, trust in the alliance with the system and the motivation to continue the interaction with it. We will discuss the hypothesis related to the later objective in the following.

2.5 Hypotheses

Based on the reviewed literature, we found that there is a substantial lack in understanding the effects of adaptive social robots for future HRI scenarios. We found that it is still uncertain, which is the best way to personalize the robot's behavior (i.e., should it be under control of the user or the robot). Additionally, it remains unclear how the different LoA changes the user trust in the HRI and the perceived competence of the system in social scenarios, as well as how these variables are related to each other. To find empirical evidence that can help answer these questions, we derived four hypotheses from previous works.

Due to the robot's initiative and control of the interaction, people will be likely to associate the robot with higher competence [17]. Since users do not have to control the robot on their own, the robot could create the impression of proactively deciding on its own, which creates unexpected experiences for the user and elicits anthropomorphic reasoning about the agent [10, pp. 873–874]. Thus, based on the theory of Epley et al. [10], we hypothesize that:

Hypothesis 1 Users perceive an *adaptive* robot as more competent than an *adaptable* robot (H_0 : adaptive and adaptable robots are perceived as equally competent).

We hypothesize that this different level of perceived competence is associated with the perceived trust or relationship with the agent.

Because the research from Rau et al. [37] did not show any significant effect on perceived trust in HRI depending on the LoA, we want to continue on the hypothesis that the LoA will affect the associated trust. Likely, the previous research did not find an effect on the trust because the robot was only a marginal partner that was not important for the task. Instead, in our work, the robot is not just a team member but an instructor and exercising partner. Therefore, trust in the alliance will be an essential feature for the relationship between the user and the robot.

Hypothesis 2 The trust in alliance to an *adaptive* robot is rated higher than to an *adaptable* robot.

Since we hypothesize that the participants in the conditions will perceive both the competence and trust differently, it is plausible to argue that the perceived competence and trust will be somehow correlated. Based on the review on trust in HRI, one can argue that users will more likely trust a system that is perceived by the users as competent [17]. Thus, we hypothesize that:

Hypothesis 3 The associated trust in HRI between the conditions is significantly mediated by the perceived competence of the system.

Fig. 2 A simplified system and interaction overview for the adaptive robot condition. Used images in this figure were not modified. Wizard Hat (https://en.wikipedia.org/wiki/File:Wizard_Hat.jpg) by Rufus van helsing (https://en.wikipedia.org/wiki/User:Rufus_van_helsing) is licensed under CC BY-SA 3.0 (<https://creativecommons.org/licenses/by-sa/3.0/>), no changes made

Additionally, low trust is often associated with the misuse or disuse of an autonomous robot [2]. Previous works hypothesized that if people do not trust a robot, they stop using it. This trust in the competence of an interaction partner to achieve the desired goal is also highly critical between a client and a therapist [19]. Perceived higher competence increases the trust in the relationship to achieve a common goal. Thus, if people do not feel that the therapist has the competence to accomplish a common goal, they do not trust the therapist, do not build rapport, and are more likely to stop the therapy or intervention.

Thus, we draw our final hypothesis for this work:

Hypothesis 4 An *adaptive* robot increases the participant's motivation to engage in a second interaction compared to an *adaptable* robot.

To investigate these hypothesis, we present in the following a system and study design that incorporates two different adaptation strategies in an exercising scenario.

3 System Design

Figure 2 shows a high-level view of the system and interaction flow. The system consists of different components that communicate in a distributed system. The system composition includes a database of different exercises for Nao; a session controller, monitoring the exercises of the user and executing the robot's behavior; a simple computer vision system using a 3D depth sensor to analyze the skeleton of the user; a position controller for the robot as well as a preference learning algorithm. The system and decision components are implemented using the framework presented in [46] and are not further detailed in this manuscript.

3.1 Exercise Database

As previously found, exercising preference is unique to each person [38]. Thus, for the aim of this study, we developed a system that provides a variety of different exercises. We have chosen 25 exercises in total from 5 different categories: strength, stretch, cardio, taichi, and meditation. This set of activities tackles one of the open issues of SARs for exercising tasks. Previous work often looked at a single type of exercises like arm movements [11,12,16]. The approach of using a spectrum of different physical activities might show that people can perform various exercises together with a robot.

Table 1 presents the list of the chosen exercises. They have been selected based on a variety of criteria: (a) the possibility to animate and execute them on Nao (i.e., Nao cannot jump), (b) the difficulty that users can perform them (i.e., exercises should not be too challenging for the participants), (c) the exercises should target the full embodiment of the robot (i.e., laying down, balancing, standing).

Moreover, we limited the set of exercises to five categories and five exercises per class due to two considerations. First, participants in our study are asked to do at least 14 exercises. Thus, we chose five categories to make sure that the user is presented at least once every combination of exercising categories, which will be important for the preference learning approach that relies on the comparison between exercising

Table 1 Names of the exercises used for the presented study. The exercises were selected to represent a variety of exercises targeting possible preferences

Strength	Stretch	Cardio	Meditation	Taiji Drills
Push up	Neck	Jumping Jacks	The boat	Golden rooster
Squats	Triceps	Front Lunge	9 breathes	Rainbow
Crunches	Hip	Side Lunge	Relaxation	Punch
Superman	Quadriceps	Boxing	Inner light	Parting kick
Bridge	Side	Mnt. Climbers	Piece sign	Lifting water

Fig. 3 Robot and user performing the Taiji drill *Parting kick* together

categories.² Second, we chose five exercises per group so that the users eventually try out some other categories after the exercises start repeating.

All of them have been animated on Nao using Choregraphe [15,34]. Figure 3 shows an example of a user exercising together with the robot.

3.2 Preference Learning Framework

Preference learning is a subfield of machine learning that aims to learn predictive models from previously observed information (i.e., preference information) [14]. In supervised learning, a data set of labeled items with preference information is used to predict preferences for new items or all the other items from a data set. In general, the task for preference learning is concerned with the problem of learning to rank.

There are many different approaches to preference learning. It can be solved using supervised learning, unsupervised learning, and also reinforcement learning. Since there exists no particular data set we could use for supervised or unsupervised learning, it is challenging to build a model that can predict preferences from previously observed information. Therefore, we are focusing on how the system can learn an initial preference relation for a given item set without any prior knowledge (i.e., the cold start problem). Thus, we are trying to solve the preference learning problem using online methods for the Multi-Armed Bandit problem or, more precisely, Dueling Bandit algorithms [56].

The dueling bandit problem consists of K ($K \geq 2$) arms, where at each time step $t > 0$ a pair of arms $(\alpha_t^{(1)}, \alpha_t^{(2)})$ is drawn and presented to a user. A noisy comparison result w_t is obtained, where $w_t = 1$ if a user prefers $\alpha_t^{(1)}$ to $\alpha_t^{(2)}$, and $w_t = 2$ otherwise. The distribution of the outcomes is presented by a preference matrix $P = [p_{ij}]_{K \times K}$, where p_{ij}

is the probability that a user prefers arm i over arm j (e.g., $p_{ij} = P\{i > j\}, i, j = 1, 2, \dots, K$).

The goal of the preference learning task is, given a set of different actions (e.g., different sport categories), to find the user's preference order for these categories by providing the user two α_i and α_j and update the user preferences based on the selection of the preference between $\alpha_i > \alpha_j$ or $\alpha_i < \alpha_j$.

Thus, the challenge is to find the user's preference by running an algorithm that balances the exploration (gaining new information) and the exploitation (utilizing the obtained information). In this work, we are using the DTS algorithm presented in [55]. Since there are several implementations to solve the dueling bandit problem, we need to answer the question of why we have chosen this specific kind of algorithm.

Two reasons mainly drive this decision, the state of the art algorithms at the time of this study were DTS, RMED and its successor ECW-RMED [23,24]. Both perform reasonably well regarding their asymptotic behavior. However, currently, we are interested in the initial phase and not interested in the long-term run of these algorithms. If one takes a look at the first steps of these algorithms, one can see a significant difference between them that likely influences the HRI experience. RMED and ECW-RMED both have an initial phase where all possible pairs are repeatedly drawn for some time (see Algorithm 1 [23,24]). From an algorithmic perspective this is reasonable, but looking at it from the viewpoint of the interaction, this would lead to systematic comparisons that could result in boredom and even annoyance when the interaction partner is seemingly interrogating the user for her/his preferences. Thus, we assume that the DTS algorithm is more useful for HRI (especially for the initial contact between the trainee and the robot coach) because it does not rely on a systematic comparison of all possible pairs.

In previous research, we both verified the applicability of this preference learning approach in an ad-hoc preference learning scenario and the influence of the robot's embodiment on the satisfaction of the preference learning results, which is not existent [45]. This assures that the embodiment does not influence the user's acceptance of the preference ranking, which is an essential prerequisite for this study. Additionally, we validated that the used algorithm performs significantly better compared to a randomly selecting preference ranking.

4 Study Design

We conducted a study with a between-subject design (*adaptive* robot vs. *adaptable* robot) where participants were randomly assigned to one of two conditions.

² Using five exercising categories results in $\binom{5}{2} = 10$ possible comparisons. Adding another category would result in $\binom{6}{2} = 15$ possible comparisons, which would increase the total experiment time.

4.1 Conditions

In both conditions, the system waits for a user to be present in the room. Depending on the distance, it asks the participant to come closer. The system introduces itself to the user, explains its behavior and asks whether the user wants to start the exercising program.

Adaptive The robot in the adaptivity condition used the algorithm described in [55]. During the introduction phase, the system explains to the user that it will do different exercises together with the user and will ask for preference feedback relating to the various exercises. At each time step, the system selects two practices based on the preference learning algorithm and executes them consecutively with the user. Afterward, the system queries the user for a preference statement. The robot acknowledges the decision by repeating the chosen exercise. The preference learning algorithm updates the user's preference database and selects the next activities based on the current user preference. This behavior happens for 14 exercises (or seven iterations of the algorithm). As an additional measurement for motivation, the system asks after the 14 workouts whether the user wants to continue exercising for two more exercises or quit the experiment. However, only one participant did not want to do two more activities. After the two additional exercises, the robot finishes the interaction. It states the user's learned preferences based on the user's feedback in ascending order, and thanks for the participation. We limited the voluntary workouts to two, due to battery concerns and overheating of joints.

Adaptable The robot in the adaptability condition did not use any preference learning algorithm and did not autonomously select the next exercises. In the introduction phase, the system explains that it offers different exercises they can do together. Before each activity, the robot verbally lists the possible exercising categories in a randomized order, and the user can choose the exercise category she or he wants to experience. Thus, the user was in control of the exercise session and could choose the exercise category she or he prefers. Also, in this condition, the human and robot did 14 exercises together, and the robot asked whether the user wants to do two additional exercises. Additionally, in this condition, only one participant did not want to do two more practices.

4.2 System and Interaction Flow

The primary interaction flow for the conditions during the exercises is as follows: based on the current user's preference database, the algorithm selects two exercises (adaptable condition: the user selects the next exercise category), then the session manager runs these exercises sequentially. During the activities, the session manager receives user skeleton information from a depth sensor and classifies if the user is doing

the practices, by comparing the joint angles with the joint angle configuration for the specific exercises. We divided activities into crucial chunks (e.g., going up and down for squats and tracking the bend of the knees). The robot starts the movement and waits for the user to follow (i.e., for squats, the robot goes into the squat position). We only check whether there is a change in the joint angles and not whether it is the correct exercises. In case the user is not doing the activities, the system will run into a timeout after one second and continues with the next step of the exercises. At the start of a new task, the user usually first looks how to do the exercises and then joins the robot in doing the exercise. No participant refused to do an activity or did the wrong training. For tasks on the ground, the skeleton tracking does not work correctly, and the system is just following the exercise scripts. However, the instructions of the robot only start when the participant is on the ground. The interested reader can find more details on the used system, classification, and exercise pattern modeling in our previous publication [46].

4.3 Wizard of Oz Strategy

In the beginning, we used the internal speech recognition of Nao. However, prototype experiments showed that speech recognition capabilities are below an acceptable recognition rate, therefore we manually inserted the user's speech input using a Wizard of Oz (WoZ) style. The wizard listens to the user's feedback on the exercising preference or selected exercise category using an installed microphone in the experimental room and forwards the user's response as fake speech recognition results using the middleware to the session manager.

Additionally, when Nao performs the exercises, it moves away from the initial position. We have implemented a simple marker-based localization strategy. However, the duration for localization and positioning creates an unsatisfying interaction experience. Since this extended time is a significant disturbance for the HRI experience, we also have implemented a WoZ position controller to move the robot to the correct position manually after each exercise.

This approach ensures that speech recognition and position work reliable and do not influence the user's trust in the system. The general instructions for the wizard were to type in the user's speech for the exercise selection and position the robot to face the human. Additionally, in cases when the robot falls during an exercise, the in-built position sensor detects this, and an automatic stand up behavior is triggered.

4.4 Participants

Due to the expensiveness of this experiment in terms of time costs (i.e., 2 h per subject, breaking robots), we limited our sampling to 20 people per group. However, due to prior test-

ing and experiments, we are expecting a large effect size for our hypothesis relevant measurements.

The sampled participants ($N = 40$; average age $M = 26.02$, $SD = 5.48$, 13 female and 7 male in the adaptivity condition; 12 female and 8 male in the adaptability condition) were mostly university students that were acquired by information on the campus and social media. The majority of the participants were naive robot user and had no background in computer engineering or programming.

4.5 Procedure

Participants arrived at the lab individually. First, they gave informed consent. Then, the experimenter led the participants to a room where they could change their clothes. Later, the experimenter told the participant to enter the lab and follow the instructions of the system. Until this point, the participants did not know that they will be interacting with a robotic system. We neglected this prior information to not bias the participants or raise false beliefs. Then the participants entered the lab without the experimenter. The interaction took approximately 50 to 60 min, and the experimenter monitored the experiment from a control room. After the interaction finished, participants had to answer a questionnaire and had a voluntary post-study interview. Finally, they were debriefed and received 8 Euros for their participation. The ethical committee of our university approved the procedure.

4.6 Measurements

We use subjective responses to questionnaires to test our hypotheses and evaluate the performance of the online learned preference ranking using the position error between the learned ranking and the ground truth of the user. The following subsections explain these different measurements.

4.6.1 Hypotheses Related Measurements

In this study, we are investigating whether different personalization methods change a user's subjective perception of the robot, the alliance to it and motivation to interact with the system. The following measures were used in this study to find evidence for our presented hypotheses. We used Cronbach's α as a measure for the internal consistency of the scales and excluded scales that are below .5 for our report [8].

Negative Attitudes Towards Robots Attitudes towards robots were measured using the Negative Attitudes Towards Robots (NARS), $\alpha = .8$, on a five-point Likert scale [33]. Negative attitudes towards robots could be a confounding factor explaining results obtained on the perception of the robot.

Physical Activity Enjoyment Participants rated their physical training enjoyment using the Physical Activity Enjoyment Scale (PAES), $\alpha = .91$ [21]. The average overall item responses calculate the overall enjoyment score.

System Usability System's usability was measured by the System Usability Scale (SUS), $\alpha = .84$, with ten items on a 5-point Likert [4].

Team Perception We measured the user's perceived team perception using scales from [32]. These scales measure the general team perception ($\alpha = .38$), the openness to suggestions from the team member ($\alpha = .94$) and the perceived cooperation ($\alpha = .39$). All scales were on a 5 point-based Likert-scale.

Perception of the Partner Participants were asked to rate the perception of the robot on the Robot Social Attribute Scale (RoSAS). This scale includes the perceived warmth ($\alpha = .85$), competence ($\alpha = .77$) and discomfort ($\alpha = .76$) on a 9 point-based Likert-scale [7].

Motivation To have an additional measure to see whether people are interested in exercising a second time with the robot, we let the participants opt-in for voluntarily exercising with the robot again without monetary compensation. Participants were asked at the end of the questionnaire to enter their email address if they want to exercise again in the following week.

Trust in Alliance Finally, we used the Working Alliance Inventory (WAI), $\alpha = .91$, as a measure commonly used in helping alliances to assess trust and belief in a common goal of helping that a therapist, clinician or coach has for another [19]. This measure has recently been used in HCI and HRI studies for assessing the alliance and trust between the human and a SAR [3,22]. We adopted it for our use case (e.g., 'What I did in today's session, gives me a new view on my exercising preferences', 'Nao and I have worked together on our common goals for this session').

4.6.2 Preference Ranking Measurements

We use two ranking error functions to evaluate the quality of the preference ranking: D_{PE} which is the position error distance and D_{DR} which is the discounted error. Given a set of items $X = x_1, \dots, x_c$ to rank and r as the user's target preference ranking and $\hat{r}$ as the learned preference ranking. Both r and $\hat{r}$ are functions from $X \rightarrow N$ which return the rank of an item x . The position error is defined as follows

$$D_{PE}(r, \hat{r}) = \hat{r}(\operatorname{argmin}_{x \in X} r(x)) - 1 \quad (1)$$

The idea of this distance measure is that we want the target item (i.e. the highest ranked item from r) to appear as high as possible in the learned preference ranking $\hat{r}$. Thus, this distance gives the number of wrong items that are predicted

Fig. 4 Box plot showing the user ratings for the NEO-FFI personality test

before the target item. The discounted error is defined as follows

$$D_{DR}(r, \hat{r}) = \sum_{i=1}^c w_i \cdot d_{x_i}(\hat{r}, r) \quad (2)$$

where $w_i = \frac{1}{\log(r(x_i)+1)}$. This distance measure gives higher ranked items from r a higher weight for the distance error d_{x_i} between the rankings, where $d_{x_i}(\hat{r}, r)$ is the position difference between the the learned preference $\hat{r}$ and the true preference r .

In other words, having a correct ordering of the high ranked values from r is more important than of the low ranked items of r .

Since the goal of this study is to learn the user's most preferred exercise during the exploration phase for a cold start problem, we consider the position error D_{PE} as the most critical measurement. In other words, the goal is to rank the most preferred item as high as possible on the learned ranking after the exploration phase. Therefore, the exact ranking of the least preferred items is not as crucial as getting the most favorite item correct.

5 Results

In this section, we present the results from our quantitative survey evaluation (see Sect. 5.1) on the participant's subjective expression regarding the interaction experience using the above-described measures. Additionally, we show the results from the preference learning algorithm to verify the applicability of DTS in out-of-the-box personalization scenarios (see Sect. 5.2). Finally, we summarize the qualitative results from semi-structured post-study interviews which highlight participant's experience of the motivational capability of robotic exercising companions and their strategies to interact with the distinct personalization methods (see Sect. 5.3).

The quantitative data were analyzed using the statistical computing language R [36]. We analysed the data for normality assumptions and used Welch's two-sample t-test if the data meet the criteria and Wilcoxon rank sum test respectively

[53,54]. To increase reproducible science, we published the data and scripts for the analysis on Github.³

5.1 Quantitative Results

Manipulation Check The data was checked for differences in possible confounding factors such as the participant's previous experience with technology, their average weekly exercising activity, personality, physical activity enjoyment and the attitudes towards robots. We found no significant difference for these variables. Previous experience ($W = 146.5$, $p = .15$, $r = .22$), exercising activity ($W = 218$, $p = .61$, $r = -0.08$), PAES ($W = 164$, $p = .34$, $r = -.96$), as well as NARS ($t(37.7) = 1.78$, $p = .08$, $d = .56$) and the rating for the personality (all $p > .5$, see Fig. 4) were not significantly different between the conditions. Thus, the manipulation seems to be successful.

The general hypothesis unrelated measurements show that participants did not evaluate the usability of the systems significantly different, $t(35.56) = .95$, $p = .35$, $d = .30$. There was also no significant different evaluation regarding the openness to follow the system's suggestions, $W = 204$, $p = .92$, $r = -.10$ (see Fig. 5). Also, participants did not feel significantly more discomfort between the conditions, ($W = 210$, $p = .80$, $r = -.26$ (see Fig. 6).

However, the system was perceived as more warm measured by the responses on the RoSAS scale, $t(36.23) = -2.47$, $p = .02$, $d = .78$. The *adaptive* system is perceived as warmer ($M = 4.08$, $SD = 1.62$) than the *adaptable* system ($M = 2.93$, $SD = 1.29$).

In the following paragraphs, we present the results for our hypotheses.

Hypothesis 1 We hypothesized that the competence is perceived as higher in the *adaptive* condition compared to the *adaptable* condition. A Welch two-sample t test confirms this hypothesis and shows a significant difference between the conditions, $t(34.55) = -2.49$, $p = .02$, $d = .79$. The *adaptive* system is indeed perceived as more compe-

³ <https://github.com/sebschne/ijsr2019/>

Fig. 5 Box plot showing the user ratings for perceived cooperation, system usability, physical activity enjoyment, and working alliance

Fig. 6 Boxplot showing the user ratings for the RoSAS

tent ($M = 6.55$, $SD = 1.67$) than the *adaptable* system ($M = 5.4$, $SD = 1.2$).

Hypothesis 2 We also hypothesized that the user's trust in the HRI is higher in the adaptive condition. The results of the WAI are depicted in Fig. 5. A Welch two-sample t test revealed significant difference between the conditions, $t(36.05) = -3.17$, $p = .003$, $d = 1.00$. The *adaptive* system has been rated significantly higher on the alliance inventory ($M = 2.8$, $SD = .93$) than the *adaptable* system ($M = 1.99$, $SD = .76$). This confirms our hypothesis H2.

Hypothesis 3 To assess whether the condition's effect on overall alliance was statistically mediated by perceived competence, we used non-parametric bootstrapping method based on the method from [35] and coded condition as adaptable = 0, adaptive = 1. We tested assumptions for a mediation analysis using the *gvlma* package and used the *mediation* package to do the analysis. This analysis confirmed that perceived competence statistically mediated the

Fig. 7 Standardized regression coefficients for the alliance between conditions and user's alliance with the robot as mediated by the user's perceived competence of the robot. The standardized regression coefficient between the conditions and the WAI, controlling for perceived competence, is in parentheses

Table 2 Participants count to voluntarily repeat the interaction

	Yes	No
Adaptive	8	12
Adaptable	2	18

alliance between *adaptive* condition and overall trust in the robot ($ACME = .48$, $p < .05$, 95% CI = .1 to .91; 10,000 resamples; see Fig. 7) with no direct effect of autonomy of the system ($ADE = .31$, $p = .09$) and significant total effect ($p < .001$).

Hypothesis 4 Hypothesis 4 states that participants in the adaptive condition would voluntarily exercise a second time with the system compared to the adaptable system. The ratio for participant's wish to voluntarily repeat the interaction is depicted in Table 2. A Fisher's Exact Test showed that participants did not significantly opted more often to exercise again with the robot ($p = .06$, FET, 95% CI = 0.01 to 1.08, odds ratio = .17).

Thus, our Hypothesis 4 is not supported, or just marginally, supported.

5.2 Preference Learning Results

The results for the online learned exercising preferences by the system are depicted in Fig. 8. The plot shows for each participant in the *adaptive* condition the user's own *ranked* preference set and the *learned* exercising preference from the system during the interaction. Figure 9 shows the measured differences between the online learned rankings and true rankings in a box plot. We show two different measurements to compare item rankings. However, the essential measurement for our evaluation is how good the system could identify the most preferred exercise during the exploration phase. The position distance error D_{PE} captures this the best because we want to have the user's most favorite exercising category on the highest position in the online learned ranking. The position distance shows the difference between

Fig. 8 Individually ranked ground truth and online learned exercising preference for the different exercising categories

Fig. 9 Box plot showing the position and discounted distances between the online learned and user ranked exercising preference ground truth

the ranked position of the user's most preferred item to the position where the online learner has ranked it during our experiment. The median for this ranking is 0, presenting evidence that for most users, the system was able to identify the user's most preferred exercise after just seven comparisons. If we consider the difference in ranking errors as a sensitivity threshold, we see that in 65% of the cases the system made no errors, in 70% cases the system made $D_{PE} \leq 1$ errors, and in 85% the system made $D_{PE} \leq 2$. The discounted median error $D_{DR} = 3.45$, which is not as straightforward to compare as the D_{PE} . However, the median is just slightly higher than ranking the most preferred item on the second rank (i.e., $D_{DR} = 3.32$).

This result is promising, presenting evidence that these kinds of learning algorithms seem to be suitable to decrease the gap between a user's preference and the system's behavior. If we take a $D_{PE} \geq 3$ as a threshold for classifying the learned ranking as failed, we see that for participants 2, 9, and 13, the learned ranking was the worst among all participants.

However, it never ranked the user's most preferred item in the last position. There are different reasons why the system did not learn the ranking correctly for these users. For participant 13, the system explored the most preferred exercise (i.e., strength) only one time. Thus the algorithm could not learn the preference for this user. For participant 9, the algorithm worked fine. The learned preferences are in line with the feedback of the user during the interaction. However, participant 9 gave inconsistent input to the system with regards to the reported ground truth after the experiment. The reason why the system could not learn the preferences correctly for participant 2 is due to errors in the speech recognition pipeline when inserting the user's preference. The other important case is when the system learned the least preferred item in the first position. On the first sight of Fig. 8, the ranking for participant 4 shows that the least favorite exercise of this participant is ranked in the first position by the learner. However, the $D_{PE} = 1$ because the most preferred category of participant 4 follows the least favorite one on rank two. This result is due to ties in the ranking because the user had no clear preference towards an exercise category. In this case, the number of iterations was too few to learn the real user preference.

5.3 Qualitative Results

To understand whether participants felt motivated by the robot and how they experienced the personalization mechanism, we conducted semi-structured post-study interviews. After participants finished the questionnaire, we asked them whether they would like to answer some interview questions. Most of the participants gave, at least, some short responses. We asked the participants the following questions

Table 3 User responses on the question *why they felt motivated by the system*.

ID	Response	Reason
PA15	“It motivated me to try out new things”	Novelty
PA17	“I think it motivated me, because I don’t like to exercise alone ”	Companionship
PA19	“It would absolutely motivate me. I exercising alone and wouldn’t have to go out and the robot does not evaluate me”	Not judging
PB06	“It would motivate a lot of people. For me, I would have to decide on my own which exercises we’re doing”	Self-determination
PB07	“I did not know taiji, so it motivated me to try out new things”	Novelty
PB09	“The robot was sweet , it was enjoyable”	Appearance
PB10	“It’s nice when somebody is around who shows you the exercises and gives you structure. Especially, when you don’t have much time”	Companionship
PB12	“It would motivate me when it is better developed. Currently, I just exercise with videos, it could replace exercising videos when it is more sophisticated. I would prefer it to a human partner or coach. The robot keeps a distance and does not judge me , it could feel more uncomfortable with a human partner”	Not judging
PB13	“the appearance of the robot motivated me to follow its instructions and trust it. If a smartphone would ask me to do some exercises, I would just swipe them away, but a robot motivates me to try out new exercises”	Appearance, embodiment, novelty
PB16	“It was more fun, because I was not alone ”	Companionship
PB18	“I would use it, because I think it’s cool. It would be a nice feature, if it could track my performance. I had bad experiences with fitness tracking devices, but a robot could be a companion for everything”	Quantifying, companionship
PB19	“Yes it helped me, because I am not easily intrinsically motivated”	Extrinsic factor

We highlight keywords in boldface that give a hint to different categories of reasons why people think they felt motivated

1. whether and why they felt motivated by the system (“Did you feel motivated by the system?”, “Can you give reason why you felt motivated by the system?”)
2. in the *adaptable* condition, we asked participants which strategy they used to select the exercises (“On which basis have you selected the exercising categories?”)
3. in the *adaptive* condition, we asked participants on which criteria they made their preference selection (“How have you decided which exercise you preferred?”)
4. how much money they would spend for such a system. (“How much would you pay for such a system?”)

Additionally, we asked participants whether they aligned their behavior (i.e., exercising execution) to the robot’s behavior and which modality was more important to them (the motion of the robot or its verbal instructions). Using an in-depth interaction video analysis, this should give insights on whether and how people synchronize their behavior with robots. Because the behavior of the robot was the same between conditions, the analysis of this question is not essential for the research question of this paper. Thus, we restrict the presentation of the results only to summarize the main evident insights regarding the exercising motivation and selection strategy, which are essential aspects of this evaluation. Therefore, the following paragraphs only show the responses for the motivation and the selection strate-

gies. The number of participants that answered a question varies because, at some points, participants finished the interview due to time constraints. Participants from the *adaptable* condition have the identifier *PAxx* and participants from the *adaptive* condition *PBxx*

Exercising Motivation Table 3 summarizes some testimonies of the participants. We highlight distinguishable keywords that provide insights into the personal reasons why users felt motivated by the system. The responses provide a multitudinous reflection on the internal heuristic to evaluate whether and why they felt motivated. The responses show that there are vast differences in the evaluation criteria for each person. Participants felt motivated by the appearance of the robot, the novelty of guiding them through new exercises, the fact that they do not feel evaluated by a robotic exercising partner, the companionship the robot can provide, or by the possibility that the robot can quantify their training progress. Participants who stated that they did not feel motivated by the systems gave recommendations and use case suggestions for the system. The suggested use case for the system would be as a reminder system, as a partner for rehabilitative exercises or as a partner for people that just started exercising. As interactive suggestions, participants proposed that the robot should be faster and emulate emotions.

From the total number of participants that took part in the interview (*adaptable*: 11, *adaptive*: 16), six participants in

the *adaptable* condition said that the system motivates them to exercise (e.g., *PA14*: Motivating, very nice during static workouts but not so good for cardio). In contrast, five participants stated that the system was not motivating or that they are intrinsically motivated and would not need it. Still, they would appreciate assistance when they are injured (e.g., *PA12*: “I don’t know. I am intrinsically motivated. For my daily life, I would not use it, perhaps if I am injured as a rehabilitation tool.”). Regarding the *adaptive* condition, eleven participants stated that the system would motivate them. At the same time, five said that they did not feel motivated by it (e.g., *PB04*: “Felt motivated but not due to the robot, because the robot does not feel any exhaustion”, *PB08*: “It misses emotion, could be more motivating. I did not take its messages seriously because I know that it is a machine and not a human. I also did not know whether it perceived me”).

Exercise Selection Strategy We asked the participants in both conditions, whether they used any strategy to select the next exercise and on which basis they chose their exercising preferences. 17 participants of the *adaptable* condition answered this question and 15 participants of the *adaptive* condition.

In the *adaptable* condition, ten participants said that they tried to select everything once to see what the system has to offer (*PA08*: “No strategy, I tried to select everything once”). Seven participants selected the exercises based on their actual exercising preferences (*PA14*: “I selected everything based on my preference. Therefore, I did not select cardio or relaxation exercises”).

In the *adaptive* condition, ten participants selected the exercises based on their current enjoyment of the task (e.g., *PB9*: “I thought about what was more fun for me and picked the exercise accordingly”), and three selected the activities based on their actual exercising preference (e.g., *PB18*: “I chose the exercises based on my choice”).

It is interesting to note that the interview responses show that the different types of personalization strategies lead to different approaches to select the exercises. While in both conditions, there were participants that chose activities based on their actual preference, participants that did not use a favorite based selection approach used mainly two different kinds of selection criteria. Participants in the *adaptable* condition more often stated that they used their curiosity as selection criteria (e.g., “I wanted to see what the system has to offer”). In contrast, participants in the *adaptive* condition did not state that they used curiosity criteria. Instead, they used enjoyment as the salient criteria for stating their exercising preference.

This presents evidence that people use different kinds of qualitative criteria to maximize their interactive experience. In *adaptable* condition, participants concentrated more on the novelty as an evaluation criterion, while in the other condition, participants used their enjoyment as an evaluation

Fig. 10 Box plots showing amount of money participants are willing to pay for the system

criterion. Though maximizing novelty can also be considered as an interaction enjoyment criterion, it is sufficiently different from the actual enjoyment evaluation. Maximizing novelty tries to optimize expected interaction experience in the future, while the other approach evaluates the currently perceived interaction experience.

Costs We asked participants how much they would pay for such a system. Their responses were ambiguous, and we found a big discrepancy in how much they are willing to pay for such a system and how much they think it would cost. While many knew that robots are expensive, the willingness to pay this amount was low. Thus, we concentrate only on the expenses they are willing to pay. The median amount of money participants in the adaptive condition would pay is 250€, while in the adaptable condition, participants would pay 300€ (see Fig. 10). Even though the median is higher for the adaptable condition, the variance is greater for the adaptive condition. Since the results are not significantly different ($t(24.53) = .3, p = .72$), we assume that the personalization strategy is currently not a salient feature to determine the value of the system. Participants rather stated that if the system is more capable of doing the exercises, they would be willing to pay more for it.

6 Discussion

This work investigated how a system’s type of personalization mechanism based on different LoA alters the user’s perception of it. It presented a study to investigate the effects of interacting with an *adaptable* or *adaptive* robot on the perceived alliance with the system and the perceived competence of it depending on different personalization strategies. Thus, it closes a gap in the research literature on the effects of varying personalization methods in HRI. The robot in our study was either indirectly controlled by a user’s preference feedback or directly controlled by the user. In the case of the

adaptive condition, we used a preference-learning method based on dueling bandits. Thus, this work also presents evidence that these kinds of algorithms are suitable for personalizing HRI experience. The results from our evaluation show that the system is, most of the time, able to learn a user preference during a short exploration phase.

We hypothesized that different LoA alters the perceived competence of the robot and alliance with it. The results present evidence that users view the robot as more competent, which is supported by a significant difference between the conditions on the RoSAS subscale. This evidence supports Hypothesis 1: An *adaptive* robot is perceived as more competent than an *adaptable* robot. Further, it confirms theoretical investigations on the LoA and shows that a system with a higher degree of automation is indeed perceived as more competent. Therefore, we could find a hint that unexpected behaviors result in different evaluations of the robot, as postulated in Epley et al. [10] theory of anthropomorphism. Expectance violations regarding the non-human agent's behavior, let people rethink their mental model of the agent's behavior or mental state and increase anthropomorphic thoughts. Whether it means that higher degrees of perceived competence also result in higher perceived degrees of anthropomorphism remains uncertain. Though, it is reasonable to argue that also, higher competence reflects a higher degree of human-likeness evaluation. Still, it would require further evaluation to discern the effects of competence and anthropomorphism.

The results on the perceived alliance with the robot also support the Hypothesis 2: Participants had more trust in the alliance with the *adaptive* robot. This result supports the hypothesis that Rau et al. [37] had, but could not find evidence to support it. Participants trusted the robot more if it was more autonomous than less autonomous. This result seems counter-intuitive. Why would users trust a robot more when they have less control over the robot? It might be that participants also could have felt overwhelmed by the exercising possibilities they had with the system. Thus, participants might have felt less burden to structure the interaction, because the *adaptive* system made the critical decision. One could say, from a system theory perspective, that the users try to reduce their uncertainty when interacting with the system [27]. Further, the mediation model supports Hypothesis 3 and provides insights on why the different conditions affected the perceived alliance. Different LoA influenced the perceived competence of the system, which in turn increased the alliance to it. Other researchers showed that anthropomorphism alters trust in an autonomous vehicle [52]. Higher anthropomorphism leads to higher confidence in the car. However, these authors have not measured the perceived competence as an independent mediator. Thus it remains an open question whether the manipulation of the anthro-

pomorphism alters the perceived competence of the system and therefore changes the associated alliance.

Finally, we could find partial evidence for Hypothesis 4. Participants in the *adaptive* condition opted more often to exercise a second time voluntarily. This result is probably due to the interest in a system that tries to personalize the interaction by itself. It might raise curiosity, and participants are interested to see what other exercises the system can offer or whether the system can effectively learn the user preference. However, this result is only marginally significant after applying a continuity correction. A bigger sample is needed to be sure whether this is a substantial effect.

Our qualitative results also reveal insights into the rationale for evaluating different personalization strategies. Using adaptive robots results in evaluation criteria based on the current state, rather than on the already established preference beliefs. This opens speculations whether an adaptive approach leads to reconsider one's opinions on exercising preference and thus helps to be open to trying new exercises that eventually better-fit one's personality. While most participants in the *adaptable* condition were using their beliefs for the exercising selection, participants in the *adaptive* condition focused more on which exercises they enjoyed. Having the results for the difference in perceived competence and trust in mind, this shows that participants are open for new suggestions and that they would trust the competence of the system to find the right exercise for them.

Regarding the preference learning results, we found that the system was able to learn online a user matching ranking after seven comparisons during the interaction in 65% of the cases. Relaxing the preference ranking error around one item, we see that the system can learn a good ranking in 70% of the cases. These results extend our previous experimental data of preference learning in an ad-hoc HRI scenario, that compares the algorithm to a random ranking [] by presenting evidence that these algorithms are suitable for online interactions. In light of the short interaction time, we hypothesize that giving the algorithm more rounds of exploration would enhance these results. However, we have only argued that the position error of the most preferred item is essential for the adaptation towards the user, and we have not investigated the effect of how a bad ranking on the lower preferred items might affect the interaction and exploration in the long run. Additionally, we only looked at a small subset of exercising categories, and perhaps the actual real exercising preference of the participants was not included in our experiment (e.g., ball sport, horse riding).

One limitation of the interpretation of the results above is the short interaction time during the study. The alliance is build up over repeated interactions between two people. Therefore, the results on the effects of partnership need to be interpreted with caution. Additionally, the scale used in this experiment is primarily designed for measuring the trust and

alliance in the client-therapist collaborations. Therefore, the results might be different if we have used a scale that is more focused on trust in the technical competence of the system. Still, the confidence in the relationship is an essential part of long-term HRI and especially for use cases where the human and the robot partner are working towards a long-term goal like increasing physical activity.

Moreover, we have not quantitatively assessed the quality of the preference learning over repeated interactions. Future work will look at long-term satisfaction regarding the learned preferences. A further limitation of our results is the sampled study population. We mainly tested healthy, young adults at university age. Thus, the results are not generalizable to a younger or older population, and not to people that require rehabilitation. Hence, future investigations also have to verify the applicability of the preference learning framework with focus user groups.

A confounding factor for our results is the significant difference in perceived warmth between the conditions. Since we have not assessed the individual's perception of warmth towards robots and had no pre-interaction evaluation, the difference could be due to differences across cohorts and not by the manipulation.

Lastly, one major limitation of our system is the use of a wizard. In our scenario, we used a wizard due to the limited speech recognition capabilities and to speed up the localization and position of the robot. As pointed out in previous research, natural language processing is a significant challenge in artificial intelligence and one of the main scenarios when using WoZ [39]. While there has been considerable progress in speech recognition, and large companies offer cloud-based services, such services can not be used due to data privacy concerns. Thus, technically, the true adaptivity of our system is limited by speech recognition capabilities. In our scenario, we restricted the possible verbal interaction between the human and the system. We faked the speech recognition input when the system was asking for the next exercising category or the preferred exercising type. This verbal response was feedback as a speech recognition result to the session controller. To get the human out-of-the-loop and increase the adaptivity of the system for real-world use cases, without requiring a human operator, one could consider using an additional interface (e.g., tablet, or smartphone) that prompts the user for feedback. It will be an open question to research how this approach will affect the user experience.

7 Conclusion

This work presented a study on different methods to personalize a SARs behavior towards a user's preference. The results of this study show that *adaptive* robots are perceived as more

competent and trustworthy than *adaptable* robots. Thus, the associated LoA indeed influences the interaction experience for the user. Further, it presents evidence that the perceived competence of the system significantly mediates the alliance with it. This mediation effect can be an essential aspect of long-term interaction with robots and needs in-depth investigations in long-term studies. The question remains whether an *adaptive* system can continuously present new and personalized behaviors so that the system will continue exciting to interact with over time. Moreover, regarding the question of how a future social robotic system can personalize its behavior for the user, our results show that the system can sufficiently identify the user's preference in a short amount of time using a DTS approach. Thus, we could show, as one of the first researchers to do this, that a qualitative comparative approach is suitable for online adaptation in HRI scenarios. Nevertheless, future research needs to verify the results in long-term investigations.

Acknowledgements Open Access funding provided by Projekt DEAL. This research/work was supported by the Cluster of Excellence Cognitive Interaction Technology 'CITEC' (EXC 277) at Bielefeld University, which is funded by the German Research Foundation (DFG).

Compliance with Ethical Standards

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Andrist S, Mutlu B, Tapus A (2015) Look like me: matching robot personality via gaze to increase motivation. In: Proceedings of the 33rd annual ACM conference on human factors in computing systems, ACM, pp 3603–3612
2. Beer JM, Fisk AD, Rogers WA (2014) Toward a framework for levels of robot autonomy in human–robot interaction. *J Hum Robot Interact* 3(2):74–99
3. Bickmore TW, Picard RW (2005) Establishing and maintaining long-term human–computer relationships. *ACM Trans Comput Hum Interact (TOCHI)* 12(2):293–327
4. Brooke J et al (1996) Sus—a quick and dirty usability scale. *Usability Eval Ind* 189(194):4–7

5. Burger JM (2013) Desire for control: personality, social and clinical perspectives. Springer, Berlin
6. Cakmak M, Chao C, Thomaz AL (2010) Designing interactions for robot active learners. *IEEE Trans Auton Ment Dev* 2(2):108–118
7. Carpinella CM, Wyman AB, Perez MA, Stroessner SJ (2017) The robotic social attributes scale (rosas): development and validation. In: *Proceedings of the 2017 ACM/IEEE international conference on human–robot interaction*, ACM, pp 254–262
8. Cronbach LJ (1951) Coefficient alpha and the internal structure of tests. *Psychometrika* 16(3):297–334
9. Endsley MR, Kiris EO (1995) The out-of-the-loop performance problem and level of control in automation. *Hum Factors* 37(2):381–394
10. Epley N, Waytz A, Cacioppo JT (2007) On seeing human: a three-factor theory of anthropomorphism. *Psychol Rev* 114(4):864
11. Eriksson J, Mataric MJ, Winstein CJ (2005) Hands-off assistive robotics for post-stroke arm rehabilitation. In: *9th international conference on rehabilitation robotics*, 2005. ICORR 2005. IEEE, pp 21–24
12. Fasola J, Mataric M (2013) A socially assistive robot exercise coach for the elderly. *J Hum Robot Interact* 2(2):3–32
13. Freedy A, DeVisser E, Weltman G, Coeyman N (2007) Measurement of trust in human–robot collaboration. In: *International symposium on collaborative technologies and systems (2007) CTS 2007*. IEEE, pp 106–114
14. Fürnkranz J, Hüllermeier E (2011) Preference learning: an introduction. Springer, Berlin, pp 1–17. https://doi.org/10.1007/978-3-642-14125-6_1
15. Gouaillier D, Hugel V, Blazevic P, Kilner C, Monceaux J, Lafourcade P, Marnier B, Serre J, Maisonnier B (2008) The NAO humanoid: a combination of performance and affordability. *CoRR abs/08073223*
16. Guneyasu A, Arnrich B (2017) Socially assistive child-robot interaction in physical exercise coaching. In: *2017 26th IEEE international symposium on robot and human interactive communication (RO-MAN)*, pp 670–675. <https://doi.org/10.1109/ROMAN.2017.8172375>
17. Hancock PA, Billings DR, Schaefer KE, Chen JY, De Visser EJ, Parasuraman R (2011) A meta-analysis of factors affecting trust in human–robot interaction. *Hum Factors* 53(5):517–527
18. Hemminghaus J, Kopp S (2017) Towards adaptive social behavior generation for assistive robots using reinforcement learning. In: *Proceedings of the 2017 ACM/IEEE international conference on human–robot interaction*, ACM, pp 332–340
19. Horvath AO, Greenberg LS (1989) Development and validation of the working alliance inventory. *J Couns Psychol* 36(2):223
20. Iolanda Leite CM, Paiva A (2013) Social robots for long-term interaction: a survey. *Int J Soc Robot* 5(2):291–308
21. Kendzierski D, DeCarlo KJ (1991) Physical activity enjoyment scale: two validation studies. *J Sport Exerc Psychol* 13(1):50–64
22. Kidd CD, Breazeal C (2008) Robots at home: understanding long-term human–robot interaction. In: *2008 IEEE/RSJ international conference on intelligent robots and systems*, pp 3230–3235. <https://doi.org/10.1109/IROS.2008.4651113>
23. Komiyama J, Honda J, Kashima H, Nakagawa H (2015) Regret lower bound and optimal algorithm in dueling bandit problem. In: *Conference on learning theory*, pp 1141–1154
24. Komiyama J, Honda J, Nakagawa H (2016) Copeland dueling bandit problem: regret lower bound, optimal algorithm, and computationally efficient algorithm. *arXiv preprint arXiv:1605.01677*
25. Leite I, Pereira A, Castellano G, Mascarenhas S, Martinho C, Paiva A (2011) Modelling empathy in social robotic companions. In: *International conference on user modeling, adaptation, and personalization*. Springer, pp 135–147
26. Leyzberg D, Spaulding S, Scassellati B (2014) Personalizing robot tutors to individuals' learning differences. In: *Proceedings of the 2014 ACM/IEEE international conference on human–robot interaction*. ACM, pp 423–430
27. Luhmann N (2000) Unsicherheitsabsorption, VS Verlag für Sozialwissenschaften, Wiesbaden, pp 183–221. <https://doi.org/10.1007/978-3-322-97093-0-6>
28. Madsen M, Gregor S (2000) Measuring human–computer trust. In: *11th Australasian conference on information systems*, Citeseer, vol 53, pp 6–8
29. Martelaro N, Nneji VC, Ju W, Hinds P (2016) Tell me more: designing HRI to encourage more trust, disclosure, and companionship. In: *The eleventh ACM/IEEE international conference on human robot interaction*, IEEE Press, Piscataway, HRI'16, pp 181–188. <http://dl.acm.org/citation.cfm?id=2906831.2906863>
30. McAllister DJ (1995) Affect-and cognition-based trust as foundations for interpersonal cooperation in organizations. *Acad Manag J* 38(1):24–59
31. Mitsunaga N, Smith C, Kanda T, Ishiguro H, Hagita N (2008) Adapting robot behavior for human–robot interaction. *IEEE Trans Robot* 24(4):911–916
32. Nass C, Fogg B, Moon Y (1996) Can computers be teammates? *Int J Hum Comput Stud* 45(6):669–678
33. Nomura T, Kanda T, Suzuki T (2006) Experimental investigation into influence of negative attitudes toward robots on human–robot interaction. *AI & Society* 20(2):138–150
34. Pot E, Monceaux J, Gelin R, Maisonnier B (2009) Choregraphe: a graphical tool for humanoid robot programming. In: *The 18th IEEE international symposium on robot and human interactive communication*, 2009. RO-MAN 2009. IEEE, pp 46–51
35. Preacher KJ, Hayes AF (2008) Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behav Res Methods* 40(3):879–891
36. R Core Team (2013) R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org/>
37. Rau PLP, Li Y, Liu J (2013) Effects of a social robot's autonomy and group orientation on human decision-making. *Adv Hum Comput Interact* 2013:11
38. Rhodes R, Smith N (2006) Personality correlates of physical activity: a review and meta-analysis. *Br J Sports Med* 40(12):958–965
39. Riek LD (2012) Wizard of oz studies in HRI: a systematic review and new reporting guidelines. *J Hum Robot Interact* 1(1):119–136
40. Ritschel H, Baur T, Andre E (2017) Adapting a robot's linguistic style based on socially-aware reinforcement learning. <https://doi.org/10.1109/ROMAN.2017.8172330>
41. Robinette P, Li W, Allen R, Howard AM, Wagner AR (2016) Overtrust of robots in emergency evacuation scenarios. In: *The eleventh ACM/IEEE international conference on human robot interaction*. IEEE Press, pp 101–108
42. Russell SJ, Norvig P (2016) Artificial intelligence: a modern approach. Pearson Education Limited, Malaysia
43. Salem M, Lakatos G, Amirabdollahian F, Dautenhahn K (2015) Would you trust a (faulty) robot?: effects of error, task type and personality on human–robot cooperation and trust. In: *Proceedings of the tenth annual ACM/IEEE international conference on human–robot interaction*. ACM, pp 141–148
44. Schneider S, Kummert F (2016) Exercising with a humanoid companion is more effective than exercising alone. In: *2016 IEEE-RAS 16th international conference on humanoid robots (humanoids)*. IEEE, pp 495–501
45. Schneider S, Kummert F (2017) Exploring embodiment and dueling bandits for preference adaptation in human–robot interaction. In: *Proceedings of the 26th IEEE international symposium on robot and human interactive communication*
46. Schneider S, Goerlich M, Kummert F (2017) A framework for designing socially assistive robot interactions. *Cognit Syst Res* 43:301–312

47. Schneider-Hufschmidt M, Malinowski U, Kuhme T (1993) Adaptive user interfaces: principles and practice. Elsevier Science Inc, New York
48. Sheridan TB, Verplank WL (1978) Human and computer control of undersea teleoperators. Tech. rep, Massachusetts Institute of Tech Cambridge Man-Machine Systems Lab
49. Tapus A, Țăpuș C, Mataric MJ (2008) User robot personality matching and assistive robot behavior adaptation for post-stroke rehabilitation therapy. *Intell Serv Robot* 1(2):169–183
50. Tsiakas K, Papakostas M, Chebaa B, Ebert D, Karkaletsis V, Makedon F (2016) An interactive learning and adaptation framework for adaptive robot assisted therapy. *PETRA* 10(1145/2910674):2935857
51. Wang N, Pynadath DV, Hill SG (2016) Trust calibration within a human–robot team: comparing automatically generated explanations. In: The eleventh ACM/IEEE international conference on human robot interaction. IEEE Press, Piscataway, NJ, USA, HRI '16, pp 109–116. <http://dl.acm.org/citation.cfm?id=2906831.2906852>
52. Waytz A, Heafner J, Epley N (2014) The mind in the machine: anthropomorphism increases trust in an autonomous vehicle. *J Exp Soc Psychol* 52:113–117
53. Welch BL (1947) The generalization of 'student's' problem when several different population variances are involved. *Biometrika* 34(1/2):28–35
54. Wilcoxon F (1945) Individual comparisons by ranking methods. *Biom Bull* 1(6):80–83
55. Wu H, Liu X (2016) Double thompson sampling for dueling bandits. In: Advances in neural information processing systems, pp 649–657
56. Yue Y, Broder J, Kleinberg R, Joachims T (2012) The k-armed dueling bandits problem. *J Comput Syst Sci* 78(5):1538–1556

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Sebastian Schneider received his B.Sc. in Bioinformatics and Genome Research, and his M.Sc. and Ph.D. in Intelligent Systems at Bielefeld University in 2009, 2011 and 2019, respectively. From 2011 to 2014, he worked at the Cognition and Robotics Lab (CoR-Lab) and Applied Informatics group at Bielefeld University. In 2014, Sebastian joined the CITEC Graduate School to work on his Ph.D. project on social robots for exercising applications. He was a visiting researcher at the Emergent Robotics Laboratory at Osaka University and the Cognitive Developmental Robotics Lab at the Center for Information and Neural Networks (CiNet) from November 2017 until January 2018. He defended his thesis in 2019 and continued to work as a postdoc at the research group Applied Informatics. From October 2019 until February 2020, he was working at the Human-Robot Interaction Lab at Kyoto University as a specially appointed researcher. His fields of research are Human-Robot Interaction (HRI) with a focus on Socially Assistive Robots (SAR), adaptive robots, and activity recognition.

Franz Kummert received the diploma and the Ph.D. (Dr.-Ing.) degree in computer science from the University of Erlangen-Nürnberg, Erlangen, Germany, in 1987 and 1991, respectively. In 1996 he received the *venia legendi* (Habilitation) in computer science from the Faculty of Technology of Bielefeld University, Germany. Since 2002 he is a apl. professor for pattern recognition. From 1987 to 1990 he worked at the research group for Pattern Recognition (Institut für Informatik, Mustererkennung) at the University of Erlangen-Nürnberg, Erlangen, Germany. Since 1991, he is with the research group Applied Informatics (Angewandte Informatik) at Bielefeld University, Germany. His fields of research are speech and image understanding and human-robot interaction. He has published various papers in these fields, and is author of a book on the control of a speech understanding system and on the automatic interpretation of speech and image signals. Dr. Kummert is Member of the Institute of Electrical and Electronics Engineers (IEEE).