



A Socially Assistive Robot Exercise Coach for the Elderly

Juan Fasola and Maja J Matarić
University of Southern California

We present a socially assistive robot (SAR) system designed to engage elderly users in physical exercise. We discuss the system approach, design methodology, and implementation details, which incorporate insights from psychology research on intrinsic motivation, and we present five clear design principles for SAR-based therapeutic interventions. We then describe the system evaluation, consisting of a multi-session user study with older adults ($n = 33$), to evaluate the effectiveness of our SAR exercise system and to investigate the role of embodiment by comparing user evaluations of similar physically and virtually embodied coaches. The results validate the system approach and effectiveness at motivating physical exercise in older adults according to a variety of user performance and outcomes measures. The results also show a clear preference by older adults for the physically embodied robot coach over the virtual coach in terms of enjoyableness, helpfulness, and social attraction, among other factors.

Keywords: Human-robot interaction, socially assistive robotics, exercise, elderly, intrinsic motivation, embodiment

1. Introduction

The growing population of aging adults is increasing the demand for healthcare services worldwide. By the year 2050, the number of people over the age of 85 will increase threefold (Centers for Disease Control and Prevention, 2003), while the shortfall of nurses and caregivers is already an issue (American Association of Colleges of Nursing, 2010; American Health Care Association, 2008; Buerhaus, 2008). Regular physical exercise has been shown to be effective at maintaining and improving the overall health of elderly individuals (Baum et al., 2003; Dawe & Moore-Orr, 1995; McMurdo & Rennie, 1993; Thomas & Hageman, 2003). Physical fitness is associated with higher functioning in the executive control processes (Colcombe & Kramer, 2003) and is correlated with less atrophy of frontal cortex regions (Colcombe et al., 2004) and with improved reaction times (Spirduso & Clifford, 1978). Social interaction, and specifically high perceived interpersonal social support, has also been shown to have a positive impact on general mental and physical wellbeing (Moak & Agrawal, 2010), in addition to reducing the likelihood of depression (George et al., 1989; Paykel, 1994; Stansfeld et al., 1997; Stice, Ragan, & Randall, 2004). Thus, the availability of physical exercise therapy, social interaction, and companionship will be critical for the growing elderly population; socially assistive robotics has the potential to help address this need.

A socially assistive robot (SAR) is a system that employs hands-off interaction strategies, including the use of speech, facial expressions, and communicative gestures, to provide assistance in accordance with a particular healthcare context. Previous SAR research by our group includes systems that were developed and tested with stroke patients (Matarić, Eriksson, Feil-Seifer, &

Authors retain copyright and grant the Journal of Human-Robot Interaction right of first publication with the work simultaneously licensed under a Creative Commons Attribution License that allows others to share the work with an acknowledgement of the work's authorship and initial publication in this journal.

Winstein, 2007; Tapus, Tapus, & Matarić, 2008), Alzheimer’s patients (Tapus, Tapus, & Matarić, 2009), and children with autism spectrum disorders (Feil-Seifer & Matarić, 2009, 2010), as well as healthy young adults (Fasola & Matarić, 2010a) and healthy elderly adults (Fasola & Matarić, 2010b, 2012). SAR systems equipped with motivational, social, and therapeutic capabilities have the potential to facilitate elderly individuals to live independently in their own homes, to enhance their quality of life, and to improve their overall health.

In this paper, we present the approach, design methodology, and implementation details of a novel socially assistive robot system that aims to motivate and engage elderly users in simple physical exercise. Our SAR system approach incorporates insights from psychology research into intrinsic motivation and contributes clear design principles developed to maximize the probability of success of the therapeutic intervention.

Furthermore, we present a multi-session user study conducted with older adults to evaluate the effectiveness of our robot exercise system across a variety of user performance and evaluation measures. The results of the study validate the system approach and its effectiveness in motivating physical exercise in older adults; the participants consistently engaged in physical exercise throughout the interaction sessions, rated the SAR system interaction highly in terms of enjoyableness and usefulness, and rated the robot coach highly in terms of helpfulness, social attraction, social presence, and companionship.

The user study also served to investigate the role of physical embodiment, a fundamental topic in human-robot interaction (HRI), by comparing the effectiveness of a physically embodied robot coach to that of a similar virtually embodied robot coach (a computer simulation of the same robot). Physically embodied agents appear to possess what Lee (2004) refers to as “social presence” to a greater extent than virtually embodied agents do. Social presence mediates how people respond to both embodied and disembodied agents and strongly influences the relative success of the social interaction. It is thus important to note that the embodiment type of a socially assistive agent can influence its effectiveness in social interaction, relationship building, gaining user acceptance and trust, and ultimately in achieving the desired health outcomes of therapeutic intervention. For these reasons, we explored the role of physical embodiment in our SAR exercise system for the elderly.

The results of the embodiment comparison show a strong preference among the participants for the physically embodied robot over the virtually embodied robot and demonstrate the positive effect that physical embodiment has on participant evaluations of both the interaction and the robot.

The rest of this paper is organized as follows: In Section 2 we discuss related work in the area of robotics for the elderly, social agent coaches, and embodiment effects. Section 3 presents our SAR system approach, design methodology, architecture, implementation details, and interaction scenario. Section 4 describes our system evaluation and embodiment comparison study design, and Section 5 presents the results of our user study conducted with older adults. We conclude the paper in Section 6 with a summary of the key research contributions of this work.

2. Related Work

2.1 Robots for the Elderly

The literature that addresses assistive robotics intended for and evaluated by the elderly is limited but growing. Representative work in the field includes robots that focus on providing assistance for functional needs, such as mobility aids and navigational guides. Dubowsky et al. (2000) developed a robotic cane/walker device designed to help individuals by functioning as a mobility aid that provides physical support for walking as well as guidance and health monitoring of a user’s basic vital signs. Montemerlo, Pineau, Thrun, and Varma (2002) designed and pilot-tested a robot that escorts elderly individuals in an assisted living facility, reminds them of their scheduled appointments, and provides informational content such as weather forecasts.

Researchers have also investigated the use of robots to help address the social and emotional needs of the elderly, including reducing depression and increasing social interaction with peers. Wada, Shibata, Saito, and Tanie (2002) studied the psychological effects of a seal robot, *Paro*, which was used to engage seniors at a day service center. The study found that *Paro*, always accompanied by a human handler, was able to consistently improve the moods of elderly participants who spent time petting and engaging with it over the course of a 6-week period. Kidd, Taggart, and Turkle (2006) used *Paro* in another study that found it to be useful as a catalyst for social interaction. They observed that seniors who participated with the robot in a group were more likely to interact socially with each other when the robot was present and powered on, rather than when it was powered off or absent.

Perhaps the robotic system for the elderly most related to our SAR exercise system is the work of Matsusaka, Fujii, Okano, and Hara (2009), who developed an exercise demonstrator robot, *TAIZO*, to aid human demonstrators during simple arm exercises with a training group. However, this robot was not autonomous, as it was controlled via key input or voice by the lead human demonstrator, and it did not have sensors by which to perceive the users; hence, it did not provide any real-time feedback, active guidance, or personalized training.

2.2 Social Agent Coaches

Social agents that aim to assist individuals in health-related tasks such as physical exercise have also been developed in the human-computer interaction (HCI) community. Bickmore and Picard (2005) developed a computer-based virtual relational agent that served as a daily exercise advisor by engaging the user in conversation and providing educational information about walking for exercise, asking about the user's daily activity levels, tracking user progress over time while giving feedback, and engaging the user in relational dialogue. Kidd and Breazeal (2008) developed a table-top robot to serve as a daily weight-loss advisor; it engaged users through a touch-screen interface, tracked user progress and the state of the user-robot relationship over time, and was tested in a six-week field study with participants at home. French, Tyamagundlu, Siewiorek, Smailagic, and Ding (2008) designed and explored the use of a virtual coach to assist manual wheelchair drivers by providing advice and guidance to help users avoid hazardous forms of locomotion.

These systems are similar to our SAR exercise system in the manner in which they provide feedback (from a social agent) and, with the exception of French's work, in the activity being monitored (physical exercise). However, our system is clearly distinct in that our robot agent not only provides active guidance, feedback, and task monitoring, but is also directly responsible for instructing and steering the task as well. Hence, our agent is both an administrator and an active participant in the health-related activity, resulting in a unique characteristic of the system: The social interaction between the robot and user is not only useful for maintaining user engagement and influencing intrinsic motivation, but is also necessary to achieving the physical exercise task.

2.3 The Effect of Embodiment

Previous studies investigating the role of embodiment within the context of human-agent interaction have demonstrated the potential positive effects that physical embodiment can have on people's level of engagement and overall perception of the agents with which they are interacting. Wainer, Feil-Seifer, Shell, & Matarić, (2006, 2007) showed that healthy adult participants engaging in a physical/cognitive task, a Towers of Hanoi table-top game, reported a strong preference for a physically embodied SAR system over similar video-only agents in terms of appeal, perceptiveness, watchfulness, helpfulness, and enjoyableness. Powers, Kiesler, Fussell, and Torrey (2007) compared interactions between robots and similar computer-simulated agents that engaged participants in a conversation about basic health habits, and found that participants rated the robots as more helpful, more lifelike, and possessing more positive personality traits than the computer-based agents. Bartneck (2003) conducted a study comparing the effectiveness of an emotionally expressive robot, *eMuu*, with its screen character version in engaging users in a

simple negotiation task, and found that participants exerted more effort and received higher task scores when interacting with the physical eMuu than with the simulated eMuu. Jung and Lee (2004) also demonstrated the positive effects of physical embodiment in relation to interactions with both a Sony Aibo robot and an anthropomorphic dancing robot, April. Bainbridge, Hart, Kim, and Scassellati (2011) found that users in a book-moving task were more likely to fulfill an unusual request and afford more personal space to the agent when interacting with a physically present robot than when interacting with a live video feed of the same robot on a computer screen. Kidd and Breazeal (2008) compared a robotic weight-loss coach (a touch screen with a physical head capable of looking at and speaking to the user) to a similar touch-screen-only device and found that participants interacting with the robotic coach chose to continue with the weight-loss program for twice as long as those interacting with the computer-only device.

While studies such as those mentioned above have previously investigated the effect of physical embodiment in human-agent interaction, most have recruited a participant pool consisting primarily of young adults. However, older adults often respond differently to technology than young adults, as studies have shown (Balakrishnan & Yeow, 2007; Kang & Yoon, 2008), and thus the observed effects do not necessarily generalize across the age span.

Embodiment studies that have targeted the elderly population include the work of Heerink, Kröse, Evers, and Wielinga (2010), who investigated the acceptance of assistive social agents by older adults. While their study was similar to our work, the robot used in their evaluation was a table-top robot (the iCat), and was either controlled via a human operator during interaction with elderly users (Wizard of Oz study), or interacted with users through a touch-screen interface. Furthermore, the interaction consisted primarily of short informational or utility interactions (e.g., medication/agenda reminders, weather forecast, companionship), lasting about 5 minutes and often involving only a single session. Lastly, their work did not explore the agent's role in actively motivating the user to engage in the task, and instead focused solely on the effect of embodiment. In contrast, our SAR system was designed to engage elderly users in fluid, highly interactive exercise sessions, completely autonomously, while providing active feedback, motivation, and guidance on the task. Also, our user study spanned multiple sessions, each lasting 20 minutes, for the dual purpose of system evaluation and comparison between two different SAR coaching embodiments.

The interactivity of the sessions is important because, as we discuss in the next section, the fluidity of the interaction can have a positive influence on the user's intrinsic motivation to engage in the task and can thereby increase his or her enjoyment level during interaction. It is thus interesting to investigate whether the positive effects introduced by the interaction characteristics might alleviate any potential negative effects resulting from embodiment characteristics. Our study addresses this possibility in capturing user evaluations of the system between study groups.

To the best of our knowledge, our user study is the first to comprehensively demonstrate the positive effect of physical embodiment in a SAR-guided healthcare scenario, wherein the SAR agent serves as both an instructor and active participant in the healthcare task with target users.

3. Robot Exercise Coach

3.1 Approach and Methodology

Our approach to designing our SAR system to help address the physical exercise needs of the elderly population was motivated by two basic axioms indicating the essential qualities that the SAR agent coach must possess: 1) the ability to influence the user's intrinsic motivation to perform the task, and 2) the ability to personalize the social interaction to maintain user engagement in the task and build trust in the task-based human-robot relationship. Following the above axioms, we developed five design principles for the SAR system; all are general and can be applied to any SAR-based therapeutic intervention.

3.1.1 Design Principles

The design principles stated that the robot coach should be as follows: 1) motivating, 2) fluid and highly interactive, 3) personable, 4) intelligent, and 5) task-driven. The following discussion elaborates on the importance of each of these qualities in the context of providing healthcare interventions and details how each was incorporated into our SAR exercise system.

1) Motivating. The coaching style and interaction methodology of our SAR exercise system was guided by psychology research in the area of intrinsic motivation. Intrinsic motivation, in contrast to extrinsic motivation, which is driven by external rewards, comes from within an individual and is based on pleasure derived from engaging in an activity. Specifically, our aim was for the robot to be capable of increasing the user's intrinsic motivation to perform the therapeutic task (physical exercise, in this case). Intrinsic motivation has been shown to be more effective than extrinsic motivation in achieving long-term user task compliance and behavior change (Dienstbier & Leak, 1976), which is the ultimate goal for any health-related intervention, technology-based or otherwise. The motivational techniques utilized by our system to accomplish this aim were derived from Csikszentmihalyi's theory of flow (Csikszentmihalyi, 1975), which Deci and Ryan (1985, p. 29) describe as asserting that people are "intrinsically motivated under conditions of optimal challenge." Toward this end, we focused on providing a variety of challenging exercise games of varying degrees of difficulty. We also focused on the fluidity of the interaction, discussed below, as well as on alternating the games at a regular pace to prevent user boredom and/or frustration.

Our design also included additional motivational techniques besides those based on flow theory. For instance, we incorporated indirect competition into the system design by having the robot periodically report the user's high score during each of the exercise games. Indirect competition, wherein the user is challenged to compete against an ideal outcome, has been shown to increase user enjoyment in an otherwise non-competitive task (Weinberg & Ragan, 1979). In addition, self-determination and user autonomy, which have also been shown to increase intrinsic motivation (Fisher, 1978), were implemented in our system by giving the user control over the exercise routine in one of the exercise games, as discussed in more detail in Section 3.3.3.

As indicated by the first basic axiom above, influencing the user's motivation was a key component in our SAR agent design; while our first design principle stresses the importance of motivation in the task scenario, all of our design principles were developed to increase user intrinsic motivation during interaction in one way or another.

2) Fluid and Highly Interactive. A primary goal of our coaching approach was to provide a fluid interaction, which required the robot to both perceive the user's activity and provide active feedback and guidance in real time, all with the aim of maintaining user engagement in the task. According to Csikszentmihalyi (1993), for any task to achieve a state of flow—or maximal enjoyment—in the user, it must establish a clear set of goals, combined with immediate and appropriate feedback. Toward this end, we developed a real-time vision algorithm to recognize user arm movements (detailed in Section 3.5), as well as coaching behaviors to continually provide the user with appropriate feedback to achieve the goals of the exercise games (discussed in Section 3.6). The result is an exercise coach that is highly interactive and responsive to the user, thus promoting a state of flow.

3) Personable. The social interaction between the user and the robot is just as important as the task interaction for achieving success in healthcare interventions. Social interaction is the primary means of relationship building, including in therapeutic scenarios. Many social intricacies contribute to the foundation of a meaningful relationship, both in human-computer interaction and in human-robot interaction. These factors include empathy, humor, references to mutual knowledge, continuity behaviors, politeness, and trust, among others (Bickmore & Picard, 2005). We place great importance on these relationship-building tools; therefore, we integrated each, in

one form or another, into the social interaction component of our robot exercise coach. Primary examples include referring to the user by name; giving praise upon successful user completion of exercise gestures, which has also been shown to increase intrinsic motivation (Vallerand, 1983; Vallerand & Reid, 1984); displaying empathy after failed gesture attempts; and demonstrating continuity to establish an ongoing relationship with the user. In our system, the robot always uses the user's name at first greeting, and also when bidding farewell at the end of a session. Having the robot refer to the user by name, along with providing direct feedback specific to the individual user's performance level and performance history during the games, was an important part of personalizing the interaction. Our SAR exercise system introduced continuity (see Section 3.6) by having the robot refer to previous sessions with the user upon introduction, reference planned future sessions at the end of interaction, and refer to past user exercise performance, such as when reporting previous high scores.

4) *Intelligent*. Trust is a key component to the success of any care-provider/user relationship, and one that is closely linked to the intelligence/helpfulness of the care provider as perceived by the user (Tway, 1994). Tway defines trust as a construct determined by three components: perception of competence, perception of intentions, and capacity for trusting. In the context of SAR-based therapy, the perceived intelligence/competence of the SAR agent by the user is necessarily correlated with the agent's demonstrated abilities in the therapeutic task (e.g., performance capability, evaluation accuracy, etc.). In our SAR exercise scenario, the effectiveness of the robot's user activity recognition procedure was paramount in shaping the user's perception of the robot's intelligence. (The algorithm is discussed in Section 3.5.) In particular, if the user were to perceive the robot as slow or ineffective in evaluating performance, this could lead to a decrease in the user's trust and perception of the robot's value/usefulness in helping to accomplish the desired healthcare goals, according to Tway. This aspect once again stresses the importance of accurate, real-time sensing in the human-robot interaction.

Intelligent robot social behaviors can also have a positive impact on user engagement. For instance, it has been shown that repetitive discourse tends to have a negative impact on user motivation, whereas increased variability, in discourse and general behaviors, tends to enhance task-based user engagement (Bickmore, Schulman, & Yin, 2010). Therefore, in our SAR system, we placed special attention on ensuring variety in the robot's utterances to minimize the perceived repetitiveness of the robot's verbal instructions/comments. Toward this end, the robot always draws from a list of phrases that emphasize the same point when speaking to the user, choosing one randomly at run time. For example, there are more than ten different ways in which the robot can praise the user after he or she completes an exercise (e.g., "Awesome!"; "Nice job!"; "Fantastic!"). Variability is also introduced in discourse during game transitions, game introductions, and upon farewell. General interaction behaviors were also designed to reduce predictability; for example, by ensuring that the order of games, game timing, and session schedules were different for each interaction. Further information regarding the methods employed to add variability to our SAR behaviors and interaction sessions is provided in Section 3.6.

5) *Task-Driven*. Perhaps the most important property of the interaction is that it be consistent in working to achieve the goals of the healthcare task, or in our case, in motivating exercise performance to achieve desired overall health benefits. As mentioned above, it is important for the task to be intrinsically motivating and enjoyable and for the coach be personable and intelligent. However, the consistency of the interaction towards accomplishing the healthcare task is fundamental for the user to trust that the system can help to obtain beneficial health outcomes, as this consistency influences the user's perception of the intentions of the SAR agent. Agent intentions, according to Tway (1994), must be seen as mutually beneficial. Without this consistency, users may perceive the robot as simply entertaining, rather than helpful.

Furthermore, it is important that the tasks not only be healthcare-driven, but that they also be successful in achieving the desired therapeutic behavior; in the case of our SAR exercise coach,

this means the interaction must elicit consistent physical exercise among the users (measurable through objective quantitative metrics), a result that was achieved in our user study with older adults as detailed in Section 5.

3.1.2 Discussion of Design Principles

Our five SAR design principles were developed in an attempt to capture and address the key components necessary for successful SAR-based therapy. Each of the design principles was developed based on findings from research in psychology, specifically in the areas of intrinsic motivation and trust, and from research in HCI and HRI, specifically regarding the effects of relational agent behaviors on user engagement. Although some previous work has addressed the design of social agents for obtaining long-term engagement and trust, most has focused on agent guidance in task scenarios in which the agent does not monitor or instruct the activity itself (e.g., Bickmore & Picard, 2005; Kidd & Breazeal, 2008).

We have presented a set of five clear and general SAR design principles that can be applied to a variety of different SAR-based therapeutic domains, including those wherein the SAR agent provides real-time instruction, activity monitoring, active guidance and feedback, task participation, and continuity behaviors spanning multiple sessions of interaction. In the following sections we provide examples of how to incorporate these design principles into practice with a SAR agent, as we elaborate on the implementation details of our SAR exercise system developed to engage older adults in physical exercise.

In Sections 4 and 5 we discuss a user study conducted with older adults to evaluate our SAR approach, design, and algorithms in motivating task engagement and performance. It is important to note, however, that the focus of the study was not on validating each of the five design principles independently. Prior work by our group and others has already demonstrated the positive effect of many aspects of the above design principles individually. For example, previous studies conducted by our lab have shown the positive influence on intrinsic motivation and user engagement of online task modification to achieve optimal challenge (Fasola & Matarić, 2010a), praise and relational discourse, adjustable user autonomy, and fluidity of interaction (Fasola & Matarić, 2012). Instead, the focus of the conducted user study was on the evaluation of our SAR exercise system and design principles as a whole, in addition to investigating the role of physical embodiment and its effect on user intrinsic motivation and user attitudes towards the system.

3.2 Interaction Scenario

Our robot exercise system consists of a socially assistive robot whose purpose is to monitor, instruct, evaluate, and encourage users to perform simple physical exercise. During the exercise sessions, the robot asks the user to perform seated arm gesture exercises. This type of seated exercise, called “chair exercise” or “chair aerobics,” is commonly practiced in senior living facilities. Chair exercises are highly regarded for their accessibility to those with low mobility, for their safety, and for their health benefits, such as improved flexibility, muscle strength, ability to perform everyday tasks, and even memory recall (Baum et al., 2003; Dawe & Moore-Orr, 1995; McMurdo & Rennie, 1993; Thomas & Hageman, 2003).

The one-on-one interaction scenario consists of the user sitting in a chair facing the robot. In this system, predating the availability of the Microsoft Kinect, the range of the robot’s arm motion in the exercises is planar and restricted to the sides of the body in order to maximize the accuracy of the robot’s real-time visual detection of the user’s arms. A black curtain is used as a backdrop, also to facilitate the robot’s visual perception of the user’s arm movements. An example of the exercise setup and human-robot interaction is shown in Figure 1.

The system allows the user to communicate with the robot through the popular Nintendo Wiimote wireless Bluetooth button control interface. There are two buttons available for the user to respond to prompts from the robot, labeled “Yes” and “No,” and one button for the user to request a rest break at any time during the interaction. The Wiimote is not used for motion sensing



Figure 1. The setup for the one-on-one interaction between user and robot coach.

of any kind; its role is solely for communication with the robot through button presses. It is also important to note that the robot conducts the exercise sessions, evaluates user performance, and gives the user real-time feedback all completely autonomously, without human operator intervention at any time during the exercise sessions.

3.3 Exercise Games

Four exercise games are available in our system: the Workout game, the Sequence game, the Imitation game, and the Memory game. Each game presents a different level of challenge and necessary skill set, in accordance with our system methodology of varying game types for increased user enjoyment and to promote intrinsic motivation in the therapeutic interaction. Following is a brief description of each game.

3.3.1 Workout Game

In the Workout game, the robot fills the role of a traditional exercise instructor by demonstrating the arm exercises with its own arms and asking the user to imitate. The robot shows only one exercise gesture at a time (involving one or both arms). Upon successful completion by the user, the robot generates a different gesture and the process repeats. The robot gives the user feedback in real time, providing corrections when appropriate (e.g., “Raise your left arm and lower your right arm” or “Bend your left forearm inward a little”), and praise in response to each successful imitation (e.g., “Great job!” or “Now you’ve got the hang of it.”). This game has the fastest pace of the four exercise games, as the users generally complete the requested gestures quickly.

3.3.2 Sequence Game

The Sequence game is similar to the Workout game in that the robot demonstrates arm exercises for the user to repeat. However, instead of showing each gesture for the user to perform only once, the robot demonstrates two gestures for the user to repeat in sequence for three repetitions (resulting in six gesture completions per sequence). The robot keeps verbal count of the number of iterations of the sequence performed in order to guide the user, while providing feedback throughout (e.g., after the user completes the first pair, the robot says “One,” and upon completion of the second pair, it says “Two”). This game requires the user to remember the gesture pair while completing the sequence, in addition to promoting periodic movements from the users. While slower-paced than the Workout game, these periodic movements are physically challenging and can cause the users to physically exert themselves more than in any of the other three games.

3.3.3 Imitation Game

In the Imitation game, the roles of the user and robot are reversed relative to the Workout game: The user becomes the exercise instructor, showing the robot what to do. The robot encourages the

user to create his/her own arm gesture exercises and imitates user movements in real time. Since the roles of the interaction are reversed, with the robot relinquishing control of the exercise routine to user, the robot does not provide instructive feedback on the exercises. However, the robot does continue to speak and engage the user by means of encouragement, general commentary, humor, and movement prompts if necessary. As the robot imitates the user's movements in real time, the interaction in this game is the most fluid and interactive of all the exercise games. Furthermore, this is the only game designed to promote user autonomy and creativity in the exercise task. Both of these characteristics are beneficial for increasing the user's intrinsic motivation to engage in the task.

3.3.4 Memory Game

The goal of the Memory game is for the user to try and memorize an ever-longer sequence of arm gesture poses, and thus compete against his/her own high score. The sequence is determined at the start of the game and does not change for the duration of the game. The arm gesture poses used for each position in the sequence are chosen at random at run time, and there is no inherent limit to the sequence length, thereby making the game challenging for users at all skill levels. Each time the user successfully memorizes and performs all shown gestures without help from the robot, the robot shows the user two additional gestures to add to the sequence, and hence the game progresses in difficulty. The robot helps the user to keep track of the sequence by counting along with each correct gesture and reminding the user of the correct poses, while demonstrating empathy upon user failure (e.g., "Oh, that's too bad! Here is gesture five again."). This game is the most cognitively challenging of all the exercise games, and primarily serves to encourage competition with the user's past performance.

3.4 System Architecture

The system architecture is comprised of six independent software modules: vision and world model, speech, user communication, behaviors, robot action, and database management. A diagram of the system architecture, showing the connections among the system modules, is provided in Figure 2. The following briefly describes the role of each system module:

Vision and World Model Module. This module is responsible for providing information regarding the state of the user to the behavior module for the robot to make task-based decisions during interaction; for example, by providing spoken or demonstrative feedback, prompts for user movement, or by imitating user arm poses. The input for the visual user activity recognition procedure is a monocular USB camera, with a frame resolution of 640x480 pixels. The presence of the user and the locations of the face, hands, and arm angles of the user are captured by the vision process and stored in the world model for subsequent use. The details of the vision algorithm for user state recognition are provided in Section 3.5.

Speech Module. The speech module is responsible for translating the speech requests from the behaviors, provided as text strings, into spoken words (synthesized audio). This module uses the commercially available NeoSpeech text-to-speech (TTS) engine to synthesize the given text into communicative speech. The speech module plays the synthesized speech through the robot's speakers, and is also responsible for synchronizing the lip movements of the robot to match the start and end of spoken utterances. To accomplish this synchronization, the module analyzes the wave file produced by the TTS engine and locates the key start and end points of silence, which correspond to pauses in the speech; for example, due to commas or periods in the input text. Then, at the appropriate times during speech playback, the speech module sends lip movement commands to the action module to simulate the opening and closing of the robot's mouth during speech. The lip synchronization procedure is used to enhance the naturalness of the interaction to promote user engagement with the robot. The alternative of keeping the robot's lips closed during

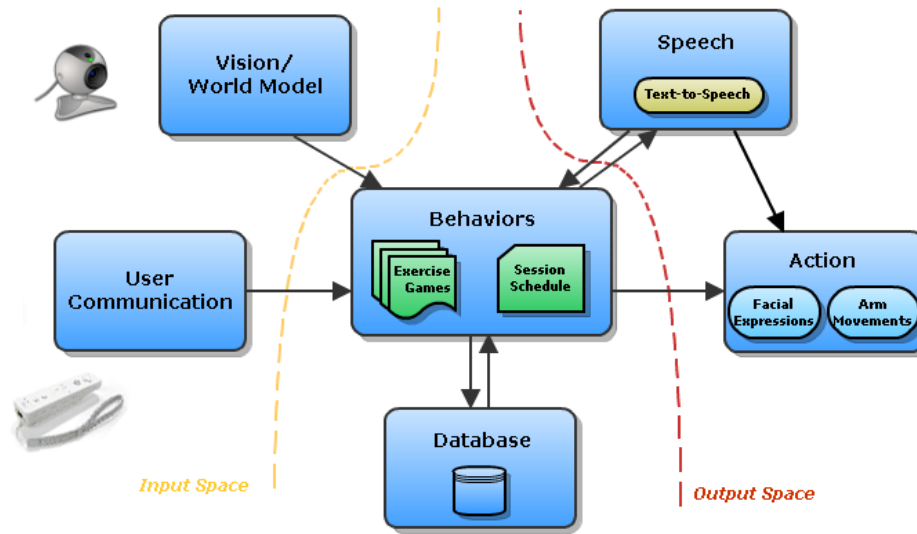


Figure 2. Diagram of system architecture, including the six system modules: vision and world model, speech, user communication, behaviors, action output, and database management.

speech might appear incongruent to the user and thus have a disengaging effect. The speech module also communicates speech state information to the behavior module, such as whether or not it is currently speaking, so that the behavior finite state machines can make transitions accordingly (e.g., by waiting for the robot to finish giving a praise comment before demonstrating the next exercise gesture).

User Communication Module. The user communication module is responsible for receiving direct input from the user and relaying the communication attempt to the behavior module. The input in our SAR system is via a Nintendo Wiimote wireless remote control device, which is capable of sending button presses wirelessly via Bluetooth. The simplicity of the input method allows the user to request breaks during interaction, respond to yes/no questions or prompts made by the robot, and also to report key information during the exercise games (e.g., notifying the robot of a gesture completion attempt during the Memory game). Each communication type is represented by a different button specifically labeled on the remote control. This method enables simple two-way communication between the robot and user, is easy to use, and avoids the sensing difficulties often found with alternative communication input methods (e.g., speech recognition).

Behavior Module. The behavior module is responsible for producing all of the coaching behaviors of the system, including steering the interaction and exercise games, providing verbal feedback, demonstrating actions, responding to user input, monitoring user activity and progress in the task, and recording important information to the database. The behavior module therefore communicates with each of the six system modules and represents the system's main loop. The behavior module also keeps track of important session variables, including session time and total duration, current game time, and session schedule.

Action Module. The action module is responsible for sending motor commands to the robot to execute. These include facial expressions (e.g., moving eyebrows and lips) and arm movements. The module was designed to promote robot/agent platform independence, and as a result, all action requests made by connected system modules are only translated to the robot platform-dependent motor commands in the action module, thereby abstracting this knowledge away from all other system modules. This hierarchy allows for seamless integration of different platforms for the social agent, without the need for modification of any of the remaining software modules. In

the user study presented in Section 4, we explored the use of both a physical humanoid robot and a virtual simulation of the same robot for system evaluation, thus demonstrating the interoperability of the system design.

Database Module. The database module is responsible for storing important user-related information regarding task performance and progress, in addition to session history. Examples include the date and time of previous sessions, exercise performance in each of the exercise games, number of breaks taken, high scores, etc. This information is available for retrieval during interaction by the behavior module, which uses the information to implement continuity behaviors and for reporting user performance and progress. The database information is also useful for obtaining quantitative metrics of user performance for post-session/user study system evaluation.

3.5 User Activity Recognition

In order to monitor user performance and provide accurate feedback during the exercise routines, the robot must be able to recognize the user's arm gestures. To accomplish this, we developed a vision algorithm that recognizes the user's arm gestures/poses in real time, with minimal constraints on the surrounding environment and the user.

Several different approaches have been developed to accomplish tracking of human motions, both in 2D and 3D, including skeletonization methods (Fujiyoshi & Lipton, 1998; Jenkins, Chu, & Matarić, 2004), gesture recognition using probabilistic methods (Waldherr, Thrun, Romero, & Margaritis, 1998), and color-based tracking (Wren, Azarbajani, Darrell, & Pentland, 1997), among others. We opted to create an arm pose recognition system that utilizes our simplified exercise setup in order to achieve real-time results without imposing any markers on the user.

To simplify the visual recognition of the user, a black curtain was used to provide a static and contrasting background for fast segmentation of the user's head and hands, the most important features of the arm pose recognition task, independent of the user's skin tone. The arm pose recognition algorithm consists of the following four major steps:

1) *Create segmented image:* The original grayscale camera frame is converted into a black/white image by applying a single threshold over the image. All pixels below the threshold are set to black, and the rest to white. The threshold is set by the experimenter to appropriately segment out the user from the background, accounting for the specific lighting of the environment. Figure 3(a) shows an original grayscale image captured from the camera, with the segmented image shown in Figure 3(b).

2) *Detect the user's face:* The OpenCV frontal face detector is used to determine the location and size of the user's face. With these values, an estimate for the shoulder positions on both sides of the body is made.

3) *Determine hand locations:* The hand locations of the user are determined by examining the extrema points of the body pixels (maximum/minimum white pixel locations along x and y directions away from body) inside the region above the chest line and to the side of the face in the segmented image (subtracting the head and estimated body regions—represented by green and blue bounding boxes in Figure 3). The algorithm applies a simple set of rules, or heuristics, to choose which extrema points correspond to the hand location for a given arm. For example, if the highest white pixel (body pixel) in the segmented image on the left side is further than the approximate shoulder-to-forearm distance, then that is the left hand location.

These extrema point heuristics are valid for this domain, as the silhouette of a person demonstrating planar arm movements (i.e., only two degrees of freedom) will always feature at least one of these extrema points at each of the hand locations, due to the geometric properties of planar two-link manipulators. This planar two-link assumption thus validates the use of extrema point heuristics performed on a binary segmentation of the image (approximated silhouette) to detect the hand locations of the user.




Figure 3. (a), (c), (d) Example face and arm angle detection results; (b) segmented image.

4) *Determine arm angles*: Once the hand location for a given arm is found, the elbow point is estimated, which in turn provides the desired arm angles. The elbow point is estimated using the white pixel (body pixel) that lies furthest from the line connecting the hand position and the shoulder position, while also not exceeding the maximum allowable distance from the hand (to enforce forearm length restriction). Examples of arm angle detection results can be seen in Figure 3(a), (c), and (d).

The vision module only searches for frontal face views, and thus detection rates depend largely on whether or not the user is facing the robot. However, high detection rates are not actually necessary for accurate gesture recognition, as the most recent detected face location is used in the arm detection procedure if a face is not found in the current frame. This substitution works well, since the user's head position generally remains stationary while the user is seated in a chair in front of the robot throughout the interaction.

The arm pose recognition algorithm runs with an average frame rate of 20 fps on a 2.4 GHz Intel Core 2 Duo processor, thus achieving our aim of real-time user activity recognition. While the arm detection procedure was not formally evaluated, the system was confirmed to be robust to different types of clothing, lighting, and user body types, with a notable arm angle estimation accuracy observed during the course of our user studies with both older and young adults.

Although a black curtain was used to aid the visual recognition procedure during our user studies, the algorithm does not rely on equipping the environment, but only on a binary user segmented image. The algorithm has been modified, for example, to receive a motion frame produced by sequential frame differencing in replacement of a color segmented image. The motion frame, once thresholded, serves as the segmentation input in the first step of the algorithm, and the user arm angles can be determined in the same manner as described. Using motion instead of color segmentation, along with Kalman filtering of the positions of the hands and elbows, provides a more general recognition approach for domain-independent use, even though the overall algorithm remains the same.

The development of the SAR exercise system and visual recognition procedure predated the availability of the Microsoft Kinect. Our future implementations of the system will utilize Kinect-type 3D vision technology and thus do away with the curtain and the planar limits of the motions.

Nevertheless, the 2D nature of the exercises was not noted as an issue by any of the participants in our user studies. Furthermore, the real-time vision algorithm described is general and can be applied to a variety of domains, including those wherein use of the Kinect is not appropriate or feasible (e.g., outdoors in sunlight).

3.6 Behaviors

The behavior module represents the main loop of the SAR exercise coach system. It communicates with all of the other system modules and manages the flow of the interaction, including transitioning between the different exercise games, managing the session and game clocks, saving user performance and session statistics to the database, and handling user-requested breaks, among other session-related tasks.

3.6.1 Interaction Flow and Behavior Management

Each of the four exercise games is managed by its own game behavior. In accordance with our principle of encouraging users to be intrinsically motivated to engage in the task, the behavior module transitions between game types every 1–2.5 minutes in order to diminish the possibility of user boredom, which would negatively affect motivation level. The behavior module chooses which game to transition to based on a predefined session schedule, which specifies the order of the games to play and their respective durations. The session schedule is defined *a priori* by the experimenter such that each session's order of games is unique. This requirement is enforced to reduce the predictability of the exercise regimen for the user, which might negatively affect his or her perception of the system (Bickmore, Schulman, & Yin, 2010). The session schedule, however, though unique for each individual session, follows approximately the same structure throughout all of the sessions. Specifically, the Workout, Sequence, and Memory games are each allotted approximately 25% of the total session time (75% for all three games), with the Imitation game being allotted 15% of the total session time and the remaining 10% being allocated for user breaks, transition behaviors, and the introduction and conclusion of the session. The Imitation game receives a lower allotment compared to the other three games due to the fact that the user is in control of the exercise. Although the Imitation game promotes user autonomy and users often find the game stimulating, in pilot studies we've observed users may get bored if the game goes on for too long. Thus the Imitation game is usually only played for 1–1.5 minutes at a time during the course of the exercise session. If a session schedule is not set, or is incomplete, the robot chooses a game at random to play at regular game change intervals. The exact durations for each game were determined empirically by taking into account multiple factors, including the total duration of the session (20 minutes), the allotted game time ratio, the time required for the robot to familiarize the user with the specific game rules, and the minimum time estimated for the user to appropriately experience the game.

The session schedule allows the robot to guide the interaction and determine which games to play when. An alternative approach would be to allow the user to choose which game to play at each game transition interval, thus perhaps increasing user intrinsic motivation to engage in the task. We have previously conducted a user study with elderly participants to investigate the role of user choice in the exercise scenario (Fasola & Matarić, 2012) and found that people's preferences regarding user choice vs. robot choice were mixed, slightly favoring having the robot decide which games to play. This can be interpreted as the users preferring to have the robot remain in the role of exercise instructor, which dictates that the robot should be responsible for deciding which exercise regimen is best. For the user study discussed in this paper, we predefined the session schedules and had the robot determine which games to play to ensure that each participant received similar session interactions to allow for unbiased comparison between them.

Though comprising most of the interaction time, the individual game behaviors are not the only behaviors available within the behavior module. Continuity behaviors, desirable for their relationship-building abilities (as discussed previously), were implemented in both the

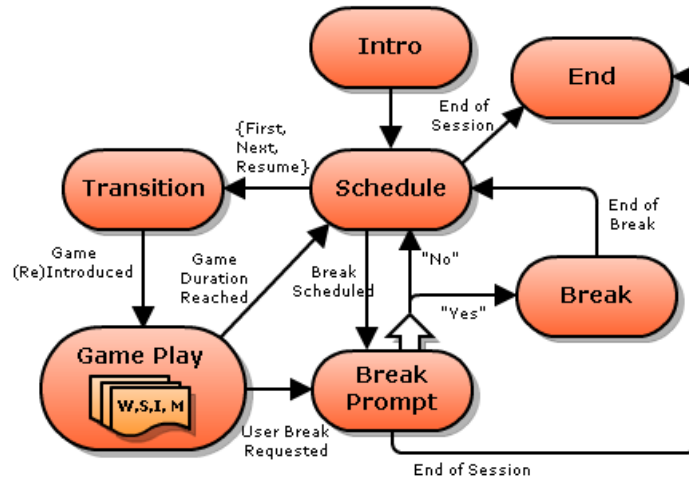


Figure 4. Diagram of the behavior module's finite state machine (FSM). The FSM illustrates the flow of the interaction during the exercise sessions. The white arrow emanating from the Break Prompt state denotes a user decision.

introduction and conclusion of the exercise sessions. These continuity behaviors are responsible for expressing the time-extended nature of the interaction. Examples of these behaviors include having the robot acknowledge past interactions, as in "It's nice to see you again, Alice; last time was fun." These also dictate how the robot re-introduces the games; the robot provides enough information to remind the user of the rules of the exercise game without re-stating everything over again. This serves to acknowledge the fact that the user has already engaged in the exercise game in a previous session.

The remaining behaviors are related to breaks that provide the user with short periods of rest during the interaction so the exercise session is not too demanding. Breaks are scheduled every 5–6 minutes (3 breaks per session). Whenever a break is scheduled, the robot first confirms with the user if a break is desired, and depending on the response, either proceeds to take a 30-second break or continues with the next exercise game. Thus users who wish to rest are allowed to do so, and those who would rather continue playing have that option as well. This presents another example of user autonomy within the interaction, in addition to promoting a level of optimal challenge for each individual user to achieve maximal enjoyment.

A diagram of the behavior module's finite state machine is given in Figure 4 to illustrate the transitions between the different sub-behaviors and the overall flow of the interaction during the exercise sessions.

3.6.2 Feedback Procedure

All four of the game behaviors employ some form of feedback and guidance throughout the interaction. For example, in the Workout game the robot gives the users verbal feedback on how they can alter their arm positions to achieve the correct arm gesture for the current exercise, in the Memory game the robot can remind the users to perform a specific gesture number, and in the Imitation game the robot can prompt the users to move their arms if they happen to stop for more than a brief period. In many cases, the robot needs to provide the same verbal feedback to the user more than once, but must do so without sounding repetitive. As discussed earlier, strict repetitiveness, apart from being a possible annoyance to the user, represents a form of predictability in the nature of the robot and may negatively affect user engagement in the task (Bickmore, Schulman, & Yin, 2010), which could ultimately be detrimental to the success of the

healthcare intervention. We avoid repetitiveness in all comments made by robot, including feedback, praise, empathy, humor, game rules, transitions, and greetings and farewells.

We aimed for variability in verbal feedback so as to avoid the perception of repetitiveness as much as possible. When sequential verbal feedback comments differ, as in “Stretch out your right arm” followed by “Now raise your left arm above your head,” the variability is inherent. However, when the robot needs to report the same feedback comment in sequence, variability must be introduced into the feedback statement itself. We accomplish this variability by introducing filler words. Example filler words include the user’s name and the words “try” and “just.” Utilizing the user’s name not only serves to provide variability, but also works to gain the user’s attention, in addition to further personalizing the specific feedback comment. The feedback procedure chooses randomly at run time whether to use the user’s name, a filler word, or both when repeating a specific feedback phrase.

4. User Study Evaluation with Older Adults

We designed and conducted a user study with older adult participants to evaluate the effectiveness of our SAR approach and system design. Another aim of the study was to investigate the role of physical embodiment in the robot exercise system. Specifically, the study compared the effectiveness and participant evaluation of our physical humanoid robot to those of a computer simulation of the same robot shown on a flat-panel display.

4.1 Study Design

Study participants were divided into two groups, physical robot embodiment vs. virtual robot embodiment, and the study consisted of a total of four 20-minute sessions of exercise interaction with the system, conducted over a two-week period.

The following subsections describe the robot platforms and the between-subjects embodiment comparison method in detail.

4.1.1 Robot Platforms

To address the role of the robot’s physical embodiment, we used Bandit, a biomimetic anthropomorphic robot platform that consists of a humanoid torso (developed with BlueSky Robotics) mounted on a MobileRobots Pioneer 2DX mobile base. The torso contains 19 controllable degrees of freedom (DOF), which include the following: 6-DOF arms (x2), 1-DOF gripping hands (x2), a 2-DOF pan/tilt neck, 1-DOF expressive eyebrows, and a 2-DOF expressive mouth. A photograph of the physical robot can be seen in Figure 5(a).

The robot’s virtual embodiment consisted of a computer simulation of Bandit shown on a 27-inch flat-panel display. The size of the display was chosen to approximate the average size display that would be available in a typical household for use with the robot exercise system, including laptop displays (15 inch), computer monitors (24 inch), and television screens (40 inch). A sample computer simulation image of the virtual robot and a photograph of the virtual embodiment on the flat-panel display are shown in Figure 5(b) and (c), respectively.

For the physical robot embodiment, the USB camera used in the visual user activity recognition procedure was placed at the waist of the torso, whereas for the virtual embodiment, the camera was attached to the top of the television display. The difference in camera location did not affect the accuracy of the visual recognition of the user’s movements.

4.1.2 Between-Subjects Design

Survey data were collected at the end of the first and fourth sessions in order to analyze participant evaluations of the robot and the interaction with the exercise system in both conditions over time. Data gathered from the first and fourth session post-session surveys were analyzed using a two-tailed independent two-sample T-test assuming unequal variances to test for significant differences among the participant evaluations of the robot and the interaction across both conditions. Survey

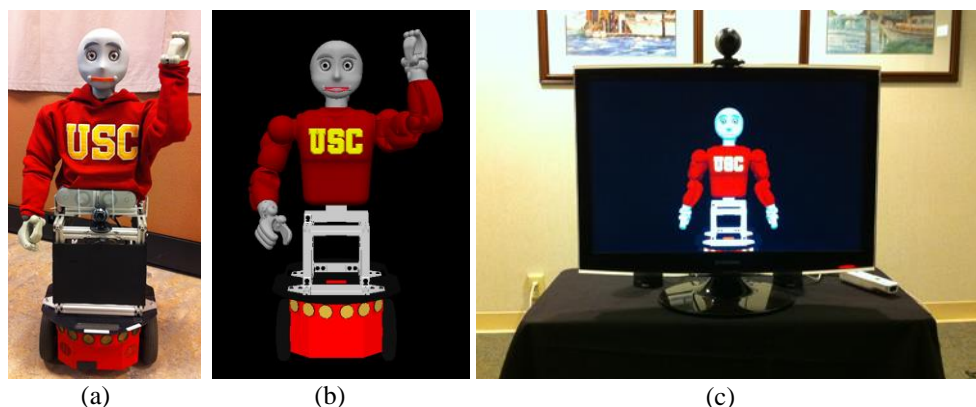


Figure 5. (a) Physical robot; (b) virtual robot computer simulation; (c) virtual robot on the screen, with camera.

results from the fourth session were used to perform the final comparison analysis, as they were less likely to contain scores influenced by the effect of novelty.

4.2 Participant Statistics

We recruited elderly individuals to participate in the study through a partnership with *be.group*, an organization of senior living communities in Southern California, using flyers and word-of-mouth. We offered a \$50 Target gift card to those willing to participate in all four sessions of the study. Thirty-seven people responded, of whom four were omitted due to inconsistent/incorrect answers to survey questions that were used to identify questionable survey results. Thus, there were a total of 33 participants whose data were analyzed. Half of the participants were placed in the physical robot group ($n = 16$), and the other half were placed in the virtual robot group ($n = 17$). The sample population consisted of 27 female participants (82%) and 6 male participants (18%). Participants' ages ranged from 68–88, and the average age was 76 ($SD = 6.32$).

4.3 Measures

4.3.1 Evaluation of Interaction

There were two dependent measures used to evaluate the interaction with the robot exercise system. The first measure was the *enjoyableness of the interaction*, collected from participant assessments of the interaction according to six adjectives: enjoyable, interesting, fun, satisfying, entertaining, boring, and exciting (Cronbach's $\alpha = .92$). Participants were asked to rate how well each adjective described the interaction on a 10-point scale, anchored by "Describes Very Poorly" (1) and "Describes Very Well" (10). Ratings for the adjective "boring" were inverted to keep consistency with the other adjectives that reflect higher scores as being more positive.

The second measure was the perceived *value or usefulness of the interaction*. Similarly, participants were asked to evaluate how well each of the following four adjectives described the interaction: useful, beneficial, valuable, and helpful (Cronbach's $\alpha = .96$). The same 10-point scale anchored by "Describes Very Poorly" (1) and "Describes Very Well" (10) was used.

4.3.2 Evaluation of Robot

Companionship of the robot was measured from participant responses to nine 10-point semantic differential scales concerning the following robot descriptions: bad/good, not loving/loving, not friendly/friendly, not cuddly/cuddly, cold/warm, unpleasant/pleasant, cruel/kind, bitter/sweet, and distant/close (Cronbach's $\alpha = .89$). These questions were derived from the Companion Animal

Bonding Scale of Poresky, Hendrix, Mosier, and Samuelson (1987). The companionship of the robot was measured to assess potential user acceptance of the robot as an in-home companion, thereby demonstrating the capability of the system to facilitate independent living.

Participants evaluated the *helpfulness of the robot* by rating four robot characteristics on a 10-point scale: useful, valuable, beneficial, and helpful (Cronbach's $\alpha = .96$). The rating scale was anchored by "Not at all" (1) and "Absolutely" (10). Using the same rating scale, the *intelligence of the robot* was measured according to the following four adjectives: competent, clever, intelligent, and smart (Cronbach's $\alpha = .93$).

To help capture the robot's social attributes, we measured both the *social attraction towards the robot* and the *social presence of the robot*. Social attraction was measured by a modified version of the Interpersonal Attraction Scale of McCroskey and McCain (1974). Participants reported their level of agreement with the following four statements: I think Bandit could be a friend of mine; I think I could spend a good time with Bandit; I could establish a personal relationship with Bandit; I would like to spend more time with Bandit (Cronbach's $\alpha = .88$). The statements were rated on a 7-point scale anchored by "Very Strongly Disagree" (1) and "Very Strongly Agree" (7). Social presence was measured by a 10-point scale anchored by "Not at all" (1) and "Very much" (10) with questions from Jung and Lee (2004) such as the following: While you were exercising with Bandit, how much did you feel as if you were interacting with an intelligent being? (Cronbach's $\alpha = .87$).

To assess perceptions of the exercise capabilities of the system, we measured participant evaluations of the robot as an *exercise partner*. Four items were used, each measured according to a 10-point scale anchored by "Not at all" (1) and "Very much" (10): How much did you enjoy exercising with Bandit?; How likely would you be to recommend Bandit as an exercise partner to your friends?; How much would you like to exercise with Bandit in the future?; How much have you been motivated to exercise while interacting with Bandit? (Cronbach's $\alpha = .93$).

4.3.3 User Performance Measures

To help assess the effectiveness of the SAR exercise system in motivating exercise among the participants, we collected fifteen different objective measures during the exercise sessions regarding user performance and compliance in the exercise task.

Five performance measures were captured during user interaction in the Workout game, including the *average time to gesture completion* (from the moment the robot demonstrates the gesture to successful user completion of the gesture), *number of seconds per exercise completed*, *number of failed exercises*, *number of movement prompts* by the robot to the user due to lack of arm movement, and *feedback percentage*. The feedback percentage measure refers to the fraction of gestures, out of the total given, for which the robot needed to provide verbal feedback to the users regarding their arm positions in order to help guide them to correct gesture completion.

For the Sequence game, we captured four objective measures: the *average time to gesture completion*, *average number of sequences completed*, *average number of gesture pairs completed*, and the *feedback percentage*. In the Memory game, we recorded the *maximum score* over all sessions and all users, *average maximum score* among individual users, and *average time per gesture attempt*. For the Imitation game, the only measure captured was the *number of movement prompts* by the robot due to lack of user arm movement.

The remaining two measures concerned user activity during the entire exercise session: the *average total number of exercises completed*, and *number of breaks* taken by the user.

4.3.4 Relation to Design Principles

The survey evaluation measures and the user performance measures together serve to evaluate our SAR system approach and design principles. Specifically, participant evaluations of each of the five SAR design principles were effectively captured by one or more of the described study measures. Table 1 provides a summary of the relations between each of the five design principles outlined in Section 3.1 and their corresponding survey and/or objective performance measures.

Table 1. Summary of the relations between the five design principles and their related study measures (evaluation and/or performance based).

Design Principle	Related Measures (Survey and/or Performance)
1) <i>Motivating</i>	<i>Enjoyableness of interaction</i> According to Csikszentmihalyi (1975), intrinsically motivating activities are characterized by user enjoyment.
2) <i>Fluid and Highly Interactive</i>	<i>Average time to gesture completion; seconds per exercise; feedback percentage</i> The fluidity and real-time interactive nature of the task is most appropriately illustrated by the speed and accuracy of the objective user performance statistics (which are tightly linked to the feedback and responsiveness of the robot).
3) <i>Personable</i>	<i>Helpfulness; companionship; social attraction; social presence</i> The robot's personable qualities, such as expressing empathy, reassurance, praise, reference to past user performance, continuity between sessions, humor, politeness, and referring to the user by name, are best characterized by the perceived helpfulness and companionship of the robot. Furthermore, more general assessments of the personable traits of the robot (and of the robot's embodiment) are captured by the social attraction and social presence measures.
4) <i>Intelligent</i>	<i>Intelligence of the robot</i> Participant evaluations of the perceived intelligence of the robot provide an appropriate one-to-one assessment of this design principle in practice.
5) <i>Task-Driven</i>	<i>Value/usefulness of interaction; robot as exercise partner; all user exercise performance measures</i> Participant evaluations of the usefulness of the interaction, and of the robot as an exercise partner, serve to illustrate the perceived effectiveness of the system in encouraging users to exercise (important for establishing user trust in the system). Furthermore, the objective performance measures serve to demonstrate the actual effectiveness of the system in eliciting physical exercise among the users (the goal of the specific healthcare task).

4.4 Hypotheses

Based on the related research on embodiment effects, six hypotheses were established for the embodiment comparison portion of this study:

Hypothesis 1: Participants will evaluate the enjoyableness of their interaction with the physical robot more positively than their interaction with the virtual robot.

Hypothesis 2: Participants will evaluate the value/usefulness of their interaction with the physical robot more positively than their interaction with the virtual robot.

Hypothesis 3: Participants will evaluate the helpfulness of the physical robot more positively than that of the virtual robot.

Hypothesis 4: Participants will be more socially attracted to the physical robot than the virtual robot.

Hypothesis 5: Participants will experience a greater sense of social presence when interacting with the physical robot than when interacting with the virtual robot.

Table 2. Results of between-subjects data comparison for all $n = 33$ older adult participants showing means and standard deviations (in parentheses).

Dependent Measure	Physical Robot	Virtual Robot	
<i>Interaction Evaluation</i>		<i>Between-Subjects Analysis</i>	
Enjoyable	7.51 (1.77)*	6.00 (2.01)	
Valuable/Useful	8.14 (1.66)*	6.19 (2.39)	
<i>Robot Evaluation</i>			
Helpful	8.11 (1.98)*	6.26 (1.98)	
Social Attraction	4.70 (1.40)*	3.61 (1.54)	
Social Presence	7.88 (0.94)*	6.47 (2.01)	

Companion	7.48 (2.07) ⁺	6.23 (1.84)	
Intelligence	8.17 (2.02) ⁺	6.76 (2.09)	
Exercise Partner	7.18 (2.17) ⁺	5.76 (2.18)	

⁺ $p < .10$, * $p < .05$

5. Results

5.1 Embodiment Comparison Results

5.1.1 Between-Subjects Comparison Results

A two-tailed independent T-test was performed on the survey data following the fourth exercise session to compare participant evaluations of the robot embodiments and of the overall SAR interaction across the two study groups. Table 2 provides the complete set of between-subjects comparison results.

Consistent with Hypothesis 1, the participants evaluated the interaction with the physical robot embodiment as more enjoyable than the interaction with the virtual robot embodiment ($t[31] = 2.29, p < .03$). Hypothesis 2 was supported by the data as well, as the participants evaluated the interaction with the physical robot as more valuable/useful than the interaction with the virtual robot ($t[29] = 2.72, p = .01$).

Regarding the direct evaluations of both robot embodiments, the participants rated the physical robot as more helpful than the virtual robot ($t[31] = 2.66, p = .01$), consistent with Hypothesis 3, and as more socially attractive ($t[30] = 2.09, p < .05$), supporting Hypothesis 4. Concerning social presence, the data were consistent with Hypothesis 5, as the participants reported feeling a stronger sense of social presence with the physical robot than with the virtual robot ($t[23] = 2.59, p < .02$).

Evaluations of the non-hypothesis-testing system performance measures were also favorable to the physical robot, though not to a significant degree, as the participants rated the physical robot as somewhat more of a companion ($t[30] = 1.81, p < .08$), more intelligent ($t[31] = 1.96, p < .06$), and a moderately better exercise partner ($t[31] = 1.87, p = .07$) than the virtual robot.

5.1.2 Performance Comparison Results

There were no significant differences found in participant performance between the two study groups, with participants from both the physical robot and virtual robot groups achieving equally high performance. For example, the average gesture completion time in the Workout game during the fourth session for the physical robot group was 2.29 seconds (SD = 0.62), compared to 2.08 seconds (SD = 0.45) for participants in the virtual robot group ($t[29] = 1.11, n.s.$), and the feedback percentage for the physical robot group was 6.0% (SD = 5.7), compared to 4.4% (SD =

4.2) for the virtual robot group ($t[29] = 0.96$, *n.s.*). Further discussion of the exercise performance statistics of the older adult study participants is provided in Section 5.2.2.

5.1.3 Discussion of Embodiment Comparison Results

The results of the between-subjects embodiment comparison analysis show a strong participant preference for the physical robot embodiment over the virtual robot embodiment, as hypothesized. More generally, the results illustrate the wide-reaching effect that SAR agent embodiment type has on the interaction and overall perception of the SAR system by the user. In particular, the study results show the influence of agent embodiment on user motivation, enjoyment, and perceptions of the value/usefulness of the interaction, in addition to showing the influence of embodiment on the perceived personable qualities of the robot, including helpfulness, social attraction, and social presence.

As previously mentioned, this study is to the best of our knowledge the first to comprehensively demonstrate the positive effect of physical embodiment in a SAR-guided healthcare scenario, wherein the SAR agent serves as both an instructor and active participant in the healthcare task with target users. Furthermore, our study contrasts with similar studies investigating the role of embodiment in HRI with older adults (e.g., Heerink et al., 2010) in that the participant evaluations in our study were based on multiple sessions of interaction with our SAR system. This is an important distinction, because as previously stated, the between-subjects comparison analysis was performed utilizing data collected following the fourth session of interaction instead of the first session in order to minimize the effect of novelty in the evaluations. However, the multiple-session study design was not only useful for attenuating the effect of novelty, but also for tracking user perceptions of the system and performance over time. For instance, it is interesting to note that although the survey results collected after the first session showed a participant preference for the physical robot over the virtual robot, the results were not statistically significant, in contrast to the results following the fourth session.

In addition, the physical robot group's ratings of the interaction and robot coach showed a positive trend across sessions, increasing by 6.2% (SD = 6.35) on average from the first session to the fourth session across all dependent measures. Conversely, the virtual robot group's ratings of the interaction and robot coach showed a negative trend across sessions, decreasing by 5.4% (SD = 6.64) on average from the first to the fourth session. The difference in the trends regarding participant ratings across sessions for the two study groups was found to be statistically significant ($t[14] = 3.56$, $p < .01$), showing a clear distinction between the study conditions regarding user perceptions of the system over time, and further demonstrating the positive influence of physical embodiment in SAR-guided interaction.

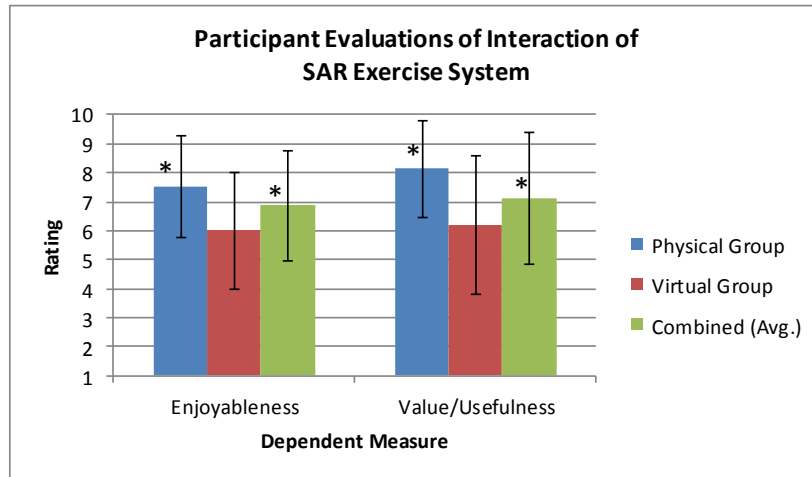
All of these results illustrate the importance of conducting multiple sessions of interaction for proper system evaluation and comparison, thus validating our study design.

5.2 SAR System Evaluation Results

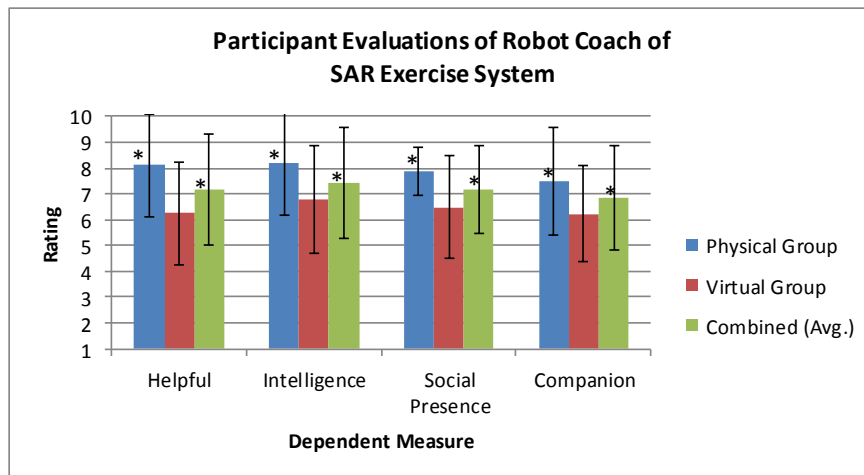
In order to evaluate the effectiveness of the SAR exercise system, we analyzed the combined data of the physical and virtual robot group's fourth session of interaction with the SAR exercise system. Therefore, the SAR system evaluation results regarding user perceptions and user exercise performance were gathered from all 33 older adult participants. The combined data were chosen for the final evaluation of the system for their embodiment-independent nature, and were constructed by averaging the evaluation ratings and user exercise performance of both study groups.

5.2.1 User Evaluations of the SAR System

To analyze the user evaluations of the SAR exercise system, we performed a two-tailed independent T-test assuming unequal variances to test for significant differences between the participant ratings of the subjective measures and a neutral evaluation rating. The neutral



(a)



(b)

Figure 6. (a) Plot of participant evaluations of the interaction of the SAR exercise system in terms of enjoyableness and usefulness; (b) plot of participant evaluations of the robot coach of the SAR exercise system in terms of helpfulness, intelligence, social presence, and as a companion. Note: Significant differences ($p < .05$) in comparison to neutral rating distribution are marked by asterisks (*).

evaluation rating distribution was obtained from a uniform sampling of the rating scale (integers from 1 to 10) for the approximate number of participants, and has a mean rating of 5.5 ($SD = 2.90$). This uniform sampling assumes no prior information regarding user perceptions of the system, and thus is deemed neutral. A plot of the interaction evaluation results, from each study group and combined, is provided in Figure 6(a).

The combined data of both study groups showed that participants evaluated the interaction with the SAR exercise system as enjoyable ($M = 6.8$, $SD = 1.90$) and valuable/useful ($M = 7.1$, $SD = 2.26$). The ratings for both measures were found to be significantly more positive than a neutral evaluation (enjoyableness: $t[49] = 2.16$, $p < .05$; usefulness: $t[55] = 2.46$, $p < .02$). These results illustrate the effectiveness of the system in promoting user intrinsic motivation, which is

characterized by enjoyment (Csikszentmihalyi, 1975), and in guiding the task-driven interaction toward achieving beneficial health outcomes for the user. Both of these factors show the successful application of our SAR design principles.

Regarding user perceptions of the SAR exercise system's robot coach, the combined data showed participants rated the robot highly and significantly more positive than neutral in terms of helpfulness ($M = 7.2$, $SD = 2.16$; $t[53] = 2.53$, $p < .02$), intelligence ($M = 7.4$, $SD = 2.15$; $t[53] = 2.98$, $p < .01$), social presence ($M = 7.2$, $SD = 1.71$; $t[46] = 2.71$, $p < .01$), and as a companion ($M = 6.8$, $SD = 2.02$; $t[51] = 2.09$, $p < .05$). The participants also rated the robot coach favorably in terms of social attraction ($M = 4.1$ (on a 7-point scale), $SD = 1.55$; $t[63] = 0.35$, *n.s.*) and as an exercise partner ($M = 6.5$, $SD = 2.26$; $t[55] = 1.43$, *n.s.*). These results illustrate that the participants perceived the robot coach as having a personable nature and being intelligent, both of which characteristics aid in the development of trust within the human-robot relationship and were design goals of our SAR system approach aimed at providing successful therapeutic interventions. A plot showing participant evaluations of the SAR system's robot coach is shown in Figure 6(b).

The results of the user evaluation of the SAR exercise system were very encouraging, as they showed a notable level of user acceptance of the system, as evidenced by the high ratings across each of the subjective measures, highlighting the effectiveness of our SAR system design principles.

5.2.2 User Exercise Performance Statistics

The collected statistics regarding participant performance in the exercise task were also very encouraging, as they demonstrated a consistently high level of user exercise performance and compliance. As previously stated, there were no significant differences in participant exercise performance between the two study groups; therefore, the results presented in this section represent the combined performance statistics of all participants, captured during the fourth session of interaction.

User compliance and performance in the Workout game were high. The average gesture completion time was 2.18 seconds ($SD = 0.54$) and the overall exercise performance averaged 5.07 seconds per exercise ($SD = 0.58$), which included time taken for verbal praise, feedback, and score reporting from the robot. The low percentage of necessary corrective feedback, averaging 5.2%, combined with zero failures and zero movement prompts during the interaction session, were all very important results, as they suggested that the participants were motivated to do well on the exercises consistently throughout the interaction. Furthermore, the quick pace of the exercise movements accomplished by the participants served to illustrate the fluid, highly interactive nature of the SAR-guided interaction; this is beneficial for increasing user intrinsic motivation and engagement in the task, and is the goal of our second SAR design principle.

The participant statistics for the remaining games also indicated high compliance and performance. A summary of all statistics regarding user performance, including those from the Sequence, Memory, and Imitation games, can be found in Table 3. As evidenced by the results, the average time to gesture completion in the Sequence and Memory games was greater than in the Workout game. This result was to be expected, as these two games incorporated a memory component, and during gameplay participants would often forget which gesture poses to complete in sequence. The increased feedback percentage in the Sequence game was also a result of this scenario, with the robot providing additional feedback and re-demonstrating the correct gestures to the users when necessary. It is important to note, however, that even though users had more difficulty in these two games, as expected, the participants were still able to achieve high performance (e.g., in the Sequence game, the participants required no feedback at all to successfully complete nearly 80% of the gestures). Other notable results include the low number of movement prompts during the Imitation game, averaging less than one per user, the high total number of exercises completed during the entire session ($M = 103.59$, $SD = 27.3$), and the low number of breaks taken by the users during interaction, all of which demonstrate a high level of participant intrinsic motivation to engage in the exercise games.

Table 3. User exercise performance statistics for all $n = 33$ older adult participants engaging with the SAR exercise system, showing means and standard deviations (in parentheses).

<i>Performance Measure</i>	<i>Avg.(std.)</i>
<i>Workout game:</i>	
Time to Gesture Completion (seconds)	2.18 (0.54)
Seconds per Exercise	5.07 (0.58)
Feedback Percentage	5.2% (4.9%)
Number of Failed Gestures	0
Number of Movement Prompts _w	0
<i>Sequence game:</i>	
Time to Gesture Completion (seconds)	5.95 (1.26)
Number of Sequences Completed	5.0 (1.41)
Number of Gesture Pairs Completed	15.1 (4.06)
Feedback Percentage	22.3% (11.8%)
<i>Memory game:</i>	
Maximum Score	6
Average Maximum Score	3.26 (1.31)
Time per Gesture Attempt (seconds)	7.64 (4.03)
<i>Imitation game:</i>	
Number of Movement Prompts ₁	0.64 (1.13)
<i>Entire Session:</i>	
Total Number of Exercises Completed	103.59 (27.3)
Number of Breaks Taken	1.08 (1.23)

5.2.3 Analysis of User Performance Results

Out of the fifteen user performance measures captured at each session, four were selected for further analysis: average completion time in the Workout game, percent feedback required in the Workout game, total exercises for the entire session, and average maximum score in the Memory game.

We began our analysis by seeking to determine whether the four measures might all be capturing a single underlying variable – user performance. As the measures use different scales (time, percentage, counts), each was standardized so that the mean was 0 and the standard deviation was 1. In addition, since the first two measures indicate better performance with a lower value and the last two with a higher value, the values of the first two measures were reversed by multiplying the standardized score by -1, so that higher values were uniformly indicative of better performance in all measures. Table 4 displays the intercorrelations of the four measures, along with the unstandardized means and standard deviations.

To evaluate performance across sessions, we performed a mixed models (multi-level) analysis of variance on the standardized scores of the measures within each session. Session is a within-subject factor. Condition, such as whether a physical or virtual robot led the task, is a between-groups measure. The effect of multiple measures on the same participants was controlled for. The analysis showed that, overall, there is no condition main effect, but there is a session main effect (performance improved over time), and there is a significant interaction of condition with session ($F[3, 465] = 4.0, p < .01$); see Table 5. The interaction effect arises from the fact that the physical robot prompted better performance in the first session but not in the later sessions. The contrast within the first session between the physical and virtual robot is $F[1, 49] = 3.98, p = 0.5$.

Table 4. Means, standard deviations, and intercorrelations among dependent measures.

Dependent Measure	Mean (Std. Dev.)	Completion Time	Feedback Percentage	Total Exercises	Maximum Score
Gesture completion time in seconds	2.54 (.89)	1.00			
Percentage of gestures where feedback was required	8.2% (7.9)	.91	1.00		
Total exercises performed in session	100 (18.6)	.78	.76	1.00	
Maximum score in Memory game	3.11 (1.35)	.24	.32	.50	1.00

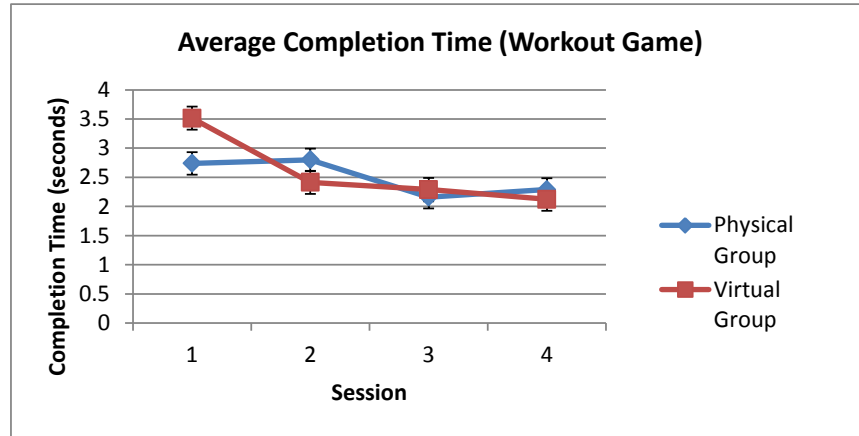
Table 5. Analysis of variance testing the fixed effect tests of condition, sessions, and measure on performance.

Variable	df	F ratio	p value
Condition (physical vs. virtual)	1, 31		
Session (1 – 4)	3, 465	44.5	< .0001
Condition * Session	3, 465	4.0	< .01
Measure (four)	3, 465		
Condition * Measure	3, 465		
Session * Measure	9, 465	2.3	= .01
Condition * Session * Measure	9, 465		

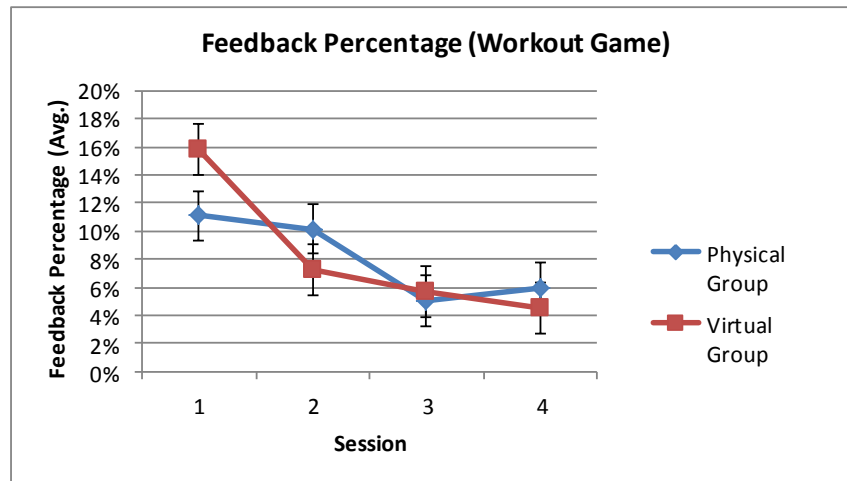
Analyzing the effects of each (unstandardized) performance measure showed that, in comparison with the virtual robot, the physical robot reduced completion time (interaction $F[3, 93] = 5.6, p < .01$) and feedback percentage (interaction $F[3, 93] = 2.9, p < .05$) in the first session but not in the remaining three sessions. Furthermore, according to a contrast, the difference between the physical and virtual robot in the first session is significant for completion time ($F[1, 82] = 7.8, p < .01$) and marginally significant for feedback percentage ($F[1, 81] = 3.3, p = .07$). The other measures do not show significant effects of condition; for example, total exercises increased notably over sessions in both conditions ($F[3, 93] = 31.99, p < .0001$). The plots for average completion time and feedback percentage are shown in Figures 7 (a) and (b), respectively.

6. Conclusions

We have presented a novel socially assistive robot system designed to motivate and engage elderly users in physical exercise in a seated aerobic exercise scenario. We detailed our SAR system approach and design methodology, including five SAR design principles derived from psychology research in the area of intrinsic motivation, which can be applied to a variety of human-robot interaction-based healthcare interventions. Implementation details of the system were also discussed, including the development of a general visual user activity recognition algorithm capable of detecting planar arm movements in real time, and procedures for providing task-based feedback, encouragement, and guidance while introducing variability to limit potential repetitiveness of the system and maintain user engagement.



(a)



(b)

Figure 7. Participant performance results across all four sessions of interaction for both study groups, showing (a) average gesture completion time (Workout game); (b) feedback percentage (Workout game). Note: Means are least squares means from the ANOVA; error bars are standard error. The statistical difference between the physical and virtual robot in the first session is significant for completion time ($p < .01$), and marginally significant for feedback percentage ($p = .07$), see text.

Furthermore, we presented a multi-session user study conducted with older adults to evaluate the effectiveness of our robot exercise system across a variety of user performance and evaluation measures. The results of the study validate our system approach and its effectiveness in motivating physical exercise in older adults; the participants engaged in physical exercise with high performance consistently throughout the interaction sessions, rated the SAR system highly in terms of enjoyableness and usefulness of the interaction, and rated the robot coach highly in terms of helpfulness, social attraction, social presence, and companionship.

In addition, we investigated the role of physical embodiment in our exercise system in the same multi-session user study by comparing the effectiveness of a physically embodied robot coach to that of a virtually embodied robot coach (a computer simulation of the same robot). Results from interaction with elderly participants ($n = 33$) were presented. The results of the

between-subjects embodiment comparison study show a strong user preference for the physical robot embodiment over the virtual robot embodiment in our SAR exercise system. Consistent with our stated hypotheses, participants reported the interaction with the physical robot as being more enjoyable and more valuable/useful than the interaction with the virtual robot. Furthermore, the participants evaluated the physical robot as more helpful, more socially attractive, and as having greater social presence than the virtual robot.

The key contributions of this work include the following: 1) the approach, design methodology, and implementation details of a fully autonomous socially assistive robot system designed to engage elderly users to perform physical exercise while providing real-time feedback, guidance, and encouragement; 2) a set of five general design principles for socially assistive systems to aid users in a healthcare task by influencing their intrinsic motivation to engage in the therapeutic intervention; 3) validation of the SAR system approach and its effectiveness in gaining user acceptance and motivating elderly users to engage in consistent physical exercise through a multi-session user study conducted with the target healthcare population; 4) the first comprehensive physical vs. virtual robot embodiment comparison study conducted to observe the effects of embodiment within a SAR-guided healthcare scenario, wherein the SAR agent serves as both an instructor and an active participant in the healthcare task with target users; and 5) demonstration of the positive effect of physical embodiment in a fluid, real-time human-robot interaction based healthcare intervention with elderly users and how embodiment can affect user perceptions across a variety of interaction metrics.

The acceptance of the SAR exercise system by elderly users, as evidenced by the high participant evaluations of the usefulness and enjoyableness of the system and consistent exercise performance in our multi-session user study, is very encouraging and illustrates the ability of SAR technology to motivate elderly users to engage in fruitful therapeutic tasks towards achieving beneficial health outcomes and improving their quality of life.

Acknowledgements

We are grateful to the be.group (formerly Southern California Presbyterian Homes) staff and participating residents, to Aras Akbari and Farrokh Langroodi for collecting the study data, and to Kwan Min Lee for providing insights into the study design and survey questions. We thank Sara Kiesler for providing numerous helpful suggestions, insightful feedback, and guidance in the data analysis. This work was supported by a grant from the Robert Wood Johnson Foundation's Pioneer Portfolio through its national program, "Health Games Research: Advancing Effectiveness of Interactive Games for Health," and the National Science Foundation under Grants IIS-0713697, CNS-0709296, and IIS-1117279.

References

- American Association of Colleges of Nursing. (2010). *Nursing shortage fact sheet* [Fact sheet].
- American Health Care Association. (2008). *Summary of 2007 AHCA survey nursing staff vacancy and turnover in nursing facilities* [Report].
- Bainbridge, W. A., Hart, J. W., Kim, E. S., & Scassellati, B. (2011). The benefits of interactions with physically present robots over video-displayed agents. *International Journal of Social Robotics*, 3(1), 41–52. doi:10.1007/s12369-010-0082-7
- Balakrishnan, V., & Yeow, P. (2007). Texting satisfaction: Does age and gender make a difference? *International Journal of Computer Science and Security*, 1(1), 85–96.
- Bartneck, C. (2003). Interacting with an embodied emotional character. In *Proceedings of the 2003 International Conference on Designing Pleasurable Products and Interfaces* (pp. 55–60). New York: ACM.

- Baum, E., Jarjoura, D., Polen, A., Faur, D., & Rutecki, G. (2003). Effectiveness of a group exercise program in a long-term care facility: A randomized pilot trial. *Journal of the American Medical Directors Association, 4*(2), 74–80. doi:10.1016/S1525-8610(04)70279-0
- Bickmore, T. W., & Picard, R. W. (2005). Establishing and maintaining long-term human-computer relationships. *ACM Transactions on Computer-Human Interaction, 12*(2), 293–327. doi:10.1145/1067860.1067867
- Bickmore, T., Schulman, D., & Yin, L. (2010). Maintaining engagement in long-term interventions with relational agents. *International Journal of Applied Artificial Intelligence, 24*(6), 648–666.
- Buerhaus, P. (2008). Current and future state of the US nursing workforce. *JAMA, 300*, 2422. doi:10.1001/jama.2008.729
- Centers for Disease Control and Prevention. (2003). *Morbidity and mortality weekly report*. Department of Health and Human Services.
- Colcombe, S., & Kramer, A. (2003). Fitness effects on the cognitive function of older adults. *Psychological Science, 14*, 125. doi:10.1111/1467-9280.t01-1-01430
- Colcombe, S., Kramer, A., Erickson, K., Scalf, P., McAuley, E., Cohen, N., ... & Elavsky, S. (2004). Cardiovascular fitness, cortical plasticity, and aging. *Proceedings of the National Academy of Sciences of the United States of America, 101*(9), 3316–3321. doi:10.1073/pnas.0400266101
- Csikszentmihalyi, M. (1975). *Beyond boredom and anxiety*. San Francisco: Jossey-Bass.
- Csikszentmihalyi, M. (1993). *The evolving self: A psychology for the third millennium*. New York: HarperCollins.
- Dawe, D., & Moore-Orr, R. (1995). Low-intensity, range-of-motion exercise: Invaluable nursing care for elderly patients. *Journal of Advanced Nursing, 21*, 675–681. doi:10.1046/j.1365-2648.1995.21040675.x
- Deci, E., & Ryan, R. (1985). *Intrinsic motivation and self-determination in human behavior*. New York: Plenum Press.
- Dienstbier, R. A., & Leak, G. K. (1976). *Effects of monetary reward on maintenance of weight loss: An extension of the overjustification effect*. Paper presented at the American Psychological Association Convention, Washington, D.C.
- Dubowsky, S., Genot, F., Godding, S., Kozono, H., Skwersky, A., Yu, H., & Shen Yu, L. (2000). PAMM - a robotic aid to the elderly for mobility assistance and monitoring. *Proceedings of the IEEE International Conference on Robotics and Automation, 1*, 570–576. doi:10.1109/ROBOT.2000.844114
- Fasola, J., & Matarić, M. J. (2010a). Robot motivator: Increasing user enjoyment and performance on a physical/cognitive task. *Proceedings of the IEEE International Conference on Development and Learning, 274–279*. doi:10.1109/DEVLRN.2010.5578830
- Fasola, J. & Matarić, M. J. (2010b). Robot exercise instructor: A socially assistive robot system to monitor and encourage physical exercise for the elderly. *19th IEEE International Symposium in Robot and Human Interactive Communication, 416–421*.doi:10.1109/ROMAN.2010.5598658
- Fasola, J. & Matarić, M. J. (2012). Using socially assistive human-robot interaction to motivate physical exercise for older adults. In T. Kanade (Ed.), *Proceedings of the IEEE, Special Issue on Quality of Life Technology, 100*(8), 2512–2526. doi:10.1109/JPROC.2012.2200539
- Feil-Seifer, D. J., & Matarić, M. J. (2009). Towards the integration of socially assistive robots into the lives of children with ASD. *Paper presented at the International Conference on Human-Robot Interaction Workshop on Societal Impact: How Socially Accepted Robots Can be Integrated in our Society*, San Diego, CA.

- Feil-Seifer, D. J., & Matarić, M. J. (2010). Using proxemics to evaluate human-robot interaction. *Proceedings of the International Conference on Human-Robot Interaction*, 143–144. doi:10.1109/HRI.2010.5453225
- Fisher, C. D. (1978). The effects of personal control, competence, and extrinsic reward systems on intrinsic motivation. *Organizational Behavior and Human Performance*, 21, 273–288. doi:10.1016/0030-5073%2878%2990054-5
- French, B., Tyamagundlu, D., Siewiorek, D., Smailagic, A., & Ding, D. (2008). Towards a virtual coach for manual wheelchair users. *Proceedings of International IEEE Symposium of Wearable Computers*, 77–80. doi:10.1109/ISWC.2008.4911589
- Fujiyoshi, H., & Lipton, A. (1998, October). Real-time human motion analysis by image skeletonization. *Proceedings of the Workshop on Applications of Computer Vision*. doi:10.1109/ACV.1998.732852
- George, L., Blazer, D., Hughes, D., & Fowler, N. (1989). Social support and the outcome of major depression. *The British Journal of Psychiatry*, 154, 478–485. doi:10.1192/bjp.154.4.478
- Heerink, M., Kröse, B., Evers, V., & Wielinga, B. (2010). Assessing acceptance of assistive social agent technology by older adults: The Almere Model. *International Journal of Social Robotics*, 2(4), 361–375.
- Jenkins, O. C., Chu, C., & Matarić, M. J. (2004). *Nonlinear spherical shells for approximate principal curves skeletonization*. (University of Southern California Center for Robotics and Embedded Systems Technical Report, CRES-04-004).
- Jung, Y., & Lee, K. M. (2004). Effects of physical embodiment on social presence of social robots. *Proceedings of Presence*, 2004, 80–87.
- Kang, N., & Yoon, W. (2008). Age- and experience-related user behavior differences in the use of complicated electronic devices. *International Journal Human-Computer Studies*, 66(6), 425–437. doi:10.1016/j.ijhcs.2007.12.003
- Kidd, C., Taggart, W., & Turkle, S. (2006). A sociable robot to encourage social interaction among the elderly. *International Conference on Robotics and Automation*, 3972–3976. doi:10.1109/ROBOT.2006.1642311
- Kidd, C. D., & Breazeal, C. (2008). Robots at home: Understanding long-term human-robot interaction. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 3230–3235. doi:10.1109/IROS.2008.4651113
- Lee, K. M. (2004). Presence, explicated. *Communication Theory*, 14(1), 27–50. doi:10.1093/ct/14.1.27
- Matarić, M. J., Eriksson, J., Feil-Seifer, D. J., & Winstein, C. J. (2007). Socially assistive robotics for post-stroke rehabilitation. *Journal of NeuroEngineering and Rehabilitation*, 4(5). doi:10.1186/1743-0003-4-5
- Matsusaka, Y., Fujii, H., Okano, T., & Hara, I. (2009). Health exercise demonstration robot TAIZO and effects of using voice command in robot-human collaborative demonstration. *The 18th IEEE International Symposium on Robot and Human Interactive Communication*, 472–477. doi:10.1109/ROMAN.2009.5326042
- McCroskey, J. C., & McCain, T. A. (1974). The measurement of interpersonal attraction. *Speech Monographs*, 41, 261–266. doi:10.1080/03637757409375845
- McMurdo, M. D., & Rennie, L. M. (1993). A controlled trial of exercise by residents of old people's homes. *Age and Ageing*, 22, 11. doi:10.1093/ageing/22.1.11
- Moak, Z. B., & Agrawal, A. (2010). The association between perceived interpersonal social support and physical and mental health: Results from the national epidemiological survey on alcohol and related conditions. *Journal of Public Health*, 32, 191–201.

doi:10.1093/pubmed/fdp093

Montemerlo, M., Pineau, J., Thrun, S., & Varma, V. (2002). Experiences with a mobile robotics guide for the elderly. *Proceedings of the AAAI National Conference on Artificial Intelligence*, 587–592.

Paykel, E. (1994). Life events, social support and depression. *Acta Psychiatrica Scandinavica*, 89, 50–58. doi:10.1111/j.1600-0447.1994.tb05803.x

Poresky, R. H., Hendrix, C., Mosier, J. E., & Samuelson, M. (1987). Companion animal bonding scale: Internal reliability and construct validity. *Psychological Reports*, 60, 743–746. doi:10.2466/pr0.1987.60.3.743

Powers, A., Kiesler, S., Fussell, S., & Torrey, C. (2007). Comparing a computer agent with a humanoid robot. *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction* (pp. 145–152). New York: ACM. doi:10.1145/1228716.1228736

Spirduso, W., & Clifford, P. (1978). Replication of age and physical activity effects on reaction and movement time. *Journal of Gerontology*, 33, 26.

Stansfeld, S., Rael, G., Head, J., Shipley, M., & Marmot, M. (1997). Social support and psychiatric sickness absence: A prospective study of British civil servants. *Psychological Medicine*, 27(1), 35–48. doi:10.1017/S0033291796004254

Stice, E., Ragan, J., & Randall, P. (2004). Prospective relations between social support and depression: Differential direction of effects for parent and peer support? *Journal of Abnormal Psychology*, 113, 155–159. doi:10.1037/0021-843X.113.1.155

Tapus, A., Tapus, C., & Matarić, M. J. (2008). User-robot personality matching and robot behavior adaptation for post-stroke rehabilitation therapy. *Intelligent Service Robotics*, 1(2), 169–183. doi:10.1007/s11370-008-0017-4

Tapus, A., Tapus, C., & Matarić, M. (2009). The use of socially assistive robots in the design of intelligent cognitive therapies for people with dementia. *International Conference on Rehabilitation Robotics* (pp. 924–929). IEEE: Kyoto, Japan. doi:10.1109/ICORR.2009.5209501

Thomas, V., & Hageman, P. (2003). Can neuromuscular strength and function in people with dementia be rehabilitated using resistance-exercise training? Results from a preliminary intervention study. *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences*, 58, M746. doi:10.1093/gerona/58.8.M746

Tway, D. C. (1994). *A construct of trust* (Unpublished dissertation). The University of Texas at Austin.

Vallerand, R. J. (1983). Effect of differential amounts of positive verbal feedback on the intrinsic motivation of male hockey players. *Journal of Sport Psychology*, 5, 100–107.

Vallerand, R. J., & Reid, G. (1984). On the causal effects of perceived competence on intrinsic motivation: A test of cognitive evaluation theory. *Journal of Sport Psychology*, 6, 94–102.

Wada, K., Shibata, T., Saito, T., & Tanie, K. (2002). Analysis of factors that bring mental effects to elderly people in robot assisted activity. *Proceedings of the International Conference on Intelligent Robots and Systems 2*, 1152–1157. doi:10.1109/IRDS.2002.1043887

Wainer, J., Feil-Seifer, D. J., Shell, D. A., & Matarić, M. J. (2006). The role of physical embodiment in human-robot interaction. *IEEE Proceedings of the International Workshop on Robot and Human Interactive Communication*, 117–122. doi:10.1109/ROMAN.2006.314404

Wainer, J., Feil-Seifer, D. J., Shell, D. A., & Matarić, M. J. (2007). Embodiment and human-robot interaction: A task-based perspective. *IEEE Proceedings of the International Workshop on Robot and Human Interactive Communication*, 872–877. doi:10.1109/ROMAN.2007.4415207

Waldherr, S., Thrun, S., Romero, R., & Margaritis, D. (1998). Template-based recognition of pose

and motion gestures on a mobile robot. *Proceedings of the National Conference on Artificial Intelligence*, 977–982.

Weinberg, R. S., & Ragan, J. (1979). Effects of competition, success/failure, and sex on intrinsic motivation. *Research Quarterly*, 50, 503–510.

Wren, C. R., Azarbayejani, A., Darrell, T., & Pentland, A. P. (1997). Pfinder: Realtime tracking of the human body. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7), 780–785. doi:10.1109/34.598236

J. Fasola, Interaction Lab, University of Southern California, Los Angeles, CA, USA. Email: fasola@usc.edu; M. J. Matarić, Interaction Lab, University of Southern California, Los Angeles, CA, USA. Email: mataric@usc.edu