

Socially Assistive Robotics for Guiding Motor Task Practice

Eric Wade^{1,*}, Avinash Parnandi^{2,†},
Ross Mead^{1,‡}, Maja Mataric^{1,§}

¹ University of Southern California,
3710 McClintock, 90089, Los Angeles,
USA

² Texas A&M University,
77843 College Station, USA

Received 2011/11/01

Accepted 2012/03/01

Abstract

Due to their quantitative nature, robotic systems are useful tools for systematically augmenting human behavior and performance in dynamic environments, such as therapeutic rehabilitation settings. The efficacy of human-robot interaction (HRI) in these settings will depend on the robot's coaching style. Our goal was to investigate the influence of robot coaching styles designed to enhance motivation and encouragement on post-stroke individuals during motor task practice. We hypothesized that coaching styles incorporating user performance and preference would be preferred in a therapeutic HRI setting. We designed an evaluation study with seven individuals post stroke. A socially assistive robotics (SAR) system using three different coaching styles guided participants during performance of an upper extremity practice task. User preference was not significantly affected by the different robot coaching styles in our participant sample ($H(2) = 2.638, p = 0.267$). However, trends indicated differences in preference for the coaching styles. Our results provide insights into the design and use of SAR systems in therapeutic interactions aiming to influence user behavior.

Keywords

socially assistive robotics · human-robot interaction · neurorehabilitation · human monitoring · motor learning · stroke

1. Introduction

Opportunities for interactions between humans and intelligent systems are becoming more pervasive. At the same time, changes in population demographics in the developing world call for the intelligent application of sensor and robotic systems to fields such as healthcare and pervasive monitoring. The multidisciplinary field of human-robot interaction (HRI) studies the broad challenges associated with placing people and robots in shared settings. Socially assistive robotics (SAR), has emerged as a promising and growing area of HRI that uses robotics for the provision and administration of motivation, encouragement, and rehabilitation for those suffering from cognitive, motor, and social deficits [1]. SAR focuses on hands-off interactions with robots in therapeutic settings. Using embodiment, verbal and non-verbal communication, personality, emotion, user models, socially situated learning, and intentionality, but no physical contact with the user, SAR systems can guide interactions with the user so as to achieve various desired outcomes [2]. In this paper we investigate the role of the coaching style in a therapeutic SAR interaction. The aim of this study was to determine if robot coaching styles, based on current literature, would influence self-reported outcomes during a robot-guided task practice session. Our primary hypothesis was that robot coaching conditions in a therapeutic HRI setting that utilize information about participant performance and input would be rated more favorably than a coaching condition not utilizing these features. Because of the focus on the therapeutic setting, this primary hypothesis required a secondary hypothesis: that the practice task we used can be considered a motor learning task. To validate the secondary hypothesis, the performance data must re-

flect what would be expected in a motor practice task: as difficulty increases, performance decreases. We outline the relevant interactive characteristics of our system design, and describe the tools used to personalize the system to the performance of individual users. We describe an evaluation study in which individuals performed a motor task practice game while being instructed, monitored, and guided by a socially assistive humanoid robot. We describe how the robot produces coaching feedback based on user actions, its own speech, and practice task conditions. We investigate the quality of the interaction based on two metrics: self-reported outcomes and user task performance. The rest of this paper is organized as follows. Section 2 provides background on assistive rehabilitation and socially assistive robotics. In Section 3 we describe the application of our system to the problem domain. In Section 4, we detail the implementation of our experimental setup and describe the study procedure. In Section 5 we present the study results and outcomes, and in Section 6, we discuss the implications of our results.

2. Background

The work described in this paper draws on established conventions from the fields of assistive rehabilitation robotics and HRI. This section therefore focuses only on background specifically relevant to our SAR system. The interested reader is referred to reviews by Fong and Tan [2, 3].

The general approach in the field of assistive rehabilitation robotics has focused on orthoses that physically interact with individuals with motor deficits. This includes lower extremity (LE) devices such as the LOKOMAT[®] [4] and ALEX (Active Leg EXoskeleton) [5]. Upper extremity (UE) devices are typically contact-based; they measure and apply forces and torques to the user's arm to assess or encourage specific motor task practice. Examples include the ARM Guide, developed by Reinkensmeyer et al. [6]. The maturity of the field is indicated by the fact that such devices have been tested in clinical trials. Lo et al. utilized

*E-mail: ericwade@usc.edu

†E-mail: parnandi@cse.tamu.edu

‡E-mail: rossmead@usc.edu

§E-mail: mataric@usc.edu




Figure 1. Socially assistive robotics architecture, adapted from Mead et al. [18].

the MIT Manus robot in a large scale clinical trial [7], and demonstrated that their robot could produce clinical outcomes in patients that were similar to those obtained using dose equivalent human-administered therapies.

The field of socially assistive robotics is newer than assistive rehabilitation robotics, and is separate from it in its methods. SAR is defined as the provision of assistance through social (not physical) interactions with robots [1]. A SAR system uses non-contact feedback, coaching, and encouragement to guide a user during the performance of a task. SAR systems can demonstrate task goals, monitor the user, and provide augmented performance feedback. The lack of physical contact means there are a minimum of safety concerns. Further advantages include lower costs and increased accessibility (relative to contact-based robot systems). A socially assistive interaction can take place in a multitude of environments, including the laboratory, the clinic, or the user's home.

SAR has shown promise in a number of domains, including tutoring, emotional expression, social skill training, daily life assistance, and physical therapy [1, 8–10]. Our research has demonstrated the efficacy of SAR in elder care [11, 12], socialization of children with autism spectrum disorder [13], stroke assessment [14, 15], and stroke rehabilitation [16].

SAR systems are particularly promising for stroke rehabilitation. Stroke affects a large percentage of the population worldwide, with over 9 million people affected annually, many of whom go on to live with motor disabilities [17]. These disabilities often include some level of paralysis, with one side being more affected than the other, and result in difficulties with gait, speech, and upper extremity task performance.

Individuals often compensate for upper extremity deficits by using the less-affected (non-paretic) arm, and voluntarily suppressing the use of the more-affected (paretic) arm. To encourage use (and recovery) of the paretic arm, therapeutic interventions for stroke typically consist of intense one-on-one practice with a trained clinician, focusing on specific real-world tasks modeled on activities of daily living (ADLs). Such task-specific practice can lead to recovery from motor deficits. A key advantage of non-contact SAR in the stroke population is that individuals can use the stroke-affected limb in the types of meaningful, unconstrained, functional tasks that are encountered in daily life, and can practice in the home setting. Applying SAR to this problem domain requires an understanding of how people with stroke respond to a robot in therapeutic settings. In the following, we describe the application of an SAR framework to the domain of stroke rehabilitation, and our method for determining how the robot's coaching style influences participant responses.

3. System Description and Application to the Problem Domain

We first orient the reader to a description of our system by describing the SAR architecture, and the relevance of its features (embodiment, presence, and communication modalities) to our research question.

3.1. Architecture

The physical setup for our SAR system consisted of a one-on-one interaction between the participant and a humanoid robot, with additional sensors used to monitor the participant and the environment (See Figure 1). A SAR control architecture enables the robot to guide the user through a series of tasks; sensor data are processed and, upon recognizing a particular event or pattern within the context of the task, the robot provides the appropriate social feedback action. This design allows for a dynamic interaction that can respond in real-time to changes in user performance.

In [18], we identified and accommodated three main objectives of a SAR control architecture for post-stroke rehabilitation: (1) an ability to accommodate varied ADL-inspired tasks without significant reconfiguration, (2) the provision of real-time task-dependent feedback to the patient, and (3) an ability to monitor user activity over time to provide more meaningful and more personal interaction behaviors. Considering the variability of inputs, tasks, and interaction modalities required for monitoring ADLs, we determined that a distributed control architecture would be beneficial for SAR scenarios. Robot Operating System (ROS; <http://www.ros.org/wiki/>) is an open-source framework for device intercommunication, utilizing distributed nodes interacting via a publish/subscribe paradigm to provide ease of use and portability [19]. Using the ROS framework, we developed collections of distributed nodes (referred to henceforth as “activity servers”) for various ADL tasks. Each activity server was responsible for independently monitoring a component of the state of the robot (e.g., joint angles), the user (e.g., motor activity), and the world (e.g., locations of objects); activity servers were also implemented to command robot verbal and nonverbal output (e.g., speech and gesture, respectively). Subsystem independence was crucial, so as to ensure that one part of the system would not cause the entire interaction to break down. In a pilot study, we observed the architecture’s ability to continue functioning despite subsystem failure on a number of occasions [18]. While some redundancy of state monitoring is required, the key observation is that the architecture is capable of maintaining the interaction in spite of individual subsystem failures.

Further details of the design and implementation of the robot control architecture are given in Mead et al. [18].

3.2. Embodiment and Presence

Physical embodiment plays a critical role in the efficacy of SAR interactions. In Fong et al., embodiment is defined as the basis for structural coupling by creating the potential for mutual perturbations between systems and their environment [2]. This is a broad definition; it includes the physical makeup of the interactive device, the presence of the device, and the physical arrangement of all actors in the interaction. In Johnson et al., it is pointed out that one of the key components of any personalized interface is the physical interface (e.g., the device and its physical settings) [20]. The physical size, shape, form, and presence of a robot affect both conscious and subconscious human reactions. Various studies, including our own, have attempted to systematically examine the role of embodiment.

Kidd et al. sought to investigate the extent to which embodiment played a role in adherence to a weight loss program [21]. In their experiment, participants were randomly assigned to three groups: the first group tracked their adherence to a weight loss program using a pen-and-paper log; the second group entered their progress into a laptop computer; the third group entered their progress into a humanoid robot interface. The researchers found that participants persisted for the least amount of time with the pen-and-paper log. They persisted longer with the laptop, and persisted statistically



Figure 2. The USC Interaction Laboratory’s Bandit humanoid robot.

significantly longer with the embodied robot. All users were logging the same types of information, so this demonstrates that the robot embodiment was more engaging than the computer, and led to more persistent adherence. Importantly, this was fundamentally an embodiment study, since the robot itself did not provide any added task-specific abilities compared to the computer. The embodiment effect of the robot in this study is due to the robot’s presence, since its behavior was minimal.

In another study, Wainer et al. investigated the role of embodiment in an agent-assisted interactive game. Participants played the classic Towers of Hanoi game while assisted by a non-humanoid, co-located robot, a remote-located robot viewed on a monitor, and a software simulation of a robot viewed on a monitor [22]. Participants felt that the co-located robot was significantly more watchful than both the remote-located robot and the simulation. It was also seen as more and significantly more enjoyable than the remote and simulated robots, respectively. Since all three systems had identical task-related function, the results can be attributed to embodiment.

Looije et al. refer to evidence of “higher appreciation” for technology that has a physical presence [23]. Their experiment sought to elicit differences in perception of a text interface, a software simulation, and an embodied robot, and found that an embodied robot may be perceived as a more “intelligent” interface, which is necessary in a persuasive setting. The authors describe the importance of natural, human-like cues including gaze, gesture, and posture. The implication is that, in an interactive, persuasive scenario, gaze, gestures, robot postures, and actions should be as natural as possible.

A recent study by Fasola and Mataric compared the responses of 66 participants across the age spectrum (half elderly over 65 years of age, half younger adults) to two types of seated exercise coaches: a SAR system and a computer simulation of the same SAR system [24]. The results indicated a statistically significant user preference for the physical robot over the computer. The participants also ranked the robot as more enjoyable, entertaining, valuable/useful, and more helpful/useful. Here again, the two embodiments (robot and computer) had identical task functions, providing further insight into the embodiment effect.

The robot used in that study was the same as the one we used in our experiments. Bandit, a humanoid torso mounted on a commercial mobile base, was developed at the USC Interaction Laboratory in collaboration with BlueSky Robotics (Figure 2). The Interaction Lab has used Bandit in a variety of SAR interactions to date, including with elderly users in seated exercises [11], children with autism spectrum disorders [1, 13], users with Alzheimer’s Disease playing a cognitive music game [12], and stroke patients [18].

Bandit’s humanoid robot torso weighs 15 pounds, is 16in. high, 14in. wide (across the chest) and 5in. thick, with 14in. long arms. Each arm has 7 degrees-of-freedom (DOF), including the 1 DOF gripper. There are 16 DOF in the torso, including the 2 DOF in the neck/head.

The head of the robot has 1 DOF expressive eyebrows, and a 3 DOF expressive mouth. The humanoid has been judged to be engaging by technical and non-technical users alike, has been assessed as a non-significant threat device, and has been repeatedly approved by the USC Internal Review Board.

In the therapeutic motor task practice setting of our experiment, Bandit interacted with the participants one-on-one through speech, gesture, head movements, and lip movements, but provided no physical demonstration of the practice task. Our study did not involve a comparison with a computer screen for two reasons. First, such embodiment studies have been conducted in the past, as discussed in the previous section. Second, in the domain of stroke rehabilitation, it is known that the engagement of the participant in the task is particularly important for cortical reorganization and recovery [25]. Since our and other work has indicated that a physically embodied robot is preferable to and more engaging than an on-screen agent, this study focused not on evaluating embodiment yet again, but on evaluating coaching styles within the same robot embodiment. Thus, in testing our hypotheses regarding robot coaching styles, we utilized the already validated engaging robot embodiment (Bandit) and enabled it to provide different coaching styles.

3.3. Verbal Interaction and Feedback

The coaching style is defined primarily by the mechanism of communication employed by the robot. In general, it is accepted that speech is highly effective for communicating emotion [2], whether making use of traditional structure and syntax or through “expressive utterances”. Johnson mentions that, in addition to physical interfaces, the communication interface (e.g., software, monitoring capability) is one of the two key components to personalization in SAR systems [20].

It has also been shown that verbalization is an important mechanism in establishing common ground in a social interaction [26]. Common ground is a theory of human-human interaction that assumes communication between people requires coordination in order to meet mutual understanding. In a therapeutic setting, this involves convincing the participant that both parties have a shared goal in the interaction. While this can be conveyed explicitly, is often more effectively done through subtler means. It is also user-specific. For instance, Kiesler cites the use of “elder speak” in therapeutic settings, noting that a robot designed for use with an older population may need to take advantage of such modified speaking styles to be persuasive [26]. Looije points out that a population’s lack of familiarity with technology may mean that verbal/communicative interfaces are most natural of untrained users [23]. For instance, the relatively limited experience with computers of the elderly populations means social communication may play an even stronger role in interactions.

The content of verbal communication can be both explicitly and implicitly motivating. An established verbal motivation mechanism used in motor control and learning is augmented feedback [27]. Augmented feedback is defined as feedback about performance not normally available during the performance of an activity. This feedback, when provided at the conclusion of a task, is known as knowledge of results (KR). KR is known to have a motivational impact in task practice; it keeps the interaction interesting by varying the invoked sensory modalities [27]. For the same reasons, it can keep the participant alert during a sometimes repeated, monotonous task and it allows participants to set personal performance goals. Explicit provision of past performance results can lead to an implicit desire to perform better. We were primarily concerned with the motivational nature of KR and use it to keep participants focused on the practice task.



Figure 3. A study participant practicing the wire puzzle task while guided by the socially assistive humanoid robot.

In addition to KR, Bandit provided congratulatory feedback according to a collection of pre-defined scripts. The verbal feedback was accompanied by non-verbal communication (e.g., gestures and head turns). When providing this congratulatory feedback, Bandit nodded its head. The amount of nodding was determined using the NonVerbal Behavior Generator (NVBG) rules [28], which relate gestures to speech content, and are based on specific keywords/structures and behavior priorities learned from user activity modeling. We determined the duration of the entire length of a phrase (stored as an audio file) and generated an overall “phrase valence” which dictated the duration and type of non-verbal behavior exhibited by the robot for that particular phrase utterance. The overall valence of a phrase is determined based on all of the NVBG rules that have been applied to that phrase. Using the NVBG rules, the weighted valence V^* of a phrase depends on the number of words n_i , their individual valence V_i , and the word priority P_i (for a more detailed description, see [18]).

$$V^* = \frac{\sum \frac{n_i V_i}{P_i}}{\sum n_i} \quad (1)$$

Bandit nodded its head with a valence-proportional amplitude at regular intervals throughout the congratulatory phrase. Thus, a participant who performed well received more positive feedback statements (and more nodding) from the robot than one who performed less well.

4. Methods

In the following, we describe our experimental design, including the participant demographics, the practice task, the robot coaching styles, the study schedule, and the outcome measures.

4.1. Participants

To validate our SAR system, we recruited seven individuals in the chronic phase (>6 months) of stroke recovery. The main inclusion criteria were: (1) concordance (right-side affected, right-side dominant pre-morbidity), (2) the ability to lift a full soda can from a desk for 5 seconds with the stroke-affected limb, and (3) transportation to and from the experiment site. The cohort consisted of four males and

Table 1. Demographics of the seven participants. For each participant, the age, dominant limb, stroke-affected side, years after stroke, and upper extremity Fugl-Meyer Score are reported.

	Age	Dom.	Par.	Yrs. after stroke	UE FMA
Participant 1	66	R	R	7	48
Participant 2	41	R	R	3	37
Participant 3	73	R	R	9	37
Participant 4	68	R	R	6	51
Participant 5	46	R	R	10	56
Participant 6	47	R	R	6	46
Participant 7	60	R	R	7	49

three females, and ranged from 41 to 73 years of age. For each participant, the age, dominant limb, stroke-affected (paretic) side, years after stroke, and upper extremity Fugl-Meyer Assessment (FMA) score are reported (Table 1). We used the FMA to quantify participant impairment levels [29]. The FMA is an impairment assessment wherein participants perform a battery of tasks, and are rated by a clinician on their ability to perform the task on a 0 to 2 point scale, with 66 points as a maximum. The participants in our study (FMA scores between 44 and 66) would be considered to have mild-to-moderate functional impairments, or mid- to high-level ability to perform tasks with the upper extremity. A trained physical therapist administered the FMA.

4.2. Practice Task

For the practice task, participants performed a wire puzzle game while being instructed, guided, and motivated by Bandit. To complete the game, participants guided a wand with a ring at its tip along a wire puzzle without making physical contact (Figure 3). Participants were instructed to perform as accurately and quickly as possible. Such wire puzzles have been used as a measure of motor performance, but not with the stroke-affected population [30, 31]. From session to session, the puzzle difficulty increased; difficulty was parameterized according to puzzle length and the number of 90° turns in the puzzle wire path. Within a session, task difficulty was augmented by changing the wand. The wand with the largest ring diameter was the most easiest, while the wand with the smallest ring diameter was the most challenging (See Figure 4). The ring diameters were 3 cm, 6 cm, 9 cm, and 12 cm (for Wands 4, 3, 2, and 1 respectively). While participants performed the task, the robot monitored the number of errors (contacts between the wand ring and the puzzle) and total movement time (the time to go from the start to the end, and back to the start again). These were then provided verbally to the participant as KR after completing a bout. The success of a bout was determined by the number of errors and the movement time on each wand/puzzle combination.

4.3. Procedure

The study consisted of three sessions separated by at least 24 hours each. Each session consisted of three 15 min task practice blocks, separated by 10 min breaks. Within each block, participants performed multiple bouts of the puzzle, where a bout is defined as a going from the start pad to the end pad, and back to the start pad again (See Figure 4). At the start of the first session, we obtained informed consent, administered the FMA, and administered a questionnaire regarding the participant's confidence in his/her upper limb use. The




Figure 4. The wire puzzle setup; decreasing wand diameters correspond to increasing levels of task difficulty.

participants then performed three 15 min task practice blocks. In each session, participants saw all three robot conditions (one per 15 min practice block); more details about the conditions are found in the following section. At the beginning and end of each 15 min practice block, we obtained self-reported fatigue scores using a 10-point visual analog scale. At the end of the session, we administered an exit questionnaire regarding the robot's performance according to the different coaching styles.

The second session was similar to the first, with the exclusion of the consenting, motor assessment, and confidence measure. Again, the participants performed three 15 min practice blocks with 10 min breaks. We obtained fatigue measures, and administered an exit questionnaire at the conclusion of the session; the same questionnaire was used in both the first and second session.

In the third session, as in the second, participants performed three 15 min practice blocks with 10 min breaks. We obtained fatigue measures, and again administered an exit questionnaire at the conclusion of the session. The exit questionnaire for the third sessions contained the same questions from the first two sessions, with additional questions regarding the overall quality of the interaction.

4.4. Robot Coaching Styles

To test our hypothesis regarding robot coaching styles, we developed three coaching styles based on recent literature indicating promising ways to keep participants engaged in rehabilitation activities. Literature states that training should be challenging, and that the trainee should be incorporated in the decision making process [32]. Seeking to test this first concept of challenge, we developed a style based on the challenge point framework (CPF) [33]. This theory implies that, during learning, there are interactions between participant skill level, conditions of practice, and inherent task difficulty. It states that the "challenge level," or the level at which a participant is learning optimally, changes as participants improve at a task. Thus, we developed a style that would increment or decrement difficulty depending on performance. This coaching style is the Challenge style. If participants had 2 successful trials at a given level, they were advanced to the next hardest wand. If they spent more than 3 trials at a wand without improvement, they were returned to the next easier wand. The success of a bout was determined by the number of errors and the movement time on each wand/puzzle combination. The success level requirements are given in Table 2. For example, for a "successful" bout with Wand 1 and Puzzle 1,

Table 2. The ‘Progression Heuristic’ for participants. The maximum allowable number of errors and movement times for progress with a given wand/puzzle combination are provided in each column.

	Puzzle 1		Puzzle 2		Puzzle 3	
	Errors	Time [s]	Errors	Time [s]	Errors	Time [s]
Wand 1	2	10	3	15	5	30
Wand 2	2	15	3	20	7	40
Wand 3	3	15	4	25	9	60
Wand 4	N/A	N/A	N/A	N/A	N/A	N/A

participants needed to make no more than two errors, and to complete the bout in 10 s or less. Thus, task difficulty was made to correspond to the participant’s performance level. The purpose of using the wire puzzle game was to have a functional, upper extremity practice task that had variable difficulty levels. The levels allow for adjusting difficulty as participants improve performance, and thus, for increasing the task challenge level. During the Challenge condition, after providing KR, the robot selected the next wand according to the progression heuristic, and then verbalized this to the participant.

Literature also indicates the importance of self-determination, wherein the participant receiving therapy is involved in the decision making process during recovery. In this case, we developed a coaching style that would query the participant about the preferred wand to use when practicing the task. The second coaching style is thus the Choice style. During the Choice condition, after providing KR, the robot allowed the participant to select the wand to use for subsequent practice.

Finally, the control condition provided a measure of how participants performed with a system that did not have the features of the first two conditions; namely, a style that did not respond to user performance or preferences. The style used in this control condition is called the Continuous style. In the Continuous condition, after providing KR, the robot selected the next wand using a cyclic schedule. Participants always began with the easiest wand, Wand 1, and worked their way up to the hardest wand, Wand 4. At this point, they were instructed to use the next easiest wand, back down to Wand 1, and continued this plan throughout the interaction.

Bandit’s verbalizations can be categorized as general task instructions (at the beginning and end of each 15 min practice block), wand selection/query (at the beginning of each bout), congratulations (at the completion of each bout), and KR (at the completion of each bout). Across all 3 conditions, the general task instructions, congratulations, and KR are identical; only the wand selection/query transcripts vary.

Instructions: At the beginning of each session, Bandit provided the participant with the task instructions and the study goals. It then prompted the participant to start the game. If the participant did not start within a specified amount of time, Bandit reiterated the instructions and provided encouragement. When the participant performed the task incorrectly, Bandit reiterated the relevant task instructions and encouraged the participant to try again.

Wand selection/query: The wand selection or query phrases are the same for the Challenge and Continuous conditions (e.g., “Please use the [green] wand for the next bout.”). In the Choice condition, the robot prompted the participant to select the wand (e.g., “Please choose which wand you would like to use for the next bout.”).

Congratulations: At the completion of a bout, Bandit uttered a congratulatory phrase (e.g., “Good job!”).

KR: Bandit also provided the participant with KR consisting of the number of errors and movement time. In this study, we provided an intense KR schedule of 100% in order to keep participants engaged in the experiment.

During task performance, Bandit did not speak, allowing the participant to focus on the practice task.

The three styles constituted the three experiment conditions, and were presented in a counterbalanced order across study participants, with each participant being exposed to each style. We hypothesized that the Choice and Challenge conditions would result in significantly more positive responses (measured with post-session surveys), and that the Continuous condition, a proxy for a control condition, would be less positively received/less preferred by participants. In the Choice condition, when Bandit queried the participant, he/she became involved in the interaction by providing his/her input and being forced to make a decision in order to proceed with the session. We expected the participants to sense that they were being involved in their intervention (rather than being ‘told’ what to do), which we expected to lead to more favorable responses (when compared to the Continuous condition). Similarly, in the Challenge condition, we expected the participants to realize that their performance affected the wand schedule dictated by the robot. Again, the sense of involvement was expected to lead to more favorable responses. In the Continuous condition, after realizing that Bandit’s wand schedule was fixed and not personalized, we hypothesized that participants would have less favorable responses.

4.5. Measures

We obtained a number of measures from the participants before, during, and after task performance in order to evaluate our hypotheses, as follows.

Performance measures: The electrically instrumented wire puzzle was used to collect the number of errors and the movement time for each bout. The number of errors (a proxy for motor control) and movement time are well-known measures for evaluating task performance after stroke [34, 35]. As a task gets harder, the number of errors and movement time increase.

Participant self-report measures: Exit surveys at the completion of each session were used to obtain participant feedback regarding the quality of the interaction. The surveys were created by the experimenters; each survey was a 2-page document that participants filled out by hand. Participants filled out the surveys alone in the test room. The experimenters only intervened to provide clarification about the survey, but did not otherwise discuss survey responses with the participants. The details of the survey questions will be presented in Section 5.

5. Results

We first describe our performance measure results. Though the main aim of the study is concerned with comparing coaching styles, we maintain a focus on the practicality of using our system in a therapeutic environment. Thus, it is important that the performance data reflect what would be expected in a motor practice task: as difficulty increases, performance decreases.

5.1. Performance Measures

We hypothesized that, for participants of mild-to-moderate impairment levels, the simplest puzzle would be the easiest, resulting in a high level of performance (few errors and short movement times). The hardest puzzle should be the most challenging, resulting in a low level of performance. Collapsing the average movement times and errors across participants validates our expected trends regarding puzzle difficulty

(See Figures 5 and 6). When evaluating performance from Puzzle 1 to Puzzle 3, the total movement times and the number of errors increase across participants.

Movement time and errors increased from Puzzle 1 to Puzzle 2 to Puzzle 3. ANOVA results indicated that movement time for Puzzle 3 was significantly higher than for Puzzles 1 and 2 ($p = .013$ and $.007$). Errors for Puzzle 3 were significantly higher than Puzzles 1 and 2 ($p = .007$ and $.013$). Further, errors for Puzzle 2 were significantly higher than Puzzle 1 ($p = .002$).




Figure 5. Aggregate errors across the three puzzle conditions.




Figure 6. Aggregate movement times across the three puzzle conditions.

5.2. Survey Results

The survey we administered consisted of four parts.

Part 1, designed to determine if participants noticed differences in the conditions, asked the following questions:

- “Were you able to differentiate the 3 different styles?”
- “How well were you able to differentiate the 3 different styles (on a scale of 1–100)?”
- “Why?”

Five of the seven participants stated they could differentiate among the conditions in all 3 sessions. One was unable to differentiate in session

1, but was able to distinguish in sessions 2 and 3. One participant was unable to distinguish in session 1, and could only distinguish among two styles in session 2 and 3.

Part 2 was designed to quantitatively measure preference across three conditions. To determine participant preference across robot conditions, we evaluated responses to the statements:

- “I felt like I could really trust the robot.”
- “I’d like a chance to interact with this robot more often.”
- “I felt distant to the robot.”

Our goal was to probe the participants’ affinity for the robot – thus, higher rankings on the first two questions (and lower ranking on the third) would imply an affinity for the robot. Since the survey was administered at the conclusion of each session, each participant rated the Challenge, Choice, and Continuous conditions 3 times over the course of the 3-session interaction.

Participants used a 1–7 Likert scale to respond to the questions. To analyze the data, we first computed an “Affinity score” by reversing the response to the third question and computing an average of the three responses. Next, we collapsed results across all three sessions for all participants. Due to the ordinal nature of the data, we utilized non-parametric statistical tests to evaluate any differences in preference for the coaching styles.

Boxplots reflecting participant affinity scores are shown in Figure 7. We performed a Kruskal Wallis one way ANOVA test to determine any group differences. There were no statistically significant differences among the different conditions ($H(2) = 2.638$, $p = 0.267$) with a mean rank of 30.71 for Challenge, 37.07 for Choice, and 28.21 for Continuous.




Figure 7. Boxplots reflecting the affinity score for the robot coaching conditions, collapsed across participants and across the 3-session interaction.

Part 3 asked participants to rank the 3 robot conditions according to the following questions:

- “Which of the robot conditions did you enjoy the most?”
- “In which condition did you feel the most frustrated?”
- “In which condition did you feel like you performed better?”
- “Which of the conditions felt more interactive?”
- “In which condition did you feel like the robot was most aware of your actions?”

Participants used a checkbox to determine which condition(s) best matched their feelings about the interactions. Participants could rate more than one condition. Therefore, for each question we present the number of times participants chose a given condition to answer the question (Table 3).

Table 3. Changes in scores for various robot coaching styles.

	Challenge	Choice	Continuous
Which of the robot conditions did you enjoy the most?	6	8	10
In which condition did you feel the most frustrated?	10	6	6
In which condition did you feel like you performed better?	5	9	6
Which of the conditions felt more interactive?	7	6	9
In which condition did you feel like the robot was most aware of your actions?	8	10	14




Figure 8. Participant error data when collapsed across sessions and participants. The abscissa represents Puzzles 1 to 3. The ordinate represents Wands 1 to 4.

6. Discussion

6.1. Performance Measures

Determination of relative difficulty levels is a necessary component of tuning the task difficulty to a participant's performance level. These results provide indications of where performance dropped off significantly. Figures 8 and 9 shows collapsed data across all participants and all trials for the number of errors and movement time. The contour plot indicates the mean number of errors and movement times for each wand/puzzle combination. Errors with a given wand increase only slightly from Puzzle 1 to Puzzle 2, but on Puzzle 3, the errors increase sharply (Figure 8). This effect is also evident in the movement time data (Figure 9).

Considering the data across the different dimensions of wands for a given puzzle, an increase in movement time (a performance decrease) occurs between Wands 3 and 4. These effects are also less pronounced, but nonetheless evident, in the error data. By comparing these results to our progression schedule (Table 2), it is evident that our schedule may be too lenient for the easier wand/puzzle combinations, but too strict on harder wand/puzzle combinations. Looking only at errors, our schedule required participants to perform at a much higher level (few allowable errors) than they typically performed. Our schedule allowed for movement times much longer than those typical of the participants.




Figure 9. Participant movement time data when collapsed across sessions and participants. The abscissa represents Puzzles 1 to 3. The ordinate represents Wands 1 to 4.

These results validate the puzzle game as a motor practice task.

6.2. Survey Results

On Part 1 of the survey, most participants were able to differentiate among all three robot coaching styles, with only one participant failing to differentiate among all 3 (citing only two that he could distinguish). This is an obvious but important step in ensuring that participant ratings on subsequent parts of the survey are measures of their preference, and not just noise due to confusion about the robot styles. Given that participants were able to distinguish among the styles, we are able to ask which style they preferred. Part 2 of the survey provides initial insight into this question, without asking directly. In other words, we use dimensions of likability in our 3 questions. From our results, there were no statistically significant differences among the scores for the different conditions. This validates the null hypothesis in our study; namely, that there are no differences among coaching conditions responding to user performance or input, and a condition that does not. However, it is interesting to note that, in Figure 7, the median values for the Challenge and Choice conditions were both higher than that of the Continuous condition (5.5, 6.0, and 5.0, respectively). Therefore, the data trend towards a preference for the Challenge and Choice conditions.

It is also interesting to observe trends in the spread of the data. One participant had very low affinity scores for the Choice condition, drastically affecting the quartile values. Overall, there was more variability in the Choice condition responses. This may be due to the strategies employed by various participants. We observed that, when given the choice, some participants constantly used the least difficult wand (Wand 1); others explored the space somewhat randomly, using different wands in no apparent order; still others followed a Challenge-like progression, starting with Wand 1 and only moving to a more difficult wand after they felt sufficiently comfortable. We suspect this variability of strategies led to different perceptions of the coaching style. Because the participant was so clearly involved in the exercise, it may have been hard to separate out which part of the coaching was inherent to the robot, and which part was inherent to the participant's own practice strategy.

Part 3 of the survey provided insight into how participant perspectives of robot styles may differ from those of the experimenters. We will address these question by question. In response to the first question, "Which of the robot conditions did you enjoy the most?",

participants ranked (in descending order) Continuous, Choice, then Challenge. It is important to note that this questions differs from those of Part 2, which specifically ask about trust, future interaction, and distance with respect to the robot.

Further insight is gained from the second question, *"In which condition did you feel the most frustrated?"*. In this case, participants ranked Challenge as the most frustrating, with Choice and Continuous both ranked lower than Challenge. This comes as no surprise, as we expected the Challenge condition to tune difficulty to the participant's performance. Thus, this trend, may indicate that the relatively low enjoyment participants experienced in the Challenge condition (from Question 1) may have been tied to a sense of frustration with the difficulty level.

For the third question, *"In which condition did you feel like you performed better?"*, Choice was ranked higher than both Challenge and Continuous. We suspect that this is due to the participants having more control over the interaction. It may also be due to the fact that self-determination has been shown to have psychosocial and motivational effects on individuals in task practice settings, which may contribute to a sense of better performance.

For the fourth question, *"Which of the conditions felt more interactive?"*, participants ranked Continuous higher than Challenge and Choice. This was a surprising result, as we expected the Choice condition to be ranked the highest. Note that, on the fifth question *"In which condition did you feel like the robot was most aware of your actions?"*, participants also ranked the Continuous condition more highly than Challenge and Choice. One possible reason could be that, in the Choice condition, the robot essentially let the participant guide the practice. As a result, participants might have perceived this as a lack of interaction or awareness. The fact that Continuous was ranked higher than Challenge on these two questions may indicate that participants had a difficult time articulating what was occurring in the Challenge condition. The progression heuristic may have been difficult to understand for participants, preventing them from determining how interactive or aware the robot was. In future studies, followup questions probing the meaning of such terms (e.g., interactivity and awareness) may shed further light on how these results can be generalized.

Overall, the survey results indicate that, while participants were able to distinguish among the coaching styles, a number of factors may have interacted to result in no statistically significant differences in participant preference across conditions. In addition to the possible reasons mentioned above, it is possible that more pronounced differences among conditions, or perhaps explicit descriptions of the aims of the different conditions, might result in observable differences in preference for different robot coaching styles.

When looking at the results from the questions as a whole, we can make more general statements regarding the nature of the interaction conditions. Challenge was responsive to participant actions, but in a way that may have been too subtle or difficult for participants to articulate. Choice may have been seen as responsive because the robot asked participants a direct question; however, the lack of acknowledgement of the participant choice may have reduced the perceived level of interaction. The Continuous condition, which did not poll the participant, was still seen as interactive. We also note that a 'comparative' effect might have arisen. For instance, participants rating the Continuous condition as more interactive may be doing so in direct comparison to the other conditions. The perception that Continuous was less interactive may be coupled to the perception that Challenge and Choice were less interactive.

7. Conclusions

We have presented a socially assistive robotics system for augmenting participant performance in a motor rehabilitation task. We discussed our goal of determining how different robot coaching styles might influence participant perception of the interaction. We then described the interactive characteristics of our system: the use of embodiment and non-verbal gestures, motivational verbalization, KR, and engaging the participant through challenge. We also detailed how, in such a framework, a system can be personalized to each participant, through the use of valence-based nodding, Challenge-based coaching style, and Choice-based coaching style. We then described the design of a specific SAR system, its implementation on a stroke rehabilitation exercise task, and its evaluation with a number of post-stroke participants. Our findings that there were no significant differences in preference across conditions may suggest that the the human element of such interactions (e.g., a participant's personality, a participants' understanding of the survey questions) may play different roles in human-human interactions and human-robot interactions. Further study of how outcomes from other fields (e.g., psychology, sociology) may interact with human-robot interactions, especially in populations exhibiting functional deficits, is warranted. While our study employed a specific single coaching scenario, the scenario was realistic and study participants were drawn from a real-world target population; these results may be useful in informing the design of interactions involving SAR systems tasked with eliciting behavioral outcomes through the use of the various interaction modalities.

Acknowledgement

The authors thank Shuya Chen for her advice and help in administering motor task assessments and the experiments. This work was supported by Award Number U01NS056256 from the National Institute of Neurological Disorders and Stroke. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NINDS or the NIH. This work is also supported by NSF CNS-0709296 grant for "CRI: IAD – Computing Research Infrastructure for Human-Robot Interaction and Socially Assistive Robotics" and NSF IIS-0713697 grant "HRI: Personalized Assistive Human-Robot Interaction: Validation in Socially Assistive Robotics for Post-Stroke Rehabilitation."

References

- [1] D.J. Feil-Seifer and M.J. Mataric, Defining socially assistive robotics. In International Conference on Rehabilitation Robotics. Chicago, IL, 2005, 465–468.
- [2] T. Fong, I. Nourbakhsh, A survey of socially interactive robots. Robotics and Autonomous Systems, Special issue on Socially Interactive Robots 42, 3-4, 2003, 143–166.
- [3] J. Tao and T. Tan, Affective computing: A review. Lecture Notes in Computer Science 3784, 2005, 981–995.
- [4] A. Duschau-Wicke, T. Brunsch, L. Lunenburger, and R. Riener, Adaptive support for patient-cooperative gait rehabilitation with the lokomat. IROS, 2008, 2357–2361.
- [5] S. Banala, S. Kim, S. Agrawal, and J. Scholz, Robot assisted gait training with active leg exoskeleton (alex). IEEE Trans Neural Syst Rehabil Eng. 17, 1, 2009 (Feb.), 2–8.

- [6] D. Reinkensmeyer, L. Kahn, M. Averbuch, M. McKenna-Cole, and W. Rymer, Understanding and treating arm movement impairment after chronic brain injury: progress with the arm guide. *J. Rehabil. Res. Dev.* 37, 6, 2000 (Nov-Dec), 653–662.
- [7] A.C. Lo, P.D. Guarino, L.G. Richards, J.K. Haselkorn, G.F. Wittenberg, D.G. Federman, R.J. Ringer, T.H. Wagner, H.I. Krebs, B.T. Volpe, C.T. Bever, D.M. Bravata, P.W. Duncan, B.H. Corn, A.D. Maffucci, S.E. Nadeau, S.S. Conroy, J.M. Powell, G.D. Huang, and P. Peduzzi, Robot-assisted therapy for long-term upper-limb impairment after stroke. *N. Engl. J. Med.* 326, 19, 2010 (May), 1772–1783.
- [8] K. Dautenhahn, C.L. Nehaniv, M.L. Walters, B. Robins, H. Kose-Bagci, N.A. Mirza, and M. Blow, Kaspar - a minimally expressive humanoid robot for human-robot interaction research. *Applied Bionics and Biomechanics* 6, 3, 2009, 369–397.
- [9] T. Shibata and K. Wada, Robot therapy - a new approach for mental healthcare of the elderly. *Gerontology*, 2010 .
- [10] A. Tapus, M.J. Mataric, and B. Scassellati, The grand challenges in socially assistive robotics. *IEEE Robotics and Automation Magazine* 14, 1, 2007 (Mar), 35–42.
- [11] J. Fasola and M.J. Mataric, Robot motivator: Improving user performance on a physical/mental task. In *Proceedings of the International Conference on Human-Robot Interaction*, 2009.
- [12] A. Tapus, C. Tapus, and M.J. Mataric, The use of socially assistive robots in the design of intelligent cognitive therapies for people with dementia. In *International Conference on Rehabilitation Robotics*, 2009.
- [13] D.J. Feil-Seifer, M.P. Black, M.J. Mataric, and S. Narayanan, Toward designing interactive technologies for supporting research in autism spectrum disorders. In *International Meeting for Autism Research*, 2009, Chicago, Il.
- [14] A. Parnandi, E. Wade, and M.J. Mataric, Motor function assessment using wearable inertial sensor. In *32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC'10)*, 2010.
- [15] E. Wade, A. Parnandi, and M.J. Mataric, Automated administration of the wolf motor function test for post-stroke assessment. In *ICST 4th International ICST Conference on Pervasive Computing Technologies for Healthcare* 2010, 2010.
- [16] M.J. Mataric, A. Tapus, C.J. Winstein, and J. Eriksson, Socially assistive robotics for stroke and mild TBI rehabilitation. In *Advanced Technologies in Rehabilitation*. Vol. 145., 2009, IOS Press, 249–262.
- [17] World Health Organization, “Burden of disease statistics,” 2010, <http://www.who.org>.
- [18] R. Mead, E. Wade, P. Johnson, A.B.S. Clair, S. Chen, and M.J. Mataric, An architecture for rehabilitation task practice in socially assistive human-robot interaction. In *19th IEEE International Symposium in Robot and Human Interactive Communication*, 2010, Viareggio, Italy.
- [19] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Ng, ROS: an open-source Robot Operating System *ICRA Workshop on Open Source Software*, 2009, Kobe, Japan.
- [20] M.J. Johnson, X. Feng, L.M. Johnson, and J.M. Winters, Potential of a suite of robot/computer-assisted motivating systems for personalized, home-based, stroke rehabilitation. *J. NeuroEngineering and Rehabilitation* 4, 6, 2007 (March), 1–17.
- [21] C.D. Kidd, Designing for long-term human-robot interaction and application to weight loss. PhD Thesis, 2008, Massachusetts Institute of Technology.
- [22] J. Wainer, D.J. Feil-Seifer, D.A. Shell, and M.J. Mataric, Embodiment and human-robot interaction: A task-based perspective. In *IEEE Proceedings of the International Workshop on Robot and Human Interactive Communication*. 2007, Jeju Island, South Korea, 872–877.
- [23] R. Looije, M. Neerinx, and F. Cnossen, Persuasive robotic assistant for health self-management of older adults: Design and evaluation of social behaviors. *Int. J. Human-Computer Studies* 68, 2010, 386–397.
- [24] J. Fasola, and M.J. Mataric, Comparing Physical and Virtual Embodiment in a Socially Assistive Robot Exercise Coach for the Elderly. Submitted to *7th ACM/IEEE International Conference on Human-Robot Interaction*. Boston, MA, USA, Mar. 2012.
- [25] J. Liepert, H. Bauder, W.H.R. Miltner, E. Taub, C. Weiler, Treatment-Induced Cortical Reorganization After Stroke in Humans. *Stroke* 31, 2000, 1210–1216.
- [26] S. Kiesler, Fostering common ground in human-robot interaction. *IEEE Int. Workshop on Robot and Human Interactive Communication*, 2005, 729–734.
- [27] R. Schmidt and T. Lee, *Motor Control and Learning*. Human Kinetics, 2005.
- [28] J. Lee and S. Marsella, Nonverbal behavior generator for embodied conversational agents. In *6th International Conference on Intelligent Virtual Agents*. 2006, Marina Del Rey, California.
- [29] A. Fugl-Meyer, L. Jaasko, I. Leyman, S. Olsson, and S. Steglind, The post-stroke hemiplegic patient. 1. a method for evaluation of physical performance. *Scandinavian J. of Rehab. Medicine* 7, 1, 1975, 13–31.
- [30] E. Mower, D.J. Feil-Seifer, M.J. Mataric, and S. Narayanan, Investigating implicit cues for user state estimation in human-robot interaction using physiological measurements. In *IEEE Proceedings of the International Workshop on Robot and Human Interactive Communication*. 2007, Jeju Island, South Korea.
- [31] F. Rose, E. Attree, A. Brooks, D. Parslow, P. Penn, and N. Ambihapahan, Training in virtual environments: transfer to real world tasks and equivalence to real task training. *Ergonomics* 43, 4, 2000 (Apr), 494–511.
- [32] S. Braun, M. Kleynen, J. Schols, T. Schack, A. Beurskens, D. Wade, Using mental practice in stroke rehabilitation: a framework. *Clin. Rehabil.* 2008 Jul;22(7):579-91.
- [33] M. Guadagnoli and T. Lee, Challenge point: a framework for conceptualizing the effects of various practice conditions in motor learning. *J. Mot. Behav.* 36, 2, 2004, 212–224.
- [34] S.B. Thies, P.A. Tresadern, L.P. Kenney, J. Smith, D. Howard, J.Y. Goulermas, C. Smith, J. Rigby, Movement variability in stroke patients and controls performing two upper limb functional tasks: a new assessment methodology. *J. Neuroeng. Rehabil.* 6,2 ,2009.
- [35] A. Handley, P. Medcalf, K. Hellier, D. Dutta, Movement disorders after stroke, *Age Ageing* 38,3 , 2009, 260–266.