

Assessing the Quality and Quantity of Social Interaction in a Socially Assistive Robot–Guided Therapeutic Setting

Eric Wade^{*†}, Jonathan Dye[†], Ross Mead[‡] and Maja J. Matarić[§]

^{*}Division of Biokinesiology

[†]Department of Mechanical Engineering

[‡]Department of Computer Science

[§]Professor of Computer Science, Neuroscience and Pediatrics
University of Southern California

Abstract—We present an application of a socially assistive robotics (SAR) system in a therapeutic setting. We examine the amount of interaction elicited by the robot in a therapeutic setting with individuals post–stroke. We examine the role of various communication modalities, and their affects on the participants’ responses. Seven participants of mild to moderate functional impairment due to stroke interacted with our SAR system during three sessions of motor task practice. The robot guided the users as they performed a wire puzzle task, while providing them with feedback about their performance. We evaluated the amount of verbalization and eye contact made with the robot. Our results indicate that users make eye contact more often than they verbalize when interacting with the robot. Further, user interactions are most frequent at the beginning of a practice session, and occur less frequently as the session progresses. When a user observes that the robot is not responding to a certain type of communication, the user limits the use of that communication modality. These insights should be useful in the design of future robot–based therapeutic interventions.

I. INTRODUCTION

Socially assistive robotics (SAR) has emerged as a promising field for the provision and administration of motivation, encouragement, and rehabilitation for those suffering from cognitive, motor, and social deficits [1]. The intersection of assistive robotics and socially interactive robotics, SAR focuses on hands–off interactions with robots in various domains, including tutoring, emotional expression, daily life assistance, and physical therapy [1] [2] [3]. In the domain of physical therapy, SAR is 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 and cognitive deficits [4]. Therapeutic interventions for stroke require intense, focused, adaptive, and repetitive task practice, focusing on specific real–world tasks modeled on activities of daily living (ADLs). Such task–specific practice can lead to recovery from motor task deficits, but requires understanding the patient’s condition.

The adaptive difficulty of these interventions is critical to the motor learning process. If a task is too easy, then the patient realizes they are not learning, and may lose interest. If the task is too difficult, it may frustrate, causing the patient to quit. In

fact, finding the right challenge point is necessary to optimize learning [5], and depends on understanding the user state.

Meanwhile, the number of trained clinicians needed to ensure recovery for this population is decreasing. This is where SAR can serve a valuable role. Given the determination of a deficit and an associated rehabilitation regimen by a trained clinician, SAR can be used as a tool for the guidance and provision of an intervention. Furthermore, the computational nature of the SAR system can be used to adapt to the user’s changing state.

In the past, we have evaluated the technological feasibility and acceptance of SAR robots in the post–stroke population [6] [7]. In this paper, we evaluate the efficacy, quality and quantity of interaction between study participants and our SAR robot during guided motor task practice. Seven individuals post–stroke performed a motor task practice game while being instructed, monitored, and guided by a socially assistive humanoid robot; we measured their performance and actions during each practice block. We also investigated relationships between their responses and robot actions.

The rest of the paper is organized as follows; in Section II, we provide background on the field of rehabilitation robotics. In Section III, we describe the study procedure. In Section IV, we present the study outcomes. Finally, in Section V, we offer discussion and conclusions of the study outcomes.

II. BACKGROUND

The field of rehabilitation robotics has benefited from recent advances in sensing and actuation. The general approach in this field has been to create orthoses that physically interact with individuals with motor deficits. The development of lower extremity (LE) devices is particularly mature; there are already commercially available devices such as the LOKOMAT[®] [8]. This device, designed for gait retraining, applies forces to the leg sections in an attempt to retrain healthy gait patterns. In response to some of the limitations of the LOKOMAT[®], ALEX (Active Leg EXoskeleton) was developed by Agarwal et al. [9]. This orthosis has increased flexibility for back motion (and a more natural gait pattern), and utilizes a force field controller

to guide the user’s legs through prescribed trajectories. Upper extremity (UE) devices have also been extensively developed. Again, the 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. to help individuals target objects using their affected limb [10]. The system was shown to produce improvements comparable to those seen after traditional rehabilitation therapies.

The field of socially assistive robotics (SAR) is a newer development, separate from the field of rehabilitation robotics. SAR is defined as the provision of assistance through social (not physical) interactions with robots [1]. A SAR robot 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 extrinsic performance feedback. The absence of physical contact means there are a minimum of safety concerns. Equally important is the ability of users of such systems to practice tasks based on activities of daily living without being physically constrained by the robot. Further advantages include lower costs, greater portability, and increased accessibility (relative to contact-based robot systems). A SAR interaction can take place in various environments, including the clinic and the participant’s home.

In our laboratory, we have shown the efficacy of SAR in a variety of domains, including elder care [11] [12], Autism Spectrum Disorders [13], stroke assessment [14] [15], and stroke rehabilitation [6]. In the context of stroke, SAR can be applied to coach and encourage the use of the stroke-affected limb in meaningful, unconstrained, functional tasks encountered in daily life. SAR can also provide personalized therapeutic interactions that can complement those provided by a therapist. In this paper, we investigate the quality and quantity of human interaction with a SAR system in a therapeutic setting.



Fig. 1. Bandit, a socially assistive humanoid robot, shown mounted on a Pioneer mobile base.

III. METHODS

A. Participants

For this study, 7 individuals in the chronic phase (> 6 months) of stroke recovery were recruited. The cohort consisted of 4 males and 3 females. All participants were pre-screened using phone interviews. The participants exhibited mild-to-moderate functional impairments (upper extremity Fugl-Meyer scores between 31 and 60), and ranged from 24 to 70 years of age. The main inclusion criteria were (1) concordance (right-side affected, right-side dominant pre-morbidity), (2) the ability to lift and hold a full soda can from a desk for 5 seconds, and (3) the ability to obtain transportation to and from the experiment site. The participants came to the study site over the course of 3 separate visits..




Fig. 2. The wire puzzle setup showing the puzzle and the four wands of varying difficulty levels.

B. Practice task

The physical setup for our SAR system consisted of a humanoid robot, an instrumented puzzle, and wearable sensors. The robot used in our experiments is the humanoid Bandit, developed at the USC Interaction Laboratory in conjunction with BlueSky Robotics (Figure 1). We have used Bandit in a variety of SAR interactions in the past [11] [12] [7]. In this study, Bandit was used to guide participants during the performance of a wire puzzle game (Figure 2). To complete the game, participants guided a wand with a ring at its tip along a wire puzzle (See Figure 3). Such wire puzzles have been used as measures of motor performance, but not with this specific population [16] [17]. Task difficulty was increased by changing the wand used by the participant to perform the task. The wand with the largest ring diameter was easiest, while the wand with the smallest ring diameter was the most difficult (See Figure 2). The purpose of using the wire puzzle game was to have a functional, upper extremity practice task with variable difficulty levels that could be adjusted as participants improved performance. It is known that adapting the challenge level to a



Fig. 3. A participant practicing the wire puzzle task while being guided by a socially assistive humanoid robot.

patient’s skill level helps to increase the potential for learning. While participants performed the task, the robot monitored the number of errors (contacts between the wand ring and the puzzle) and total movement time for one bout (the time to move from the pad at the start of the puzzle to other pad at the end, and back to the first pad). Error and timing information were provided verbally to the participant as knowledge of results (KR) at the completion of each bout. Finally, participant wrist motion was monitored using an inertial measurement unit (IMU) developed in our laboratory [18]. The IMU captured 3–axes each of accelerometer, rate gyroscope, and magnetometer data. The entire setup is shown in Figure 3.

C. Schedule

The study consisted of three days of task practice, each separated by at least 24 hours. At the beginning of the first session, we obtained informed consent, administered the Fugl–Meyer motor assessment, and administered a questionnaire regarding the participant’s confidence in their upper limb use. The participant then performed three 20–minute task practice blocks, separated by 10–minute breaks. At the beginning and end of each 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 as a coach.

The second session was similar to the first, with the exclusion of the consenting, motor assessment, and confidence measure. Again, the participants performed three 20–minute practice blocks with 10–minute breaks. We obtained fatigue measures, and administered an exit questionnaire, identical to that used after session 1, at the conclusion of the session.

In the third session, as in the second, participants performed three 20–minute practice blocks with 10–minute breaks. We

| Feature | Description |
|--------------|---|
| robot speech | Robot verbalizes. |
| point | Robot points at a wand. |
| human speech | Participant verbalizes. |
| look | Participant makes eye contact with robot. |
| complete | Participant completes 1 bout. |
| cheat | Participant cheats. |
| wand | Participant picks up a wand. |

TABLE I
FEATURES ENCODED FROM SESSION DATA.

obtained fatigue measures, and again administered an exit questionnaire at the conclusion of the session. The Session 3 exit questionnaire contained the same questions as were used in Sessions 1 and 2, but included additional questions regarding the overall quality of the intervention.

D. Measures

We obtained a number of measures from the participants during task performance in order to study the interaction with the robot and the impact of the practice schedule on individual performance.

- Performance: The electrically instrumented wire puzzle was used to collect the number of errors and the movement time for each bout.
- Fatigue: A visual analog fatigue scale was used to obtain self–reported fatigue measures at the beginning and end of each 20–minute practice block.
- Psychometrics: Exit surveys at the completion of each session were used to obtain user feedback regarding the quality of the interaction.
- Interaction: Video from each session was coded according to various features of interest.
- Confidence: The only pre/post measure used was a survey developed to determine an individual’s confidence in using the stroke–affected limb in a variety of real–world scenarios.
- Motion quality: We used an IMU sensor on the wrist to obtain measures of motion quality.

We use a subset of these measures to discuss the interaction between the user, the robot, and the task in the next section.

IV. RESULTS

To describe the interaction between the user and robot, we coded and analyzed video data of the sessions. We utilized a video coding software package, ANVIL, to extract a number of features; the feature list is shown in Table I.

A. Amount of Interaction

Using these features, we were able to develop quantitative measures of the human–robot interaction during task performance. We define interaction as a robot action (e.g., speaking) and a subsequent user action (e.g., verbalization or eye contact) that follows within a finite amount of time. As a first step towards investigating the amount of interaction between the user and the robot, we posed the following

question:

Question 1: *When the robot speaks, how many responses does it elicit from the user?* To answer this, we determined the total number of instances of verbalization or eye contact by the user after the onset of robot speech. For verbalization, we used all instances, regardless of content. For eye contact, we used all instances of the user looking at the robot. For each participant, we averaged the number of instances of verbalization and eye contact across their sessions. The results are shown in Figure 4. For each participant (P1 through P7), it appears that, excluding one participant (P5), there are fewer instances of verbalization than eye contact. To obtain further insight, we averaged the instances of verbalization and eye contact across all participants. The results, also shown in Figure 4 (individual and group data), indicate that, as a mode of interaction, eye contact occurred significantly more frequently than verbalization ($p < 0.001$). When collapsed across participants, instances of verbalization had mean $\mu = 6.13$ and deviation $\sigma = 11.67$. Instances of eye contact had $\mu = 14.54$ with $\sigma = 10.49$. Thus, the bulk of the social interaction consisted of looking at the robot, when compared to verbalizations. This inspired us to look more closely at the behavior of the participants, and motivated a second question.




Fig. 4. Total instances of verbalizations and eye contact after robot speech (individual and group data).




Fig. 5. Percentage of user verbalization or eye contact resulting from robot speech (individual and group data).

Question 2: *When the user is verbalizing or looking at the*

robot, what percentage of those instances result from robot speech? In other words, how much time did participants spend verbalizing or making eye contact with the robot because it was ‘speaking to’ them relative to the amount of time spent verbalizing or looking without an apparent stimulus (perhaps out of curiosity, confusion, etc.). To answer this, we looked for all instances of verbalization or eye contact, and compared them to instances of robot speech. The results are shown in Figure 5.

We found that instances of verbalization due to robot speech had $\mu = 23.9$ with $\sigma = 34.33$. Instances of eye contact that were due to robot speech had $\mu = 66.66$ with $\sigma = 14.38$. In this case, the percentage of eye contact elicited by the robot was greater than the percentage of verbalization elicited, though not significantly ($p = 0.08$).

B. Temporal nature of interaction

Based on these results, we can make the general statement that participants look at the robot ‘more’ than they speak to the robot. Is this due to the robot’s behavior? Or is it due to the format of the intervention? To answer this question, we investigated the temporal nature of the user verbalizations and eye contact responses. Specifically, we binned the number of responses in each minute to evaluate when, in the course of a practice block, the user was responding (So, the total number of responses in the first 60s of an interaction went into the bin for minute 1). For instance, equal numbers of instances of verbalization across all bins would indicate that a participant responds with fairly regular frequency. A large number of responses in the first few bins relative to fewer responses in later bins would indicate that a participant responded more often in the early parts of the session, than later on.

For a given participant, we averaged the number of responses in the first minute of each 20-minute practice block (thus, if a participant came for 3 days, he had 9 practice blocks, and 9 data points per minute). Then, averaging across participants, we determined the instances of verbalization and eye contact during a given minute of the interaction. The data, averaged across all participants, are shown in Figure 6. As the initial results indicated, eye contact occurred more than verbalization.




Fig. 6. Mean instances of verbalizations and eye contact, by minute.

Individual performances provide additional insight. Figure 7 shows that many occurrences of eye contact take place within the first 3 minutes, at which point there is a clear decline, and then a flattening out of the eye contact frequency. This is less evident in the results for Figure 8; however, excluding P5, we see a similar trend, albeit at lower amplitude.

V. DISCUSSIONS & ANALYSIS

A. User responses

The first question we posed gives insight into the interaction from the robot’s perspective. As stated, each practice session lasted up to 20 minutes. Since the number of elicited verbalizations had $\mu = 6.13$ with $\sigma = 11.67$, there were 0.31 instances of verbalization per minute of interaction. Since the number of instances of eye contact had $\mu = 14.54$ with $\sigma = 10.49$, there were 0.73 instances of eye contact per minute of interaction. This makes sense when we consider the nature of the interaction. While the robot spoke to the participants, it did not perform speech recognition, and as a result, did not understand their responses. The lack of speech recognition capability quickly became obvious to the participants. Thus, though participants did continue to have normal social responses to the robot’s speech, including making eye contact with the robot, they had far fewer instances of verbalization over time.




Fig. 7. Temporal occurrences of user visual responses.




Fig. 8. Temporal occurrences of user verbal responses.

We also analyzed the interaction from the user’s perspective. For verbalization, we found $\mu = 23.9$ with $\sigma = 34.33$. This indicates that less than a quarter of the verbalizations resulted from robot actions. Verbalizations may be largely due to personality differences across participants. For instance, two participants never spoke, while another verbalized throughout the interaction, whether or not the robot was speaking. With eye contact, the percentage of instances due to robot actions had $\mu = 66.66$ with $\sigma = 14.38$. Thus, the percentage of eye contact elicited by robot actions was greater than the percentage of elicited verbalization.

These results are further supported by the temporal data, in terms of the relative amount of verbalizations versus instances of eye contact. We note that the first 2 minutes of each practice block consisted of the robot providing task instructions to the users. Thus, the user naturally looked at the robot during this period (Figure 7). Note also in Figure 8 that the results for P5 inflate the mean value of verbalizations. Using the leverage criteria, we find that P5 is indeed an outlier (leverage point = 0.968). Removing the effect of the outlier shows an even greater difference between the amount of verbalization and instances of eye contact (Figure 9).




Fig. 9. Temporal occurrences of user verbal responses, outlier removed.

By looking at the participant–robot interaction from both the robot and user perspectives, we can reason that the difference between the amount of eye contact and verbalization is due to the fact that the robot provides information to the participant, but does not solicit or understand their responses. While this may not be problematic given the user is practicing and improving at the task, a study designed to obtain information about the human–robot interaction should focus on tasks that require more interaction. The wire puzzle task requires the user to focus most attention and effort on the puzzle, and not necessarily on the person (or in this case, the robot) providing instruction and feedback.

B. Limitations

Some issues worth further exploration are related to the features we used to describe the human–robot interaction. A number of taxonomies already exist for describing interactions between humans, agents, and robots. Having observed the interactions between our participants and the robot, we are

aware that some of the verbalizations were due to confusion, and may not have occurred in close temporal proximity to the activity that elicited them. For instance, the robot might provide instructions, and a user, after having tried to complete the puzzle for more than 60 seconds, finally expressed confusion verbally.

Finally, the social interaction we measure for participants could be tied to other actions, not considered in this analysis. For instance, individuals might verbalize surprise at their performance after the robot provides KR. To further investigate these two types of interaction, we plan to analyze the content of user feedback to examine the nature of responses elicited by the robot.

VI. CONCLUSION

We have presented a SAR framework for use in guided motor task practice for individuals post-stroke. The results of our study indicate that the type of interaction elicited by the robot depends both on the types of robot actions and on the practice task. We found that a task such as the wire puzzle, that requires a high level of focused attention, may be promising for intense task practice, but may not elicit rich interactions with the robot. Nonetheless, we were able to use the robot to guide and motivate individuals as they performed the task, an important requirement for long term motor task practice. Going forward, we will evaluate the ability of our system to make use of user feedback and performance to provide an intervention based on these initial results.

VII. ACKNOWLEDGEMENT

The authors would like to thank Shuya Chen for her advice and help in the administration of functional assessments for the participants. 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. Matarić, “Defining socially assistive robotics,” in *International Conference on Rehabilitation Robotics*, Chicago, IL, Jun 2005, pp. 465–468.
- [2] T. Shibata and K. Wada, “Robot therapy - a new approach for mental healthcare of the elderly.” *Gerontology*, Jul 2010.
- [3] K. Dautenhahn, C. Nehaniv, M. Walters, B. Robins, H. Kose-Bagci, N. Assif-Mirza, and M. Blow, “Kaspar - a minimally expressive humanoid robot for human-robot interaction research,” *Applied Bionics and Biomechanics*, vol. 6, no. 3, pp. 369–397, 2009.
- [4] W. H. Organization, “Burden of disease statistics,” 2010, <http://www.who.org>.
- [5] M.A. Guadagnoli and T.D. Lee, “Challenge point: a framework for conceptualizing the effects of various practice conditions in motor learning,” *J. Mot. Behav.*, vol. 36, no. 2, pp. 212–224, 2004.
- [6] M. J. Matarić, A. Tapus, C. J. Winstein, and J. Eriksson, “Socially assistive robotics for stroke and mild tbi rehabilitation,” in *Advanced Technologies in Rehabilitation*, P. L. T. W. G. R. Andrea Gaggioli, Emily A. Keshner, Ed. IOS Press, 2009, vol. 145, pp. 249–262.
- [7] R. Mead, E. Wade, P. Johnson, A. St.Clair, S. Chen, and M. J. Matarić, “An architecture for rehabilitation task practice in socially assistive human-robot interaction,” in *19th IEEE International Symposium in Robot and Human Interactive Communication*, Viareggio, Italy, Sep 2010.
- [8] A. A. Duschau-Wicke, T. Brunsch, L. Lunenburger, and R. Riener, “Adaptive support for patient-cooperative gait rehabilitation with the lokomat,” *IROS*, pp. 2357–2361, 2008.
- [9] S. Banala, S. Kim, S. Agrawal, and J. Scholz, “Robot assisted gait training with active leg exoskeleton (alex).” *IEEE Trans Neural Syst Rehabil Eng.*, vol. 17, no. 1, pp. 2–8, Feb. 2009.
- [10] D. Reinkensmeyer, L. Kahn, M. Averbuch, A. McKenna-Cole, B. S. B, and W. Rymer, “Understanding and treating arm movement impairment after chronic brain injury: progress with the arm guide.” *J Rehabil Res Dev.*, vol. 37, no. 6, pp. 653–662, Nov–Dec 2000.
- [11] J. Fasola and M. J. Matarić, “Robot motivator: Improving user performance on a physical/mental task,” in *Proceedings of the International Conference on Human-Robot Interaction*, Mar 2009.
- [12] A. Tapus, C. Tapus, and M. J. Matarić, “The use of socially assistive robots in the design of intelligent cognitive therapies for people with dementia,” in *International Conference on Rehabilitation Robotics*, Jul 2009.
- [13] D. J. Feil-Seifer, M. P. Black, M. J. Matarić, and S. Narayanan, “Toward designing interactive technologies for supporting research in autism spectrum disorders,” in *International Meeting for Autism Research*, Chicago, IL, May 2009.
- [14] E. Wade, A. Parnandi, and M. J. Matarić, “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.
- [15] A. Parnandi, E. Wade, and M. J. Matarić, “Motor function assessment using wearable inertial sensor,” in *32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC’10)*, 2010.
- [16] E. Mower, D. J. Feil-Seifer, M. J. Matarić, 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*, Jeju Island, South Korea, Aug 2007.
- [17] F. Rose, E. Attree, B. Brooks, D. Parslow, P. Penn, and N. Ambihapahan, “Training in virtual environments: transfer to real world tasks and equivalence to real task training.” *Ergonomics*, vol. 43, no. 4, pp. 494–511, Apr 2000.
- [18] E. Wade and M. J. Matarić, “Design and testing of lightweight inexpensive motion-capture devices with application to clinical gait analysis,” in *Proceedings of the International Conference on Pervasive Computing*, Aug 2009, pp. 1 – 7.