2007 IEEE International Conference on
Robotics and Automation
Roma, Italy, 10-14 April 2007

Hands-Off Therapist Robot Behavior Adaptation to User Personality for Post-Stroke Rehabilitation Therapy

Adriana Ţăpuş, *Member, IEEE,* Cristian Ţăpuş, *Member, IEEE,* and Maja J Matarić, *Senior Member, IEEE*

***Abstract*— This paper describes a hands-off therapist robot that monitors, assists, encourages, and socially interacts with post-stroke users in the process of rehabilitation exercises. We developed a behavior adaptation system that takes advantage of the users introversion-extroversion personality trait and the number of exercises performed in order to adjust its social interaction parameters (e.g., interaction distances/proxemics, speed, and vocal content) toward a customized post-stroke rehabilitation therapy. The experimental results demonstrate the robot's autonomous behavior adaptation to the user's personality and the resulting user improvements of the exercise task performance.**

***Index Terms*— Rehabilitation Robotics, Socially Assistive Robotics, Social Human-Robot Interaction, Learning and Adaptive Systems**

I. Introduction

THE recent trend toward developing a new generation of robots that are capable of moving and acting in human-centered environments, interacting with people, and participating and helping us in our daily lives has introduced the need for building robotic systems able to learn how to use their bodies to communicate and react to their users in a social and engaging way. Social robots that interact with humans have thus become an important focus of robotics research.

Nevertheless, Human-Robot Interaction (HRI) for *socially assistive* applications is still in its infancy. Socially assistive robotics, which focuses on the *social interaction*, rather than the *physical interaction* between the robot and the human user has the potential to enhance the quality of life for large populations of users, such as the elderly [30], people with physical impairments and in rehabilitation therapy (e.g., post-stroke patients) [7], [28], people with cognitive disabilities and social and developmental disorders (e.g., children with autism, children with attention deficit/hyperactivity disorder (AD/HD)) [23], [24], [25].

In our work, the target user population is post-stroke patients. Stroke is the leading cause of serious, long-term disability among American adults, with over 750,000 people suffering a new stroke each year [19]. Stroke patients are unable to perform movements with the affected limb, even though the limb is not completely paralyzed. This loss of function, termed learned disuse, can improve with rehabilitation therapy during the critical post-stroke period. The best strategy of any post-stroke rehabilitation program is the repetitive practice of exercises, which can be passive and active. In the passive exercises (also knows as hands-on rehabilitation), the patient is helped by the human (or robot) therapist to move the affected limb, while in the active exercises, the patient performs the exercises with no physical hands-on assistance. The vast majority of existing work into rehabilitation robotics focuses on hands-on robotic systems (e.g., [4], [5], [22]). However, recent results from physical therapy research show that such therapy may not be the most effective means of recovery from stroke, and is certainly not the only necessary type of much-needed treatment [7].

Our work focuses on hands-off therapist robots that assist, encourage, and socially interact with patients during their active exercises. We previously demonstrated [7], [11], [12], [27], [28], through real-world experiments with stroke patients, that the physical embodiment (including shared physical context and physical movement of the robot), encouragement, and monitoring play key roles in patient compliance with rehabilitation exercises. Recently, we also investigated the role of the robot's personality in the hands-off therapy process, by focusing on the relationship between the level of extroversion/introversion (as defined in Eysenck Model of personality [10]) of the robot and the user [28].

Building robotic systems capable of adapting their behavior to user personality, user preferences, and user profile so as to provide an engaging and motivating customized protocol is a very difficult task, especially when working with vulnerable users. Different learning systems for human-robot interaction have been proposed in the literature [3], [20], but none of them includes the human profile, preferences, and/or personality in the model. To the best of our knowledge, no work has yet tackled the issue of robot behavior adaptation as a function of user personality in the assistive human-robot interaction context. In our work, we address this issue and propose a behavior adaptation system based on reinforcement learning. The robot incrementally adapts its behavior as a function of the users introversion-extroversion level and of the amount of exercises he/she has performed, aiming toward a more individualized and appropriately challenging/nurturing therapy style that will help to improve user task performance. Our robot behavior adaptation system monitors

Manuscript received September 15, 2006; This work was supported by USC Women in Science and Engineering (WiSE) Program and the Okawa Foundation.

Dr. Adriana Ţăpuş is with the Robotics Research Lab/Interaction Lab, Department of Computer Science, University of Southern California, Los Angeles, USA (phone: +1 (213) 740 6245: fax: +1 (213) 821 5696; e-mail: tapus@robotics.usc.edu)

Dr. Cristian Ţăpuş is with the Mojave Research Lab, Department of Computer Science, California Institute of Technology (Caltech), Pasadena, USA (e-mail: crt@caltech.edu)

Prof. Maja J. Matarić is with the Robotics Research Lab/Interaction Lab, Department of Computer Science, University of Southern California, Los Angeles, USA (e-mail: mataric@usc.edu)

1-4244-0602-1/07/$20.00 ©2007 IEEE.

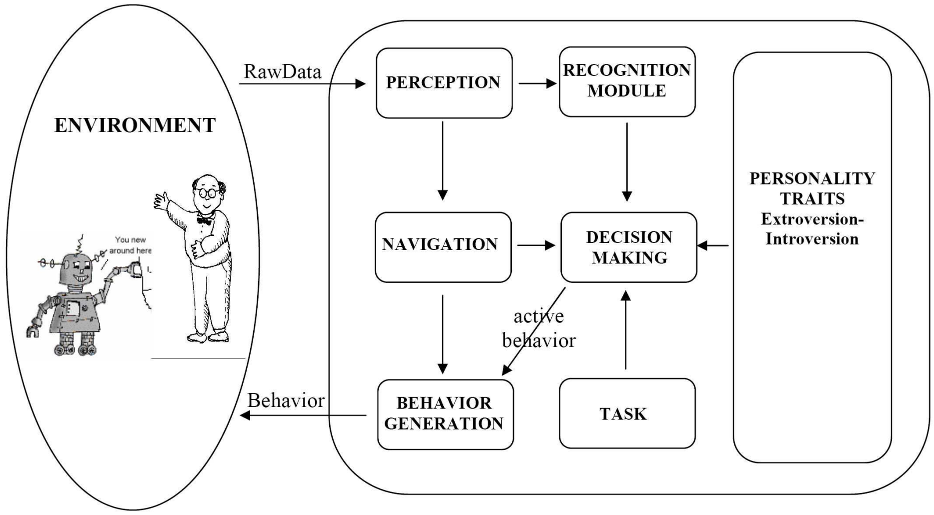

Fig. 1. HRI information processing using the personality model of the user

the number of exercises performed by the human/patient, which indicates the level of progress, and changes the robots behavior in order to maximize this level.

The rest of the paper is structured as follows. Section II presents a brief overview of our interaction design. In Section III the proposed behavior adaptation system is described. Section IV is dedicated to the description of the experimental test-bed and the experimental setup. Experimental results are presented in Section V. Finally, in Section VI, we conclude the paper and discuss future work.

II. INTERACTION DESIGN

To date, none of the existing robotic systems for socially assistive applications integrate the personality dimension in their behavioral model. Inspired by Bandura's model of reciprocal influences on behavior [2], we believe that it is helpful to incorporate the personality dimension in order to improve human-robot interaction (HRI) and behavior selection. Figure 1 depicts our general behavior control architecture, which integrates the Eysenck model [10].

The extraversion-introversion dimension is based on the observed inter-correlations between traits such as sociability, activity, impulsiveness, liveliness and excitability, all of which strongly influence behavior (see Figure 2).

In our interaction design, we chose to use two of those traits: sociability and activity, which can be most readily emulated in robot behavior. We chose to express these two personality traits through three main parameters that define the therapist robot behavior: interaction design / proxemics, speed, and verbal and para-verbal communication. These are described in more detail below.

A. Sociability

Sociability is the trait that most clearly expresses a person's level of extroversion-introversion. A large body of research in social psychology has shown that individual

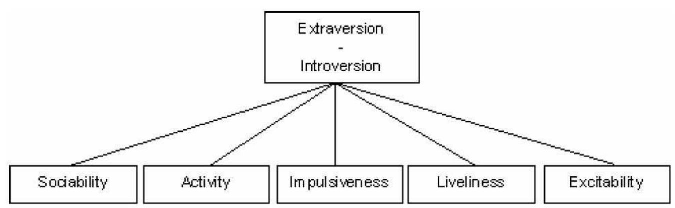

Fig. 2. Hierarchical level of Eysenck's extraversion-introversion personality

behavioral differences are most apparent in social situations [6], [8], [16]. In [14], Harkins, Becker and Stonner empirically illustrated that both the presence of others and their social activities are typically more enjoyed by extraverts than by introverts. In [9], Eysenck described the extravert as sociable, friendly, talkative and outgoing. In contrast, the introvert is quiet, introspective, and prefers small groups of intimate friends. We posit that these are directly related to verbal and non-verbal communication patterns. Hence, we identified proxemics and vocal features (i.e., content, volume, and speech rate) as relevant aspects to be embodied in the robots behavior. Each is described below.

1) Proxemics: The interpersonal space in human interactions has been widely studied in social psychology. Hall [13], pioneer of the field of proxemics, identified four general interaction spaces: **Intimate** (up to 0.3m; involves physical contact), **Personal** (between 0.3-1.3m; typically used for family and friend interaction); **Social** (about 1.3-3m; used in business meetings and in public spaces); **Public** (beyond 4m, e.g., the distance between an audience and a speaker) (see Figure 3).

In this work, we focus only on personal and social interaction spaces. Neither the intimate space, nor the public space is appropriate for our application; the former implies contact and the latter involves no interaction. In [13], Hall analyzed and found a strong link between human sense of

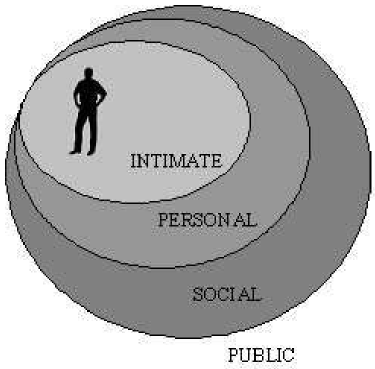

Fig. 3. Interaction zones / proxemics: intimate, personal, social, and public

space and human behavior and personality type. We posit that extraverted individuals, who like social interactions, may prefer to have the robot physically closer than introverted individuals, who may perceive the robot as invading their space. Therefore, proxemics can be encoded as function of the individual extraversion-introversion level.

2) Verbal and Para-Verbal Communication: Both vocal content and paralinguistic cues, such as volume and speech rate, play important roles in human interactions, and express personality and emotion [1], [21], [29]. The similarity-attraction principle, which assumes that individuals are more attracted to other people manifesting the same personality as theirs, has been studied in HCI (e.g., [18]). The interaction scripts that we designed displayed extroverted and introverted personality types through the choice of words and paralinguistic cues. More details about the different interaction scripts are given in the Experimental Design section (Section IV).

B. Activity

In addition to sociability, we also considered the activity trait. Eysenck, in [9], [10], linked the human introversion-extroversion personality trait with the activity level and showed that people with high activity scores are generally energetic and favor physical activity, while individuals with low scores tend to be physically inactive. Therefore, high activity is an extravert characteristic, while low activity tends to characterize introversion. In our system, the activity of the robot is correlated/matched to the users movement and sociability, and its represented through robots movement speed. Some of the movements that the robot performs are: turning around itself, moving around the participant, and moving toward and away so as to vary the proxemics. These movements are performed at different speeds as a function of the introversion/extroversion of the participant.

III. The Behavior Adaptation System

The main goal of our robot behavior adaptation system consists of optimizing the three main parameters (interaction distance/proxemics, speed and vocal content) that define the therapist robot behavior, so as to adapt to the users personality and improve his/her task performance. The system monitors the users task performance and the time spent between exercises, and changes the robot's behavior in order to maximize the patients level of progress.

We formulated the problem as policy gradient reinforcement learning (PGRL) and developed a learning algorithm. The functioning of PGRL algorithm is the following: (a) parameterization of the behavior; (b) approximation of the gradient of the reward function in the parameter space; and (c) moving towards a local optimum. More details about the algorithm can be found in [15], [26]. Other reinforcement learning techniques, such as Q-Learning, learn an action-value function. Nevertheless, Q-learning, designed for Markov decision processes, cannot directly be applied to our problem since there is no obvious notion of state.

Figure 4 shows the pseudo-code of the *policy gradient reinforcement learning* (PGRL) algorithm. Our n-dimensional policy gradient algorithm implemented for this work starts from an initial set of parameters $\pi = \{\pi_1, \pi_2, ..., \pi_n\}$ (where n = 3 in our case) and generates p random policies in the vicinity of π. The perturbations $\pi^{Pert} = \{\pi_1{}^p, \pi_2{}^p, \ldots, \pi_n{}^p\}$ are calculated by randomly adding either $+\epsilon_i$, 0, or $-\epsilon_i$ to the initial policy. Each ϵ_i is individually chosen for each parameter and is small relative to π_i. The new perturbed parameter set is tested with a user and the reward function is evaluated. Next, the partial derivative in each of the n-dimensions is estimated. This is realized by grouping each $\pi_i{}^p$ into one of three sets for each dimension n, as shown in Figure 4 (steps (14), (15), and (16)). The average rewards (i.e., $Avg_{+\epsilon,j}$, $Avg_{+0,j}$, and $Avg_{-\epsilon,j}$) for each parameter for the three cases (i.e., $+\epsilon_i$, 0, or $-\epsilon_i$) are calculated. These three averages give an estimate of the benefit of altering the i parameter by $+\epsilon_i$, 0, or $-\epsilon_i$.

The gradient in dimension i, A_i, is considered to be 0 if the average reward for the unperturbed parameter is greater than the other two average rewards, and is considered to be the difference between the average rewards for the perturbed parameters otherwise. The gradient A is normalized and multiplied by the system step size η, so that the adjustment will remain a fixed size for each iteration. Finally, the parameter set π is adjusted by adding A.

The reward function that we used to evaluate the behavior of the robot was one of the major challenges in implementing the adaptive algorithm. The main issues with computing the reward function and running an adaptive algorithm in our case were the following: (1) the events that mark the interaction between the robot and the participant are discrete, thus computing the reward function can only occur at discrete moments in time; (2) the evaluation of the reward function has to take into consideration both: (a) the fact that as the participant performs the exercises it will incur fatigue, which will slow him/her down, regardless of the personality of the robot or other parameters that we considered, and (b) the fact that the robot is adapting and acting differently can distract the patient, slowing down his/her response.

With these factors in mind, we designed the reward function as follows: (1) we counted the number of exercises performed by the patient during a given period of time, and

$\pi = \{\pi_1, \pi_2, \ldots, \pi_n\} \leftarrow$ Initial policy composed of the initial set of n parameters
$\pi^{Pert} = \{\{\pi_1{}^p, \pi_2{}^p, \ldots, \pi_n{}^p\} \quad | \quad p \in \{1..k\}\} \leftarrow$ Perturbed parameter sets derived from π
$\epsilon = \{\epsilon_1, \epsilon_2, \ldots, \epsilon_n\} \leftarrow$ Parameter step size vector
$\eta \leftarrow$ System step size
while (not done) **do**
  for p = 1 **to** k
    for i = 1 **to** n
      $r \in \{-1, 0, 1\} \leftarrow$ Randomly chosen
      $\pi_i{}^p = \pi_i + \epsilon_i * r$
    end for
  end for
  Evaluate the reward by running the system using the parameter set π^{Pert}
  for i = 1 **to** n
    $Avg_{+\epsilon,i} \leftarrow$ average reward for all π^p with positive perturbation (for which r was 1) in dimension i
    $Avg_{0,i} \leftarrow$ average reward for all π^p with no perturbation (for which r was 0) in dimension i
    $Avg_{-\epsilon,i} \leftarrow$ average reward for all π^p with negative perturbation (for which r was -1) in dimension i
    if $(Avg_{0,i} > Avg_{+\epsilon,i}\ and\ Avg_{0,i} > Avg_{-\epsilon,i}$
      $A_i \leftarrow 0$
    else
      $A_i \leftarrow (Avg_{+\epsilon,i} - Avg_{-\epsilon,i})$
    end if
  end for
  $A \leftarrow \frac{A}{|A|} * \eta$
  for i = 1 **to** n
    $A_i = A_i * \epsilon_i$
  end for
  $\pi = \pi + A$
end while

Fig. 4. Pseudo-code for the Policy Gradient Reinforcement Learning (PGRL) Algorithm

(2) we adjusted the threshold value to reduce the effects of fatigue and distraction caused by the adaptation procedure.

Similar applications of the adaptive algorithm deal with either more consistent environments, or can evaluate the reward function on a continuous basis, making adaptation seem more real-time. For applications like the one described in [15], the reward function is clearly determined as the speed of the robot. Furthermore, in that case the same robot was used in the same environment for the entire set of tests, reducing the number of uncontrolled perturbations that can affect the computation of the reward function.

Fig. 5. Robot Test-bed

IV. Experimental Setup

A. Robot Test-bed

Our experimental test-bed, shown in Figure 5, consists of an ActivMedia Pioneer 2-DX mobile robot base, equipped with a SICK LMS200 laser rangefinder used to track and identify people in the environment by detecting reflective fiducials worn by the users.

An OHAUS SCOUT Pro Scale, a simple electronic scale measuring weight, connected to a RS232 interface, was used to compute the number of exercises performed by the user.

B. Experimental Design

Two experiments were designed to test the adaptability of the robot's behavior to the participant's personality and preferences. Our experimental design attempted to adapt and match robot's behavior and interaction style to different user's personality traits.

In each experiment, the human participant stood or sat and faced the robot. The experimental task was intended as a functional exercise similar to those used during standard stroke rehabilitation and consisted of moving pencils from one bin on the left side of the participant to another bin on his/her right side. The bin on the right was on a scale, in order to measure task performance.

The participants were asked to perform the task for 15 minutes, but they could stop the experiment at any time. At the end of each experiment, the experimenter presented a short debriefing. The adaptation algorithm was running for the entire duration of the experiment, but it was automatically activated only when the participant was performing below the set threshold. The adaptation window was set to five minutes.

So far, our system has been validated with only healthy participants. In order to be able to obtain more relevant results, the healthy volunteers used their non-dominant limb (their weaker side) while doing the specified tasks. They were also encouraged to establish a social relationship with the robot based on its personality and act as they would normally do when interacting with a person with the same characteristics.

Before starting the experiments, the participants were asked to fill two questionnaires: (1) a general introductory questionnaire in which personal details such as gender, age, occupation, and educational background were determined

TABLE I
INITIAL PARAMETERS FOR THE BEHAVIOR ADAPTATION ALGORITHM

Robot Behavior Parameters		Initial Values	Step Size ϵ
Therapy Style and Robot's Personality as expressed through vocal content and para-verbal cues		Id=1	1
Interaction distance / proxemics	Extroverted	0.7m	0.5m
	Introverted	1.2m	0.5m
Speed	Extroverted	0.1m/s	0.1m/s
	Introverted	0.1m/s	0.05m/s

and (2) a personality questionnaire based on the Eysenck Personality Inventory (EPI) [9] for establishing the user's personality traits.

Our learning algorithm was initialized with parameter values that were in the vicinity of what was thought to be acceptable for both extroverted and introverted individuals, based on our previous study [28]. These values are described in Table I.

On the post-experiment survey, the participants were asked to provide their preferences related to the therapy styles, interaction distances, robot's speed, and robot's vocal cues (i.e., gender and accent). The participants were also asked to rate their impressions on the robots personality on a 7-point Likert scale (i.e., from strongly agree to strongly disagree).

1) Experiment 1: The goal of the first experiment was to test the adaptability of the robot behavior to the user personality-based therapy style preference. Four different scenarios were designed for both extroverted and introverted personality types. These therapy styles ranged from coach-like therapy to encouragement-based therapy for extroverted personality types and from supportive therapy to nurturing therapy for introverted personality types (see Table II). The words and phrases for each of these scenarios were chosen carefully and in concordance with encouragement language used by human therapists. The coach-like therapy script was composed of very strong and aggressive language (e.g., "Move! Move!", "You can do more than that!"). Higher volume and faster speech rate were used in the pre-recorded transcript voice, based on the evidence that those cues are associated with high extroversion [17]. The aggressiveness of words, the volume, and speech rate diminished along with the robot's movement towards the nurturing therapy style of the interaction spectrum. Therefore, the nurturing therapy script contained only empathetic, gentle, and comforting language (e.g., "I'm glad you are working so well.", "I'm here for you.", "Please continue just like that", "I hope its not too hard"). The voice used had lower volume and pitch.

A set of 3 interaction distances and speeds were chosen for each introverted and extroverted personality types. These are detailed in Table III.

2) Experiment 2: To construct an engaging and motivating customized protocol, in the second experiment we wanted to ensure the robot was able to adapt to user preferences. People are more influenced by certain voices and accents than others. Two main scenarios were designed, one for extroverted individuals and one for introverted, respectively.

TABLE II
THE CHOICE OF THERAPY STYLES AS A FUNCTION OF THE USER-PERSONALITY

Parameter	Extroverted			
Therapy Style	Id=1	Id=2	Id=3	Id=4
	Coach-like	Very Challenging	Stimulating	Encouragement-based
	Introverted			
	Id=1	Id=2	Id=3	Id=4
	Supportive	Educative	Comforting	Nurturing

TABLE III
THE CHOICE OF INTERACTION DISTANCES/PROXEMICS AND ROBOT'S SPEED PARAMETERS AS A FUNCTION OF THE USER-PERSONALITY

Parameter	Extroverted			Introverted		
Interaction Distance/ Proxemics (m)	Id=1	Id=2	Id=3	Id=1	Id=2	Id=3
	0.7	1.2	1.7	1.2	1.7	2.2
Speed (m/s)	0.1	0.2	0.3	0.1	0.15	0.2

The scenario for the extroverted group was challenge-based while the scenario for the introverted individuals was more nurturing, in conformity with our previous study [17]. We pre-recorded the same scenario with 2 males (one with accent - French native speaker, and one without accent American native speaker) and 2 females (one with accent - Romanian native speaker, and one without accent American native speaker) (see Table IV).

The choice of interaction distances/proxemics and robot movement speeds was the same as in the first experiment (see Table III).

V. EXPERIMENTAL RESULTS

The pilot experimental group consisted of 11 participants (6 male and 5 female). The participants ranged in age between 19 and 35, 27% were from a non-technological field, and 73% worked in robotics or other technological field.

A. Experiment 1

The results obtained in the first experiment are mainly shown in Figure 6. Figure 6 shows, for each participant, the percentage of time spent interacting with each of the four therapy styles of the robot. Bars represent the percentages with respect to the total time spent doing the exercises. Crosses represent the preference of the participant, as specified in a post-experiment survey. As illustrated in Figure 6, the robot adapted to match the preference of the participant in almost every case. The only exception was the interaction

TABLE IV
THE CHOICE OF THERAPIST ROBOT'S PERSONALITY AS EXPRESSED THROUGH ENGLISH ACCENT AND VOICE GENDER AS A FUNCTION OF THE USER-PREFERENCES

Parameter	Id=1	Id=2	Id=3	Id=4
Therapist Robot's Personality as Expressed through English Accent and Voice Gender	Female with accent	Male with accent	Male without accent	Female without accent

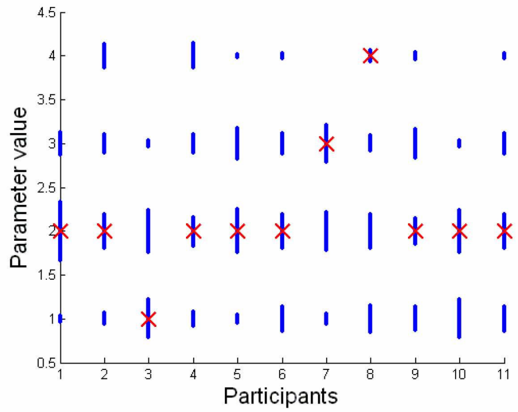

Fig. 6. The percentage of time that the 11 participants interacted with each of the four therapy styles of the robot (for extroverted and introverted participants, as described in Table II). Crosses represent the participants' interaction distances preferences.

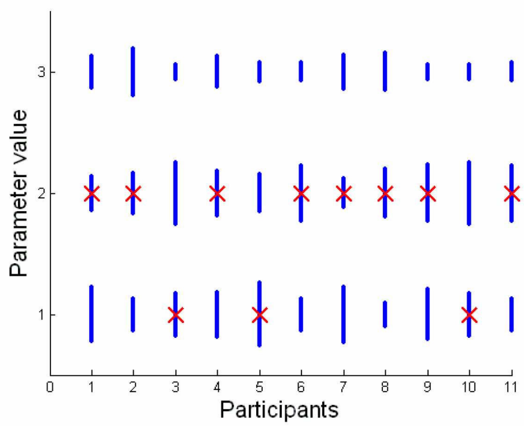

Fig. 7. The percentage of time that the 11 participants interacted with the robot at a certain distance (for extroverted and introverted participants, as described in Table III). Crosses represent the participants' interaction distance preferences.

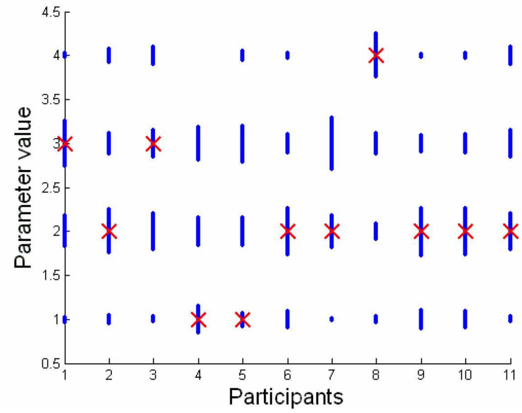

Fig. 8. The percentage of time that the 11 participants interacted with each of the four therapist robot's personality as expressed through accent and voice gender (gender and English accent, as described in Table IV). Crosses represent the participant's preferences.

with participant number 8. Despite the fact that time spent in the preferred training style of the participant was shorter than time spent in other training styles, the robot converged to it at the end of the exercise period. The reason for this slight inconsistency was caused by the fact that the initial state of the robot was in a training style that was furthest from the preference of the participant. The fact that we only allowed perturbations to neighboring training styles, combined with the relatively short duration of the exercise, jointly contributed to this result.

The results show that the robot could adapt its behavior to both introverted and extroverted participant therapy style preferences.

The adaptability of the robot matched the preferences of the participant even in the case of the preferred distance for interaction, as shown in the Figure 7. Both the extrovert and introvert personalities chose a distance that matched the personal state rather than the social space for their interaction with the robot.

The robot also succeeded in adapting its speed to user preferences. Introverted users preferred the lower speeds of the robot and extroverted participants preferred the higher speeds of the robot's movement.

B. Experiment 2

In the second experiment, the results were again consistent with our hypothesis that the robot is able to adapt and match the participant's preferences. Nevertheless, we encountered two special cases, which we discuss next. First, participant number 5 was not particularly influenced by the variations in voice and accents and, as part of the post-experiment survey, mentioned the fact that his second preference was in fact the one in which the robot spent most of the time out of the four choices. Second, participant number 7 had a preference for a male therapist robot but did not care whether it spoke with an accent or not. This is in fact consistent with choices 2 and 3 in which the robot spent 90% of the total time of the exercise (see Figure 8).

For this case the distance between the robot and the participant (see Figure 9) did not match perfectly the preference of the participant. We believe that this was caused by the fact that the number of adaptation steps was rather small compared to the size of the state space of the parameters and that this parameter had a smaller impact on the participant reaching his/her goal.

In this second experiment, we obtained similar results for the speeds as the ones in the first experiment.

VI. Conclusions and Future Works

This paper demonstrated the on-line adaptation of the robot's behavior (i.e., interaction therapy styles, interaction distance/proxemics, and speed) to user personality and preferences. Our experiments show that, with our proposed model, the robot can adapt to user personality and preferences and can help users to improve their task performance, by changing its own behavior. Future work includes exercises over longer periods of time which will allow the robot to

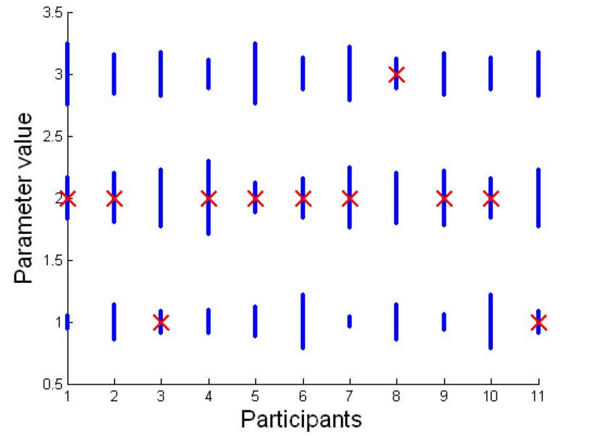

Fig. 9. The percentage of time that the 11 participants interacted with the robot at a certain distance (for extroverted and introverted participants, as described in Table III). Crosses represent the participants' interaction distance preferences.

better adapt as well as investigating various reward functions that would enable us to understand the impact of each parameter over the success rate of the adaptation process.

Even if socially assistive robotic technology is still in its early stages of development, the next decade promises assistive robotic platforms and systems that will be used in hospitals, schools, and homes in therapeutic programs that monitor, encourage, and assist their users. It is therefore important that potential users, well beyond the technical community, become familiar with this growing technology and help shape its development toward its intended positive impact on numerous lives.

References

[1] W. Apple, L. A. Streeter, and R. M. Krauss. Effects of pitch and speech rate on personal attributions. *Journal of Personality and Social Psychology*, 37:715–727, 1979.

[2] A. Bandura. *Principles of behavior modification*. Holt, Rinehart & Wilson, New York, USA, 1969.

[3] M. Berlin, J. Gray, A. L. Thomaz, and C. Breazeal. Perspective taking: An organizing principle for learning in human-robot interaction. In *Proc. of the 21st National Conference on Artificial Intelligence (AAAI)*, 2006.

[4] B. R. Brewer, R. Klatzky, and Y. Matsuoka. Feedback distortion to overcome learned nonuse: A system overview. *IEEE Engineering in Medicine and Biology*, 3:1613–1616, 2003.

[5] C. G. Burgar, P. S. Lum, P. C. Shor, and M. Vander Loos. Development of robots for rehabilitation therapy: the palo alto va/stanford experience. *Journal of Rehabilitation research and Development*, 37(6):639–652, 2000.

[6] N. S. DiCaprio. *Personality theories: A guide to human nature*. Holt, Rinehart & Wilson, New York, USA, 1983.

[7] J. Eriksson, M. J. Matarić, and C. Winstein. Hands-off assistive robotics for post-stroke arm rehabilitation. In *Proc. IEEE International Conference on Rehabilitation Robotics (ICORR'05)*, pages 21–24, Chicago, Il, USA, June 2005.

[8] R. B. Ewen. *An Introduction to Theories of Personality*. Lawrence Erlbaum Associates, 2003.

[9] H. J. Eysenck. *The structure of human personality*. Methuen, London, UK, 1953.

[10] H. J. Eysenck. Dimensions of personality: 16, 5 or 3? criteria for a taxonomic paradigm. *Personality and Individual Differences*, 12:773–790, 1991.

[11] D. Feil-Seifer and M. J. Matarić. Defining socially assistive robotics. In *Proc. IEEE International Conference on Rehabilitation Robotics (ICORR'05)*, pages 465–468, Chicago, Il, USA, June 2005.

[12] R. Gockley and M. J. Matarić. Encouraging physical therapy compliance with a hands-off mobile robot. In *Proc. of the First International Conference on Human-Robot Interaction (HRI'06)*, pages 150–155, Salt Lake City, USA, March 2006.

[13] E. T. Hall. *Hidden Dimension*. Doubleday, Gordon City, NY, USA, 1966.

[14] S. Harkins, L. A. Becker, and D. Stonner. Extraversion-introversion and the effects of favorability and set size on impression formation. *Bulletin of the Psychonomic Society*, 5:300–302, 1975.

[15] N. Kohl and P. Stone. Policy gradient reinforcement learning for fast quadrupedal locomotion. In *Proc. IEEE International Conference on Robotics and Automation (ICRA'04)*, volume 3, pages 2619–2624, New Orleans, USA, May 2004.

[16] L. W. Morris. *Extraversion and Introversion: An Interactional Perspective*. Hemisphere Pub Corp., 1979.

[17] H. Nakajima, Y. Morishima, R. Yamada, S. Brave, H. Maldonado, C. Nass, and S. Kawaji. Social intelligence in a human-machine collaboration system: Social responses to agents with mind model and personality. *Journal of the Japanese Society for Artificial Intelligence*, 19(3):184–196, 2004.

[18] C. Nass and M. K. Lee. Does computer-synthesized speech manifest personality? experimental tests of recognition, similarity-attraction, and consistency-attraction. *Journal of Experimental Psychology: Applied*, 7(3):171–181, 2001.

[19] National Institute of Neurological Disorders and Stroke. *Post-Stroke Rehabilitation Fact Sheet*, January 2006.

[20] Monica Nicolescu and Maja J. Matarić. Learning and interacting in human-robot domains. *IEEE Transactions on Systems, Man, Cybernetics, special issue on Socially Intelligent Agents - The Human in the Loop*, 31(5):419–430, 2001.

[21] J. Pittman. *Voice in social interaction: An interdisciplinary approach*. Sage, Thousand Oaks, CA, USA, 1994.

[22] D. J. Reinkensmeyer, C. T. Pang, J. A. Nesseler, and C. C. Painter. Web-based tele-rehabilitation of the upper extremity after stroke. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 10(2):102–108, 2002.

[23] B. Robins, K. Dautenhahn, R. te Boekhorst, and A. Billard. Robotic assistants in therapy and education of children with autism: Can a small humanoid robot help encourage social interaction skills? *Universal Access in the Information Society (UAIS)*, 4(2):105–120, December 2005.

[24] B. Scassellatti. How social robots will help us to diagnose, treat and understand autism. In *Proc. of the 12th International Symposium of Robotics Research (ISSR'05)*, San Francisco, CA, USA, October 2005.

[25] B. Scassellatti. Quantitative metrics of social response for autism diagnosis. In *Proc. of the 14th IEEE International Workshop on Robot and Human Interactive Communication (RO-MAN'05)*, Nashville, USA, August 2005.

[26] R. Sutton, D. McAllester, S. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in Neural Information Processing Systems*, 12:1057–1063, 2000.

[27] A. Tapus and M. J. Matarić. Towards socially assistive robotics. *International Journal of the Robotics Society of Japan (JRSJ)*, 24(5):14–16, 2006.

[28] A. Tapus and M. J. Matarić. User personality matching with hands-off robot for post-stroke rehabilitation therapy. In *Proc. International Symposium on Experimental Robotics (ISER'06)*, Rio de Janeiro, Brazil, July 2006.

[29] K. J. Tusing and J. P. Dillard. The sounds of dominance: Vocal precursors of perceived dominance during interpersonal influence. *Human Communication Research*, 26(1):148–171, 2000.

[30] M. Walton. Meet paro, the therapeutic robot seal. In *CNN*, 2003.