Long Term Learning and Online Robot Behavior Adaptation for Individuals with Physical and Cognitive Impairments

Adriana Tapus, Cristian Tapus and Maja Matarić

Abstract In this paper, we present an online adaptation approach and a long-term learning approach for socially assistive robotic (SAR) systems that aim to provide customized help protocols through motivation, encouragements, and companionship to users suffering from physical and/or cognitive changes related to stroke, aging and Alzheimer's disease.

1 Introduction

A recent trend in robotics is to develop a new generation of robots that are capable of moving and acting in human-centered environments, interacting with people, and participating in our daily lives. This has introduced the need for developing robotic systems able to learn how to use their bodies to communicate and to 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.

Research into Human-Robot Interaction (HRI) for socially assistive applications is in its infancy. Socially assistive robotics [4] is an interdisciplinary and increasingly popular research area that brings together insights from a broad spectrum of fields, including robotics, health, social and cognitive sciences, and neuroscience, among others.

It is estimated that in 2050 there will be three times more people over the age 85 than there are today [1]. Most of the ageing population is expected to need physical and/or cognitive assistance. As the elderly population continues to grow, new research has been dedicated to developing assistive systems aimed at promoting ageing-in-place, facilitating living independently in one's own home as long

Prof. Adriana Tapus is with ENSTA-ParisTech, Paris, France, e-mail: adriana.tapus@ensta.fr · Dr. Cristian Tapus is a Research Scientist at Google Inc., Mountain View, CA, USA. e-mail: crt@google.com · Prof. Maja Matarić is with the University of Southern California, Computer Science Department, Los Angeles, USA. e-mail: mataric@usc.edu

as possible, and helping caregivers and doctors to provide long-term rehabilitation/cognitive stimulation protocols. The first efforts towards having socially assistive robotic systems for the elderly have been focused towards constructing robot-pet companions aimed at reducing stress and depression [5], [12], [7], [6], and [8]. In addition to the growing elderly population, other large user populations represent ideal beneficiaries of socially interactive assistive robotics. Those include individuals with physical impairments and those in rehabilitation therapy, where socially assistive technology can serve to improve not only mobility [13], [2] [3] but also for outcomes in recovery. Finally, individuals with cognitive disabilities and developmental and social disorders (e.g., autism [9]) constitute another growing population that could benefit from assistive robotics in the context of special education, therapy, and training.

In order to be able to aid the target user populations, an effective socially interactive assistive robot must understand and interact with its environment, exhibit social behavior, and focus its attention and communication on the user in order to help the user achieve specific goals. Social behavior plays an important role in the assistance of people with special needs. An adaptive, reliable and user-friendly hands-off therapist robot can provide an engaging and motivating customized therapy protocol to participants in laboratory, clinic, and ultimately, home environments, and can establish a very complex and complete human-robot relationship. Therefore, such robots must be endowed with human-oriented interaction skills and capabilities to learn from us or to teach us, as well as to communicate with us and understand us. Hence, the work proposed here will focus on robot behavior adaptation to user's personality, preferences and disability level, aiming toward a long-term customized therapy protocol for stroke rehabilitation and other elderly specific application domains.

This paper presents two learning approaches, one based on on-line adaptation and the other based on long-term learning, for socially assistive robots designed for helping stroke patients and people suffering of age-related cognitive impairments (i.e., dementia). The rest of the paper is structured as follows. Section 2 illustrates the robotic test-bed. Section 3 describes the online learning and behavior adaptation approach and its validation in a rehabilitation-like context and Section 4 describes the long-term learning approach and its validation in a study with patients with dementia. Section 5 concludes the paper.

2 Experimental Platform

The experimental testbed used was a custom-designed humanoid torso robot mounted on a mobile robot base (Figure 1). The mobile base was an ActivMedia Pioneer 2DX robot equipped with a speaker, a Sony Pan-Tilt-Zoom (PTZ) color camera, and a SICK LMS200 eye-safe laser range finder. The anthropomorphic setup involved a humanoid Bandit II torso, consisting of 22 controllable degrees of freedom, which included: 6 DOF arms (x2), 1 DOF gripping hands (x2), 2 DOF pan/tilt neck, 2 DOF pan/tilt waist, 1 DOF expressive eyebrows, and a 3 DOF expressive mouth.

All actuators were servos allowing for gradual control of the physical and facial expressions. We are interested in utilizing the humanoid's anthropomorphic but not highly realistic appearance as a means of establishing user engagement, and comparing its impact to our prior work with non-biomimetic robot test-beds [11].

Fig. 1: Robot test-bed: Bandit II humanoid torso mounted on the Pioneer mobile base

3 Study 1

3.1 Robot Learning and Behavior Adaptation to User Personality and Preferences

The main goal of the first implemented methodology was to develop a robot behavior adaptation system that allows for dynamically optimizing three main interactional parameters (in our case: interaction distance/proxemics, speed, and vocal content) so as to adapt to the user's personality toward improving the user's task performance. These parameters defined the behavior (and thus personality) of the "therapist" robot. Task performance is measured as the number of exercises performed in a given period of time; the learning system changed the robot's personality, expressed through the robot's behavior, in an attempt to maximize the task performance metric.

A learning algorithm based on policy gradient reinforcement learning (PGRL) was developed. The n-dimensional policy gradient algorithm implemented for this work starts from an initial policy $\pi = \{\theta_1, \theta_2, \ldots, \theta_n\}$ (where n = 3 in our case). For each parameter θ_i we also defined a perturbation step ε_i to be used in the adaptation process. The perturbation step defined the amount by which the parameter may vary to provide a gradual migration towards the local optimum policy. The use of PGRL required the creation of a reward function to evaluate the behavior of the robot as parameters changed to guide it toward the optimum policy. The algorithm consisted of the following steps: (a) parametrization of the behavior initial policy π; (b) approximation of the gradient of the reward function in the parameter space; (c) movement towards a local optimum.

The reward function was monitored to prevent it from falling under a given threshold, which would indicate that the robot's behavior at the time did not provide the user with an ideal exercise scenario. This triggered the activation of the PGRL adaptive algorithm phase to adapt the behavior of the robot to the continually-changing factors that determined the user's task performance. More details about this work can be found in [11].

3.2 Experimental design for learning in the physical exercise context

We endeavored to develop an experimental design for a study involving stroke patients, and validate it first with non-patients, in a lab setting, in order to test the adaptation algorithm. In the experimental design, the participant stands or sits facing the robot. The experimental task is a common object transfer task used in post-stroke rehabilitation and consists 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 is on an electronic scale in order to measure the participant's task performance. The system monitors the number of exercises performed. The participants are asked to perform the task for a fixed amount of time (15 minutes for healthy adults, 6 minutes for stroke patients), but they can stop the experiments at any time. At the end of each experiment session, the experimenter presented a short debriefing. Before starting the experiments, the participants are asked to complete two questionnaires: (1) a general introductory questionnaire in which personal details such as gender, age, occupation, and educational background were determined and (2) a personality questionnaire based on the Eysenck Personality Inventory (EPI) for establishing the user's personality traits.

The learning algorithm is initialized with parameter values that are in the vicinity of what is thought to be acceptable for both extroverted and introverted individuals, based on the user-robot personality matching study described in [10]. The PGRL algorithm evaluates the performance of each policy over a period of 60 seconds. The reward function, which counts the number of exercises performed by the user in the past 15 seconds is computed every second and the results over the 60 seconds "steady" period are averaged to provide the final evaluation for each policy. The threshold for the reward function that triggers the adaptation phase of the algorithm is adjusted to account for the fatigue incurred by the participant. The threshold and the time ranges are all customizable parameters.

In the post-experiment survey, the participants are asked to provide their preferences related to the therapy styles or robot's vocal cues, interaction distances, and robot's speed from the values used in the experiments.

We designed four different scenarios for extroverted and introverted personality types; the 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. We chose to use pre-recorded speech and selected words and phrases for each of these scenarios in concordance with encouragement language used by professional rehabilitation therapists. The challenge-based therapy script is composed of assertive language (e.g., “Keep going!” and “You can do more than that!”). Extoversion is also expressed with higher speech volume and faster speech rate. The aggressiveness of words, volume, and speech rate are adjusted to diminish along with the robot’s movement towards the nurturing therapy style of the interaction spectrum. In contrast to the challenge-based script, the nurturing therapy script contains 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 it’s not too hard”). The speech uses lower volume and pitch. The transition from one personality-based therapy style to another is done smoothly (see algorithm above) in order to avoid any jarring influence on the human-robot interaction. We chose a set of three interaction distances and speeds for each introverted and extroverted personality type.

3.3 Experimental results for the physical exercise context

We performed the above-described experiment in a lab-like setting, with 12 non-patient, healthy adult participants (7 male, 5 female). The participants ranged in age between 19 and 35; 27% were from a non-technological field, while 73% worked in a technology-related area.

As shown in Figure 2, the robot adapted to match the preference of the participant in almost every case. The only exception was the interaction with participant 8. Despite the fact that the time spent in the preferred training style of that participant was smaller than the time spent in other training styles, the robot converged to it at the end of the exercise period. This was caused by the fact that the initial state of the robot was in a training style that was furthest from the participant’s preference.

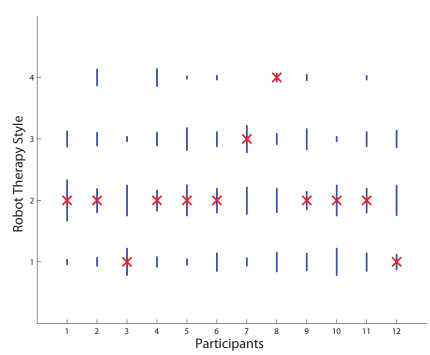

Fig. 2: The percentage of time that the 12 participants interacted with each of the four therapy styles of the robot. The crosses represent the participants’ preferences.

The pilot results we obtained support our hypothesis that the robot could adapt its behavior to both introverted and extroverted participants. Further details about this work can be found in [11].

4 Study 2

4.1 Robot Learning and Behavior Adaptation to User Ability/Performance

The second learning and robot behaviour adaptation methodology was designed for the interaction between the robot and a user with dementia and/or Alzheimer's disease, with the main goal of helping users improve or maintain their cognitive attention through encouragements in a music-based cognitive stimulation game.

This approach consists of two parts: supervised learning and adaptation. The robot models the level of game challenge that can be: (a) Difficult: no hints; (b) Medium: when the song excerpt starts say "push the button" but do not indicate which button to push; and (c) Easy: when the song excerpt starts say which button to push. The supervised learning system learns the Accepted Variation Band (AVB) for each game level and for each disability bucket (mild, moderate, and severe), as a function of the user's task performance. The learning phase is followed by an adaptation phase, where the robot adapts its behavior so as to minimize the user's reaction time and maximize the correctness of the user's answers. If the user's task performance is below the Accepted Variation Band, the user is performing better than during the learning phase. The user is then promoted to the next level of game difficulty (if not already at the Difficult level). If the user's task performance is above the Accepted Variation Band, the user is not performing well enough. The user is then helped by having the game difficulty level decreased (if not already at the Easy level).

4.2 Experimental design for cognition exercise context

The experiment consists of repeated sessions, during which the user and the robot interact in the context of a cognitive game. The first session is the orientation, in which the participant is 'introduced' to the robot. The robot is brought into the room with the participant, but is not powered on. During this introduction period, the experimenter or the participant's nurse/physical/music therapist explains the robot's behavior, the overall goals and plans of the study, and what to expect in future sessions. The participant is also asked about his/her favorite songs from a variety popular tunes from the appropriate time period; those songs are later used in the subsequent sessions. At the end of the session, the Standardized Mini-Mental State

Examination (SMMSE) cognitive test is administered so as to determine the participant's level of cognitive impairment and the stage of dementia. This test provides information about the cognitive (e.g., memory recall) level of impairment of the participant for use in initializing the game challenge level. The data determine the participant's initial mental state and level of cognitive impairment, and serve as a pre-test for subsequent end-of-study comparison with a post-test.

This experiment is designed to improve the participant's level of attention and consists of a cognitive game called Song Discovery or Name That Tune. The participant is asked to find the right button for the song, press it, say the name of the song, and sing along. The criteria for participation in the experiment (in addition to the Alzheimer's or dementia diagnosis) include the ability to read large print and to press a button. The participant sits in front of a vertical experimental board with 5 large buttons (e.g., the Staples EASY buttons). Four buttons correspond to the different song excerpts (chosen as a function of the user's preference) and the last button corresponds to the SILENCE or no song excerpt condition. Under each button, a label with the name of the song (or SILENCE) is printed. The robot describes to each participant the goal of the game before each session, based on the following transcript: "We will play a new music game. In it, we will play a music collection of 4 songs. The songs are separated by silence. You will have to listen to the music and push the button corresponding to the name of the song being played. Press the button marked "SILENCE" during the silence period between the songs. The robot will encourage you to find the correct song." Each participant is first asked by the music therapist or the robot to read aloud the titles of the songs and to press a button. Some additional directions are given. The participant is also directed to press the SILENCE button when there is no music playing. After a review of the directions, the participant is asked by the robot to begin the music game. The music compilation is composed of a random set of song excerpts out of the four different songs that form the selection and the silence condition. The entire music compilation lasts between 10 and 20 minutes, and is based on the user's level of cognitive impairment: the larger the impairment, the shorter the session. A song excerpt can be vocal, instrumental, or both. The order of song excerpts is random. The experiment was repeated once per week for a period of 8 months in order to capture longer-term effects of the robot therapist. A within-subject comparison was performed to track any improvement over multiple sessions. No between-subject analysis was done due to the small sample size and large differences in cognitive ability levels.

4.3 Experimental results for the cognitive exercise context

The initial pilot experimental group consisted of 9 participants (4 male, 5 female), from our partner Silverado Senior Living care facility. All the participants were seniors over 70 years old suffering of cognitive impairment and/or Alzheimer's disease. The cognitive scores assessed by the SMMSE test were as follows: 1 mild, 1 moderate, and 7 severe. Due to the total unresponsiveness of 6 of the severely affected participants, only 1 severely cognitively disabled participant was retained for the rest of the study, resulting in a final participant group composed of 3 participants (all female).

Fig. 3: Human-robot interacting during the music game: the robot gives hints related to the music game, the user answers, and the robot congratulates and applauds the correct answer

We constructed the training data and built a model for each cognitive disability level and for each game level. The participants played each game level 10 times (stages) in order to construct a robust training corpus.

The results obtained over 6 months of robot interaction (excluding the 2 months of learning) suggest that the elderly people suffering of dementia and/or Alzheimer's can sustain attention to music across a long period of time (i.e., on average 20 minutes for mildly impaired participants, 14 minutes for moderately impaired participants, and 10 minutes for severely impaired participants) of listening activity designed for the dementia and/or Alzheimer's population. Figures 4a, 4b, and 4c illustrate the evolution of the game difficulty over time, as well as response incorrectness and reaction time for user_id 1.

Outcomes are quantified by evaluating task performance and time on task. Based on the results we obtained, it can be concluded that the SAR system was able to adapt the challenge level of the game it was presenting to the user in order to encourage task improvement and attention training. Figure 4a shows the evolution in time of the game level for user_id 1. The participant started at the easy game level and remained there for several sessions. The participant then started to perform better and diminished the reaction time and reduced the number of incorrect answers, which, in turn, resulted in a game level evolution from the easy level to difficult. Starting from the 22nd trial, the participant consistently remained at the highest level of difficulty in the game (see Figure 4a). Figures 4b and 4c depict the evolution of the reaction time and the number of incorrect answers. The decrease of those metrics indicates improvement on the task. Similar improvement was observed for all participants.

The participants recognized the songs and identified the silence periods with the same probability. Hence, the analysis of the "no answer" situation among our elderly participants provides us with additional information. From our experiments, we noticed that the average rate of absence of response to silence was higher than the average rate of absence of response to songs, and that this phenomenon in-

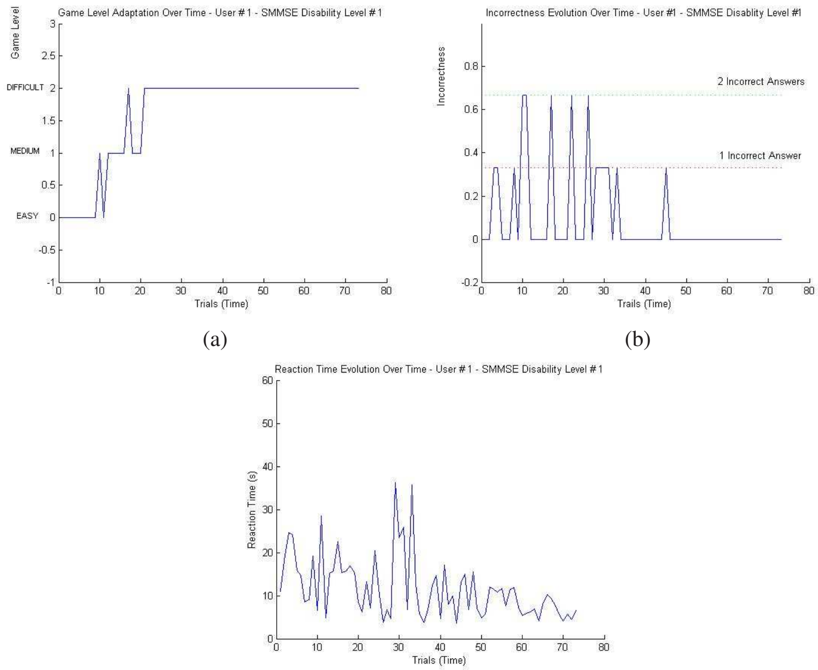

(a) (b)

(c)

Fig. 4: Results: (a) Game Level Adaptation and Evolution Over Time (6 months) for User Id 1; (b) Incorrectness Evolution Over Time (6 months) for User Id 1; (c) Reaction Time Evolution Over Time (6 months) for User Id 1

creased with the severity of the cognitive impairment. Our conjecture is that music stimulates the interest and responsiveness of the participants. Another interesting observation that deserves more study is the users' ability to participate simultaneously in different tasks (multitasking): the participants were able to sing and push the correct buttons at the same time. This is notable in particular for participants with cognitive disability, since multitasking requires dividing attention.

In summary, our social robot was able to improve or maintain the cognitive attention of users with Alzheimer's Disease in a specific music-based cognitive game. The robot's capability of adapting its behavior to the individual user's level of disability helped to improve the user's task performance in the cognitive game over time.

5 Conclusions

This research has aimed to develop adaptation and learning methods for socially assistive therapist robots that can provide customized physical rehabilitation and/or cognitive stimulation. We have presented results from two different adaptation and learning approaches, validated with healthy adults as well as with elderly users with Alzheimer's Disease. Our results are encouraging in light of our pursuit toward creating personalized socially assistive technologies that aim to improve human quality of life.

Acknowledgements This work was supported in part by the National Academies Keck Futures Initiative (NAKFI), by the USC Alzheimer's Disease Research Center (ADRC), by the NSF IS-0713697 grant, and by the USC WiSE Program. The infrastructure for this research was supported by the NSF Computing Research Infrastructure grant CNS-0709296. We are also grateful to our partner: Silverado Senior Living - The Huntington, Alhambra, CA, USA.

References

1. American Alzheimer Association. About alzheimer's disease statistics. *American Alzheimer Association*, November 2007.
2. 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.
3. 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.
4. 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.
5. Nourbakhsh I. Fong T. and Dautenhahn K. A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42(3-4):143–166, 2003.
6. C. Kidd, W. Taggart, and S. Turkle. A sociable robot to encourage social interaction among the elderly. In *IEEE International Conference on Robotics and Automation (ICRA)*, Orlando, USA, May 2006.
7. A. Libin and J. Cohen-Mansfield. Therapeutic robocat for nursing home residents with dementia: Preliminary inquiry. *Am J Alzheimers Dis Other Demen*, 19(2):111–116, 2004.
8. P. Marti, L. Giusti, and M. Bacigalupo. Dialogues beyond words. *Interaction Studies*, 2008.
9. F. Michaud and C. Theberge-Turmel. Mobile robotic toys and autism. In A. Billard, K. Dautenhahn, L. Canamero, and B. Edmonds, editors, *Socially Intelligent Agents - Creating Relationships with Computers and Robots*. Kluwer Academic Publishers, 2002.
10. 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.
11. A. Tapus, C. Tapus, and Maja J Matarić. User-robot personality matching and robot behavior adaptation for post-stroke rehabilitation therapy. *Intelligent Service Robotics*, 1(2):169–183, April 2008.
12. M. Walton. Meet paro, the therapeutic robot seal. In *CNN*, 2003.
13. H. Yanco. Evaluating the performance of assistive robotic systems. In *Proc. of the Workshop on Performance Metrics for Intelligent Systems*, Gaithersburg, MD, USA, 2002.