Interactive Learning and Adaptation for Robot Assisted Therapy for People with Dementia

Konstantinos Tsiakas
HERACLEIA Lab
CSE Department
University of Texas at Arlington
konstantinos.tsiakas@mavs.uta.edu

Cheryl Abellanoza
Cognitive Neuroscience of
Memory Lab
Department of Psychology
University of Texas at Arlington
cheryl.abellanoza@mavs.uta.edu

Fillia Makedon
HERACLEIA Lab
CSE Department
University of Texas at Arlington
makedon@uta.edu

ABSTRACT

In this paper, we present an adaptive cognitive music game designed to monitor and improve the attention levels of people with dementia. The goal of this game is to provide a customized protocol based on user needs and preferences, following the Reinforcement Learning (RL) framework. The game adjusts its parameters (e.g., difficulty level) so as to help the user complete the task successfully, while keeping them engaged. The main contribution of this paper is an interactive learning and adaptation framework that enables and facilitates the adaptation of the robot behavior towards new users, providing a safe, tailored and efficient interaction.

CCS Concepts

•Computing methodologies → Reinforcement learning; Learning from implicit feedback; •Human-centered computing → Human computer interaction (HCI);

Keywords

Interactive Reinforcement Learning, Policy Adaptation, Robot Learning and Behavior Adaptation, Robot Assisted Therapy, Music Therapy

1. INTRODUCTION

Alzheimer's disease is a form of dementia thought to be caused by neurodegeneration due to plaques and tangles that collect abnormally in the brain [15]. Alzheimer's symptoms include the chronic, incurable loss of memory and cognitive function, as well as problems with communication and mood [20]. It is estimated that the number of Alzheimer's disease patients worldwide ranges from about 30 to 40 million [1].

Given this worldwide impact, much focus has been placed on researching pharmacological and biological treatments; however, these treatments often fail to combat memory loss or to address the quality of life for the patient [21, 7]. Pre-

vious work has been done in the investigation of alternative therapy modalities. In [16] the authors present a therapeutic *robocat* to see how patients with dementia would respond. Their interactions with the robocat led to less agitation, and more positive experiences. Similarly, in [29] the authors present *Paro*, a robotic seal that interacted proactively and reactively with patients, leading to improvements in various symptoms.

There is also a growing trend in research on cognitive and multisensorial therapies, which allow patients to experience more personalized, complete care, while also offering higher rates of responsiveness to treatment. For instance, in [11], the authors argue that music therapy was well received by patients, resulting to significant decreases in severity, frequency, and disruptiveness of symptoms (compared to controls) as measured by the Neuropsychiatric Inventory (NPI-NH). Further, music therapy positively influenced not only cognitive states but also related mood states, like depression and anxiety [18].

This indicates that there is a need for personalized and tailored robotic assistants that interact with people suffering from dementia. The goal of such systems is to improve symptoms and quality of life for such people, as well as to provide a personalized protocol by providing users with motivation, encouragement, and companionship [26]. Such therapies are also being assessed for other chronic diseases that affect cognition and mood (e.g., Autism Spectrum Disorder, cancer). Given this expanding body of research, a possible computational paradigm to quantify the effects of such therapies is presented in [22].

In this work, we present an adaptive cognitive music game that employs a NAO robot¹ that monitors, instructs, and adapts to user abilities, aiming to encourage task improvement and attention training. We focus on the adaptation capabilities of such a system, proposing an interactive learning and adaptation framework that enables the robot to adapt its behavior towards each specific user, following the Reinforcement Learning paradigm.

The paper is organized as follows: In Section 2, we present several approaches for robot learning and behavior adaptation in the area of robot assisted therapy, showing the motivation of our work. In Section 3, we present our proposed adaptation framework. In Section 4, we illustrate the framework by presenting the cognitive music game.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PETRA '16, June 29-July 01, 2016, Corfu Island, Greece

© 2016 ACM. ISBN 978-1-4503-4337-4/16/06...\$15.00

DOI: <http://dx.doi.org/10.1145/2910674.2935849>

¹<https://www.aldebaran.com/en/cool-robots/nao>

2. ROBOT LEARNING AND ADAPTATION

Robot Assisted Therapy has been extensively applied to assist users with cognitive impairments [8]. There are works that indicate that personalized robotic assistive systems can establish a productive interaction with the user, improving the effects of a therapy session [26]. Two significant attributes of interactive systems or *agents* that interact with users during a training or therapy session, are *personalization* and *safety*. An interactive agent should be able to adapt to user behavior preferences and needs. More importantly, it should be robust to user behavior changes, while also ensuring a safe interaction [10].

Reinforcement Learning is an appropriate framework for modeling and optimizing the behavior of an agent that interacts with a user [17, 28]. However, an agent that interacts with many users should be trained to cope with various users. One approach is to initially train the agent against simulated users, explore different policies, and then deploy the system with real users, enabling the agent to refine its policy for different preferences and abilities of all different users.

Another approach is to build different user models (or *personas*) that capture possible preference profiles. Consequently, it is assumed that the model to which a particular user belongs to is known prior to the interaction. However, this approach is more appropriate for long-term HRI applications, where the agent adapts over repeated interactions with the users [4]. In this work, we focus on agents that are able to interactively reuse and modify a learned policy during a short-term interaction, without requiring an existing model of the specific user.

3. INTERACTIVE LEARNING AND ADAPTATION FRAMEWORK

As previously mentioned, an important feature of an interactive agent is the ability to adapt to new users. If an agent is trained to interact with a specific user, the learned policy will not be as efficient to a different user, since each action may have different effects on them. Moreover, in real-world applications, the user is not a deterministic environment, since user dynamics (user intentions, preferences, abilities, etc.) are subject to change over time.

Our work moves towards the definition and implementation of a framework for interactive learning and adaptation [27]. Based on this framework, an interactive agent is able to adapt a learned policy towards a new user, by exploiting additional communication channels (i.e., feedback and guidance) provided during the interaction. More specifically, the framework utilizes implicit feedback provided by the primary user to refine its learned policy towards the current user. However, a key challenge is to ensure a safe interaction while adapting the agent's behavior to a different user.

Our framework supports the participation of a secondary user, as a supervisor, that can guide the interaction in its early steps, avoiding unsafe interactions. The supervisor can either physically or remotely supervise the interaction. A user interface can be used to provide the supervisor with useful information, to help them monitor the interaction and enhance their own decision making, before altering the agent's policy. The goal of this framework is to enable agents to learn as long as they interact with primary and secondary users, adapting and refining their policy dynamically.

4. ADAPTIVE COGNITIVE MUSIC GAME

In this section, we present the adaptive cognitive music game that we will use for our experiments. We describe the game setup and procedure, as well as how we model the robot's behavior following the RL paradigm, in order to train the system for specific users. Then, we illustrate the proposed framework for dynamical adaptation to new users, exploiting additional communication channels (i.e., feedback and guidance).

4.1 Definition of the Cognitive Game

To illustrate our proposed framework, we present a variation of a cognitive music game called *Name that Tune*, as presented in [26]. This game is designed to improve the participant's level of attention. Participants have five buttons in front of them; the four buttons correspond to a song excerpt and the last one to silence (no song excerpt). Each button has the corresponding song title written on it (or 'SILENCE' for the last button). The user listens to a music collection of these four songs and is asked to identify the appropriate song excerpt (or silence), find the correct button, and push it as soon as possible. The system measures reaction time and correctness to evaluate user's performance.

A training session consists of 3 stages. Each stage consists of $N = [4-10]$ song excerpts, including silence (each song excerpt is followed by silence). The order of the song excerpts is random. Each song excerpt is played for a predefined time or until the user pushes the corresponding button. After each song excerpt, a silence excerpt follows.

The goal of the proposed system is to adjust the game parameters so as to encourage task improvement and attention training. There are three difficulty levels, based on the hint the robot provides. At the EASY level, the robot tells the user which button to push. At the MEDIUM level, the robot tells the user to push a button, without indicating which one needs to be pushed. At the HARD level, no hint is provided. The robot models the difficulty level based on the user's performance during the interaction; this allows the robot to assist the user to finish the task, while keeping the user engaged.

4.2 Robot Learning

Figure 1: The game as an Reinforcement Learning task. The state-action space describes the interaction during the game.

In order to model robot’s behavior, we follow the Reinforcement Learning approach [25]. We formulate the task as a Markov Decision Process (MDP). The state space includes all the relevant information needed from the robot to decide its next action. The state features describe the current game state (stageID, songID, difficulty) as well as the user’s current performance (reaction time and correctness). Based on these state features, the system must find the appropriate action of each state, so as to maximize user performance and engagement.

In [26], the authors proposed an adaptation mechanism that evaluates user performance after each stage and uses this information to adapt forthcoming sessions. Our formulation enables dynamical adaptation during the interaction, making it a challenging problem due to the large state-action space (Fig. 1).

4.3 Behavior Adaptation

As we already mentioned, an interactive system that interacts with many users should have the ability to adapt its behavior towards different users. In this section, we illustrate our framework [27] from Section 3, for the adaptive music game. The framework is shown in Fig. 2.

The goal of this framework is to enable dynamical adaptation of robot’s behavior, during each interaction, to different users. The system uses a learned policy and utilizes two additional communication channels to adapt the current policy efficiently and safely. We follow two Interactive Reinforcement Learning approaches; *Learning from Feedback* [14] and *Learning from Guidance* [5]. We argue that a proper combination of these two techniques and their integration to the adaptation mechanism can facilitate the safe adaptation of the agent towards the current user, exploiting human knowledge provided through these two communication channels.

Figure 2: Adaptation framework. This framework extends the RL setup, including two additional communication channels, used for adaptation

We argue that feedback can be considered as a personalization factor. We want to exploit user feedback so as to adapt the game to user-specific needs. One way to personalize this process is to measure user engagement, which can indicate how active the user is when interacting with the game, as well as how effective the game is. We propose to

use the Muse EEG headset², a commercially available tool that measures electrical activity at the scalp as it relates to various cognitive and mood states. This type of sensor can be used to measure how engaged a person is while that person is completing some sort of games or music tasks [12, 19]. It is shown that increased frontal alpha and theta activity are predictive measures of engagement; sensors like Muse can be used to measure this type of activity [2] and also in self-calibrating protocols [13].

Guidance is provided in the form of corrective or suggested actions, by the therapist. In this way, the agent adapts to the user preferences and needs and offers the therapist the ability to intervene in order to ensure safe interactions. Moreover, the therapist can set their own therapeutic goals by dynamically altering the agent policy during the interaction.

5. DISCUSSION AND FUTURE WORK

Since this is an ongoing work, several aspects of the game should be considered. For instance, the type of music that is selected for the game may play a role in symptom reduction. In [9], it was found that personalized music choices showed a significant reduction in patient agitation, but classical music did not show the same reduction. This makes sense; as memory is highly individualized, so are the connections that users will make [24]. As such, the personalization of the game, as well as the personalization of the music, would create the best therapeutic tool.

Moreover, the gameplay itself should be carefully regarded. People who suffer from Alzheimer’s disease and dementia are open to more memory interference than their healthy counterparts [6]. Since this population has cognitive impairments in remembering more recent information (as compared to older information), the game should be easy to teach multiple times, with many built-in prompts for each step. To assess symptoms, performance for more complex steps in the game (e.g., having the robot play all music sequences first, then asking the patient to recall the order in which those sequences were played) can be measured to see if memory problems are stabilizing or worsening. Also, having removable labels for the buttons can help the therapist gauge how much the patient is able to retain while learning the game. Also, longitudinal studies of these therapies are greatly needed. In [3], it was shown that music therapy was helpful in reducing symptoms in a short-term period; however, this reduction was not long-lasting. Future research can examine how to expand the longevity of these important symptom reductions.

Considering the adaptation framework, we need to investigate how feedback can be used to efficiently modify a learned policy. Feedback must be handled as a policy modifier and not as an additional reinforcement signal, since it may not alter the policy. Guidance can be seen as a human-guided exploration mechanism [23]. A significant aspect of guidance is the workload of the therapist, that should reduce over time, indicating that the agent converges to an optimal policy. Active Learning methods can be used to learn, based on state information (i.e., state uncertainty and importance), when the therapist should intervene.

To conclude with, we proposed an interactive learning and adaptation framework for dynamically adaptive robot as-

²<http://www.choosemuse.com/>

sisted therapy. We presented our use case, a cognitive music game. Our next steps include the implementation of the game and a case study with participants to evaluate the game itself, as well as the proposed framework.

6. ACKNOWLEDGMENTS

This material is based upon work supported by NSF under award numbers CNS 1338118, 1035913 and by the educational program of NCSR Demokritos in collaboration with the University of Texas at Arlington.

7. REFERENCES

- [1] Global health and aging - the burden of dementia. <https://www.nia.nih.gov/research/publication/longer-lives-and-disability/burden-dementia>, 2011.
- [2] M. Abujelala, C. Abellanoza, A. Sharma, and F. Makedon. Brain-ee: Brain enjoyment evaluation using commercial eeg headband. In *Proceedings of the 9th ACM International Conference on Pervasive Technologies Related to Assistive Environments*. ACM, 2016.
- [3] A. Baker. An investigation of long-term effects of group music therapy on agitation levels of people with alzheimer’s disease. volume 11, pages 330–338, 2007.
- [4] K. Baraka and M. Veloso. Adaptive interaction of persistent robots to user temporal preferences. In *Social Robotics*, pages 61–71. Springer, 2015.
- [5] F. Cruz et al. Interactive reinforcement learning through speech guidance in a domestic scenario. In *Proc. IJCNN*, pages 1–8. IEEE, 2015.
- [6] M. Dewar et al. Insights into spared memory capacity in amnesic mci and alzheimer’s disease via minimal interference. volume 78, pages 189–199, 2011.
- [7] S. Douglas, I. James, and C. Ballard. Non-pharmacological interventions in dementia. volume 10, pages 171–179, 2004.
- [8] P. Georgakopoulos, M. Chatzidimitriou, and M. Tsolaki. A review of interventions with assistive technologies for patients with cognitive impairment. *Handbook of Research on Innovations in the Diagnosis and Treatment of Dementia*, page 118, 2015.
- [9] L. Gerdner. Effects of individualized versus classical “relaxation” music on the frequency of agitation in elderly persons with alzheimer’s disease and related disorders. volume 12, pages 49–65, 2000.
- [10] N. Giullian et al. Detailed requirements for robots in autism therapy. In *Proc SMC 2010*, pages 2595–2602. IEEE.
- [11] M. H. Hsu, R. Flowerdew, M. Parker, et al. Individual music therapy for managing neuropsychiatric symptoms for people with dementia and their carers: a cluster randomised controlled feasibility study. *BMC geriatrics*, 2015.
- [12] N. Hurless et al. Music genre preference and tempo alter alpha and beta waves in human non-musicians. *The Premier Undergraduate Neuroscience Journal*, 2013.
- [13] T. Karydis et al. Performance characterization of self-calibrating protocols for wearable eeg applications. In *Proceedings of the 8th ACM International Conference on Pervasive Technologies Related to Assistive Environments*, page 38. ACM, 2015.
- [14] W. B. Knox and P. Stone. Reinforcement learning from simultaneous human and mdp reward. In *Proc. AAMAS*, pages 475–482, 2012.
- [15] K. Kurbatskaya et al. Upregulation of calpain activity precedes tau phosphorylation and loss of synaptic proteins in alzheimer disease brain. *Acta Neuropathologica Communications*, 2016.
- [16] A. Libin and J. Cohen-Mansfield. Therapeutic robocat for nursing home residents with dementia: preliminary inquiry. *American journal of Alzheimer’s disease and other dementias*, 19(2):111–116, 2004.
- [17] H. Modares et al. Optimized assistive human-robot interaction using reinforcement learning. 2015.
- [18] L. Ozdemir and N. Akdemir. Effects of multisensory stimulation on cognition, depression and anxiety levels of mildly-affected alzheimer’s patients. *Journal of the neurological sciences*, 283(1):211–213, 2009.
- [19] A. Plotnikov et al. Exploiting real-time eeg analysis for assessing flow in games. In *Advanced Learning Technologies (ICALT), 2012 IEEE 12th International Conference on*, pages 688–689. IEEE, 2012.
- [20] M. Prince, M. Guerchet, M. Prina, B. U. P. Association, et al. *World Alzheimer report 2013: Journey of caring: an analysis of long-term care for dementia*. Alzheimer’s Disease International, 2013.
- [21] S. Samson et al. Efficacy of musical interventions in dementia: methodological requirements of nonpharmacological trials. *Annals of the New York Academy of Sciences*, 1337(1):249–255, 2015.
- [22] B. Sandak et al. Computational paradigm to elucidate the effects of arts-based approaches and interventions: Individual and collective emerging behaviors in artwork construction. *PloS one*, 10(6):e0126467, 2015.
- [23] E. Senft, P. Baxter, and T. Belpaeme. Human-guided learning of social action selection for robot-assisted therapy. In *4th Workshop on Machine Learning for Interactive Systems*, 2015.
- [24] S. Smith and E. Vela. Environmental context-dependent memory: A review and meta-analysis. volume 8, pages 203–220, 2001.
- [25] R. S. Sutton and A. G. Barto. *Introduction to reinforcement learning*, volume 135. MIT Press Cambridge, 1998.
- [26] A. Tapus. Improving the quality of life of people with dementia through the use of socially assistive robots. In *Advanced Technologies for Enhanced Quality of Life, 2009. AT-EQUAL’09.*, pages 81–86. IEEE, 2009.
- [27] K. Tsiakas. Facilitating safe adaptation of interactive agents using interactive reinforcement learning. In *Companion Publication of the 21st International Conference on Intelligent User Interfaces*, pages 106–109. ACM, 2016.
- [28] K. Tsiakas, M. Huber, and F. Makedon. A multimodal adaptive session manager for physical rehabilitation exercising. In *Proceedings of the 8th ACM International Conference on Pervasive Technologies Related to Assistive Environments*. ACM, 2015.
- [29] K. Wada, T. Shibata, et al. Robot therapy for elders affected by dementia. *Engineering in Medicine and Biology Magazine, IEEE*, 2008.