

Implementation and Evaluation of a Motivational Robotic Coach for Repetitive Stroke Rehabilitation

Ross, Martin K.; Broz, Frank; Baillie, Lynne

DOI

[10.1145/3610978.3640646](https://doi.org/10.1145/3610978.3640646)

Publication date

2024

Document Version

Final published version

Published in

HRI 2024 Companion - Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction

Citation (APA)

Ross, M. K., Broz, F., & Baillie, L. (2024). Implementation and Evaluation of a Motivational Robotic Coach for Repetitive Stroke Rehabilitation. In *HRI 2024 Companion - Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction* (pp. 911-915). (ACM/IEEE International Conference on Human-Robot Interaction). IEEE. <https://doi.org/10.1145/3610978.3640646>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



Implementation and Evaluation of a Motivational Robotic Coach for Repetitive Stroke Rehabilitation

Martin K. Ross
School of Mathematical and
Computer Sciences
Heriot-Watt University
Edinburgh, United Kingdom
mkr30@hw.ac.uk

Frank Broz
Intelligent Systems
Delft University of Technology
Delft, Netherlands
f.broz@tudelft.nl

Lynne Baillie
School of Mathematical and
Computer Sciences
Heriot-Watt University
Edinburgh, United Kingdom
l.baillie@hw.ac.uk

ABSTRACT

Repetitive, individual exercises can improve the functional ability of stroke survivors over the long term. With the aim of providing extra motivation to adhere to repetitive, individual rehabilitation, this paper presents a robotic coach for stroke rehabilitation. Our system uses the Pepper robot and performs one of twelve data-driven coaching policies. The policies were learned from human-human observations of professional stroke physiotherapists and provide high-level personalisation based on user information and training context. A within subjects evaluation of the system was conducted in-person involving short interactions with 3 stroke survivors. The system was able to engage the target end users and there were indications that decreased workload could be possible when using the system compared to exercising alone.

CCS CONCEPTS

• Computer systems organization ~ Embedded and cyber-physical systems ~ Robotics ~ Robotic control • **Human-centered computing** ~ **Human computer interaction (HCI)**

KEYWORDS

Coaching, Personalisation, Rehabilitation, Stroke

ACM Reference format:

Martin K. Ross, Frank Broz and Lynne Baillie. Implementation and Evaluation of a Motivational Robotic Coach for Repetitive Stroke Rehabilitation. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '23 Companion)*, March 11-14, 2024, Boulder, CO, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3610978.3640646>

1 INTRODUCTION

In rehabilitation after stroke, adherence tends to decrease after discharge from hospital due to, among other reasons, a lack of motivation [1]. A personalised robotic coach could provide extra

motivation to adhere to an exercise routine and therefore increase the potential of functional improvements over the long term.

Previous work has shown the potential of using a Socially Assistive Robot (SAR) to motivate users during physical exercise (e.g. [2]–[4]). However, past systems tend not to offer technical advice on a specific skill or are limited in their personalisation.

We have built on the work of [5], using a novel development and personalisation approach to implement and evaluate a robotic rehabilitation coach on the Pepper platform. We use the term ‘personalised’ in reference to high-level personalisation to groups of users but acknowledge that continual, low-level adaption to individuals would need added to meet the full personalisation requirement [6]. This is the first work to explore this personalisation approach in rehabilitation and complements the results of [7] who performed a similar study in sports training.

2 BACKGROUND

2.1 Stroke Rehabilitation

Stroke is a sudden and devastating medical condition leaving more than half of survivors with permanent disabilities [9], and almost half feeling abandoned after they leave hospital [10]. Individual rehabilitation after a stroke is often done alone, contains repetitive exercises and is performed frequently over a long period of time. Motivation for rehabilitation can be affected by the actions of a therapist [12] [14]. In this work, we explore if a robotic coach used during individual exercise could have a similar effect.

Examples of past work in this area include a fully autonomous robotic coach (Baxter) leading stroke survivors through a wire puzzle task [4], and a gamified robotic system (Pepper) being rated highly by participants in a series of stroke rehabilitation tasks over 15 interactions [8]. These studies point to the effectiveness of a SAR providing motivation during physical activity. However, the evaluation conducted by [4] focussed more on the efficacy of the wire puzzle task than the robot, and the personalisation conducted by [8] was set up ahead of time by the researchers.

2.2 Personalisation of SARs

Personalisation has been suggested as a requirement of a robotic coach [6]. One promising method of achieving personalisation to



This work is licensed under a Creative Commons Attribution International 4.0 License.

HRI '24 Companion, March 11–14, 2024, Boulder, CO, USA.
© 2024 Copyright is held by the owner/author(s).
ACM ISBN 979-8-4007-0323-2/24/03.
<https://doi.org/10.1145/3610978.3640646>

groups of users is to learn from human demonstrations. In a collaborative packing task, Nikolaidis et al. showed that by clustering human demonstrations into similar styles and applying inverse reinforcement learning over those clusters, it was possible to learn a reward function that was representative of each user type [9]. In this work, we evaluate the application of a similar strategy to the more interactive, open scenario of coaching.

By combining data collection methods adapted from sports coaching literature with computational techniques and mathematical modelling, Ross et al. defined a process to formalise human knowledge in the form of ‘coaching policies’ usable for robotic control [5]. A policy refers to a mapping from states to actions. Starting from observations of professional squash coaches and stroke physiotherapists (Ross’ study was cross-domain), they used Nikolaidis et al.’s clustering algorithm to generate 12 unique coaching policies that can be viewed as ‘behaviour graphs’ on GitHub.¹ They then utilised qualitative data obtained through semi-structured interviews with the observed coaches and physiotherapists to make actionable suggestions as to which policies were likely to be more appropriate for which groups of users. The robotic coach evaluated in this work uses these suggestions to select from Ross et al.’s 12 policies and achieve personalisation based on user information and training context.

This type of personalisation goes beyond that previously explored in the context of robotic coaching. Most past works have focussed on customising the utterances of a system with the name and performance history of its user (e.g. [10], [11]). While findings suggest that a personalised robotic coach can increase adherence to [10], and enjoyment of [11], interactions with a robotic exercise coach, we went a step further. In this work, we attempted to predict the style of interaction (i.e. the behaviours used by the robot during a session) that each user would prefer.

3 SYSTEM DESCRIPTION AND IMPLEMENTATION

3.1 Overview of the System

The robotic coaching system was implemented on a Pepper robot using NAOqi Python API. It receives data from a human operator using Wizard of Oz (WoZ) techniques, which it uses to formulate appropriate behaviours to coach users during individual rehabilitation sessions. These sessions comprise sets of repetitions of a particular exercise, sandwiched between introduction and feedback sequences from the robot. By performing a range of coaching behaviours similar to those performed by a human physiotherapist, the robot leads its user through their rehabilitation session. Behaviours are primarily animated utterances spoken by the robot, but also include demonstrations via the robot’s movements. For example, the robot might perform a pre-instruction behaviour, praise, or ask a question while demonstrating the correct arm position for a certain exercise. Throughout the session, Pepper’s tablet screen is used to display subtitles of Pepper’s utterances and images of the current exercise.

3.2 System Implementation

Figure 1 shows the architecture of the robotic coach. Following are details on the implementation of each part of the architecture.




Figure 1: Architecture diagram of the robotic coach.

3.2.1 Processing Layer. The processing layer is composed of two main blocks: the controller and the coaching policy.

The controller coordinates everything in the system and communicates with the interface and tracking layers. It is implemented using a behaviour tree, a structure often used for robotic control [12]. The behaviour tree structure drives the format of each session, which mimics what was seen during observations of human coaches and physiotherapists in [5]. The format involves: introduction by the robot; instruction by the robot to perform a set of an exercise; feedback during the set; and feedback on the set just performed. Within the leaf nodes of the tree (where execution of coaching actions occurs) are calls to the coaching policy module.

The coaching policy selects an action for the robot to perform in its current state. The policy is formulated as a $T(a^t, a^{t+1})$ matrix where a denotes an action and $t+1$ denotes the next timestep. The available actions are 13 of the most frequently used behaviour categories in the observation instrument used by [5]. Example actions are given in Table 1. The state space consists of the system’s previous action combined with the stage of the interaction. At each timestep, the system selects an action from the next state distribution as defined by the T matrix.

The choice of policy to execute is made at the start of each coaching session. It is based on the user information categories identified by coaches and physiotherapists in [5]: the user’s level of impairment (self-rated), number of interactions with the robotic coach (i.e. length of the relationship), motivation for conducting rehabilitation (self-rated), and type of session. Each of these information categories was split into a ‘high’ and ‘low’ value, and if the value matched the recommendations given by [8], the policy received a higher score. The policy with the highest score was used by the system during the coaching session. Thus, the behaviour of the robotic coach is based on the raw data obtained from the HHI observations conducted by Ross et al. [5] but the interaction style is personalised to each user.

Once an action has been selected by the policy module, the controller formats it into a robotic behaviour (i.e. an utterance) that incorporates data from the tracking layer.

¹ <https://github.com/M4rtinR/Behaviour-Graph-Visualisations>

Table 1: Examples of actions and robotic behaviours.

Action	Example Behaviour
Pre-Instruction	“In this exercise, you can lean forward with your body if you feel comfortable doing so.”
Concurrent Instruction (Positive)	“Slow, controlled movements.”
Post Instruction (Negative)	“Your table-top circles weren’t quite as good there.”
Questioning	“How did your towel slides feel there? Touch the back of my hand if they felt good or the top of my head if you think more work is needed.”
Positive Modelling	Demonstrates the correct way to perform shoulder openers.
Praise	“Good”

3.2.2 Tracking Layer. This layer consists of a human operator who indicated completion of an exercise repetition during the session. This is the only part of the system that was run using WoZ techniques and was designed this way to remove the potential for failure of a vision system or sensor. If an appropriate sensing system was developed it could easily be plugged in to the system to replace the human operator and would allow the system to gather much more information usable in its feedback.

3.2.3 Interface Layer. The user directly interacts with this layer. Communication between the user and the system is done through Pepper using the robot’s touch sensors and in-built text-to-speech technology. Actions selected by the processing layer are performed by the robot and lead the user through their session. The specific utterances and demonstrations that correspond to each action were selected at random from four available options for each combination of action, stage of interaction and user performance. This resulted in over 1,000 possible utterances.

The differences in the coaching policies mean that the robot’s behaviour is personalised to the user. For example, some policies contain lots of instruction actions (likely to be chosen for a user who is early in their rehabilitation journey), whereas others would have the robot ask more questions to have the user figure things out for themselves or recall how an exercise was performed in a previous session.

4. PILOT STUDY EVALUATION

4.1 Participants

Independently living stroke survivors (N=3) were recruited with the help of local charities using the following inclusion criteria:

1. Have a stroke-related arm impairment.
2. Have the required cognitive ability to provide informed consent for the study.
3. Are living independently (i.e. in their own home, not in hospital or a care home).

Each participant entered a raffle for a £50 Amazon voucher.

4.2 Conditions

Three conditions were evaluated during the study: two coaching conditions and one baseline condition. In the **Data Selected Policy (DSP)** condition, participants interacted with the robot executing the coaching policy chosen using the method described in Section 3.2.1. In the **Non-Personalised Policy (NPP)** condition, the robot executed a randomly selected policy from the other 11 policies that were not the best match for the participant’s information and training context. Comparing these two conditions allowed us to discover the effect of high-level personalisation. The selection of the random policy was performed at the beginning of the interaction. These two coaching conditions are in contrast with the **No Coaching Policy (NCP)** baseline condition in which the robot told the user which exercise to perform, and when to perform each set (see Section 4.4) but gave no coaching behaviours. It was the closest condition to a regular individual rehabilitation session.

4.3 Measures

The following measures were used in the study. The CBS-S and IMI used 7-point scales and the NASA TLX used a 21-point scale.

4.3.1 Coaching Behaviour Scale for Sport (CBS-S) [13]. The “technical skills” subscale of the CBS-S was used to measure participants’ opinions on the coaching provided by the robot.

4.3.2 Intrinsic Motivation Inventory (IMI) [14]. The interest/enjoyment, perceived competence, perceived choice and value /usefulness subscales of the IMI were used to assess the effect of each condition on the participants’ intrinsic motivation for conducting an individual rehabilitation session.

4.3.3 NASA Task Load Index (TLX) [15] was used as a measure of workload during a rehabilitation session with a robotic coach.

4.4 Study Design

The sessions for 2 participants were conducted in the sports centre at the university, and for the other at a local respite centre. The setup used in the study is shown in Figure 2. All necessary equipment was provided and sanitised between sessions. At least 1 week prior to the study, participants were sent an information sheet outlining the study procedure and a consent form. The study received full ethical approval from the university’s ethics board.




Figure 2: The experimental setup shown from the researcher’s perspective at the back of the room.

Table 2 shows an overview of the procedure. A within subject design was used with the order of conditions counterbalanced and interactions split across 2 separate days. In each session the robot asked participants to perform 2 sets of 4 exercises (chosen with input from physiotherapists: external rotations with cane, shoulder openers, towel slides, and table-top circles) in a random order. The first set consisted of 10 repetitions and the second of 5.

Table 2: Summary of the evaluation procedure used.

Phase	Activity	Time (mins)
<i>Intro</i>	Researcher explained the study and answered questions. Participant filled in the consent form & demographic questionnaire.	10
<i>Sess. 1</i>	Participant conducted sets of exercises led by the robotic coach.	10-25
<i>Qs 1</i>	Participant completed the questionnaires.	10
Day 1 total: 30-45 mins		
<i>Sess. 2</i>	As in <i>Sess. 1</i> using a different condition.	10-25
<i>Qs 2</i>	As in <i>Qs 1</i> . Researcher set up the next condition. Participant took a break.	30-40
<i>Sess. 3</i>	As in <i>Sess. 1</i> using the third condition.	10-25
<i>Qs 3</i>	As in <i>Qs 1</i> .	10
<i>Wrap-up</i>	Researcher answered any final questions the participant had.	5
Day 2 total: 1 hr 5 mins – 1 hr 45 mins		

4.5 Results

One of the participants opted not to complete the final session (NPP condition) but as they did not experience any problems with the system, and gave permission for their data to be disseminated, they have been included in the analysis shown in Figure 3.

While no statistical significance can be shown due to the small number of participants, the NCP condition performed worse than both coaching conditions in 5 of the 6 NASA TLX measures. Participants perceived the same exercises as more mentally and physically demanding, while perceiving their performance on the exercises as lower when the robot did not offer any coaching.

There were minimal differences between conditions in the IMI scores given by participants. Meanwhile, each condition scored well in different aspects of perceived coaching effectiveness, as measured by the CBS-S. Interestingly, the NCP condition was perceived to give more advice, but the NPP condition scored highest in reinforcement and technical feedback, whilst the DSP condition was perceived as the best at giving immediate feedback.

5. DISCUSSION AND CONCLUSION

This paper has presented the evaluation of a novel robotic coach for individual rehabilitation after stroke. We acknowledge the lack of participants as a limitation. However, the system was able to engage three stroke survivors in rehabilitation with limited WoZ input from a human operator. This shows the effectiveness

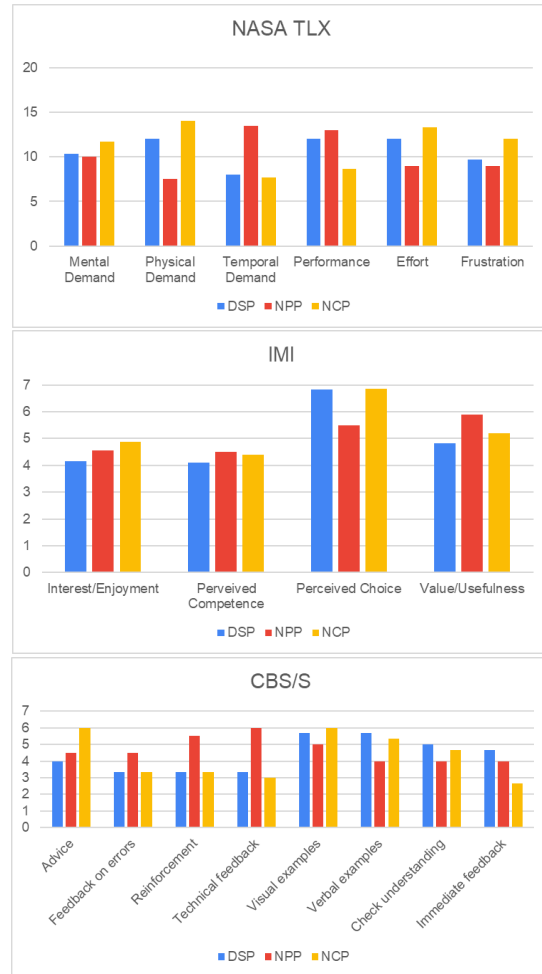


Figure 3: Results from the 3 scales. Values are mean.

of the implementation approach proposed by Ross et al. [5], and adds to the results obtained in squash coaching by [7].

Interestingly, personalisation of the system’s coaching policy did not result in improvements in motivation or perceived coaching effectiveness. It is important to note that the NPP condition also used a policy based on human coaching data. The rigid experimental setup and limited number of participants resulted in few differences in participants’ user information and training context and therefore there was little personalisation possible in the DSP condition.

Future work will learn from studies such as [16] by further adapting the selected policy to individuals over time using reinforcement learning. Evaluations of the system with more participants and over the long-term is also a clear next step.

ACKNOWLEDGMENTS

The authors would like to thank Chest Heart & Stroke Scotland and Leuchie House for their help with participant recruitment. This work is funded by the Engineering and Physical Sciences Research Council (EPSRC). Grant ID: EPSRC DTP18.

REFERENCES

- [1] M. T. Jurkiewicz, S. Marzolini, and P. Oh, "Adherence to a Home-Based Exercise Program for Individuals After Stroke," *Topics in Stroke Rehabilitation*, vol. 18, no. 3, pp. 277–284, 2011, doi: 10.1310/tsr1803-277.
- [2] L. Sussenbach *et al.*, "A robot as fitness companion: Towards an interactive action-based motivation model," *IEEE RO-MAN 2014 - 23rd IEEE International Symposium on Robot and Human Interactive Communication: Human-Robot Co-Existence: Adaptive Interfaces and Systems for Daily Life, Therapy, Assistance and Socially Engaging Interactions*, pp. 286–293, 2014, doi: 10.1109/ROMAN.2014.6926267.
- [3] K. Winkle, S. Lemaignan, P. Caleb-Solly, U. Leonards, A. Turton, and P. Bremner, "Couch to 5km Robot Coach: An Autonomous, Human-Trained Socially Assistive Robot," in *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, in HRI '20. Cambridge, United Kingdom: Association for Computing Machinery, Mar. 2020, pp. 520–522. doi: 10.1145/3371382.3378337.
- [4] E. Wade, A. R. Parnandi, and M. J. Mataric, "Using Socially Assistive Robotics to Augment Motor Task Performance in Individuals Post - Stroke," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011, pp. 2403–2408.
- [5] M. K. Ross, F. Broz, and L. Baillie, "Observing and Clustering Coaching Behaviours to Inform the Design of a Personalised Robotic Coach," in *Proceedings of the 23rd International Conference on Human-Computer Interaction with Mobile Devices and Services*, Virtual (originally Toulouse, France): ACM, Sep. 2021. doi: 10.1145/3447526.3472043.
- [6] K. Winkle, P. Caleb-Solly, A. Turton, and P. Bremner, "Social Robots for Engagement in Rehabilitative Therapies: Design Implications from a Study with Therapists," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction - HRI '18*, 2018, pp. 289–297. doi: 10.1145/3171221.3171273.
- [7] M. K. Ross, F. Broz, and L. Baillie, "Individual Squash Training is More Effective and Social with a Humanoid Robotic Coach," in *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, Aug. 2023.
- [8] R. Feingold Polak and S. L. Tzedek, "Social Robot for Rehabilitation: Expert Clinicians and Post-Stroke Patients' Evaluation Following a Long-Term Intervention," in *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, in HRI '20. Cambridge, United Kingdom: Association for Computing Machinery, Mar. 2020, pp. 151–160. doi: 10.1145/3319502.3374797.
- [9] S. Nikolaidis, R. Ramakrishnan, K. Gu, and J. Shah, "Efficient Model Learning from Joint-Action Demonstrations for Human-Robot Collaborative Tasks," in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI '15*, ACM, Mar. 2015, pp. 189–196. Accessed: Mar. 02, 2020. [Online]. Available: <https://dl.acm.org/doi/10.1145/2696454.2696455>
- [10] N. Céspedes *et al.*, "A Socially Assistive Robot for Long-Term Cardiac Rehabilitation in the Real World," *Frontiers in Neurobotics*, vol. 15, 2021, Accessed: Jul. 28, 2022. [Online]. Available: <https://www.frontiersin.org/articles/10.3389/fnbot.2021.633248>
- [11] J. Fasola and M. J. Mataric, "Using Socially Assistive Human-Robot Interaction to Motivate Physical Exercise for Older Adults," in *Household Service Robotics*, vol. 100, no. 8, 2014, pp. 487–516. doi: 10.1016/B978-0-12-800881-2.00022-0.
- [12] M. Iovino, E. Scukins, J. Styurd, P. Ögren, and C. Smith, "A Survey of Behavior Trees in Robotics and AI," *arXiv:2005.05842 [cs]*, May 2020, Accessed: Apr. 01, 2021. [Online]. Available: <http://arxiv.org/abs/2005.05842>
- [13] J. Côté, J. K. Yardley, J. Hay, W. A. Sedgwick, and J. Baker, "An Exploratory Examination of the Coaching Behavior Scale for Sport," *AVANTE*, vol. 5, no. 3, pp. 82–92, 1999.
- [14] R. M. Ryan, "Control and information in the intrapersonal sphere: An extension of cognitive evaluation theory. - PsycNET," *Journal of Personality and Social Psychology*, vol. 43, no. 3, pp. 450–361, 1982.
- [15] S. G. Hart and L. E. Staveland, "Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research," in *Advances in Psychology*, vol. 52, P. A. Hancock and N. Meshkati, Eds., in Human Mental Workload, vol. 52. , North-Holland, 1988, pp. 139–183. doi: 10.1016/S0166-4115(08)62386-9.
- [16] A. Silva and M. Gombolay, "ARXIV version: Encoding Human Experts' Domain Knowledge to Warm Start Reinforcement Learning," in *AAAI*, 2021. Accessed: Feb. 11, 2021. [Online]. Available: <http://arxiv.org/abs/1902.06007>