REPAIR platform: Robot-aidEd PersonAlIzed Rehabilitation

Christian Tamantini, Alessandro Umbrico, and Andrea Orlandini

Institute of Cognitive Sciences and Technologies, National Research Council of Italy, 00196 Rome, Italy

Abstract. Physical rehabilitation is essential for restoring functionality and improving the quality of life for individuals affected by neurological or musculoskeletal conditions. Rehabilitation robots emerged as key-enabling technology to deliver intensive treatments and objectively quantify patients' motor performance. In the context of Healthcare 5.0, personalization of the treatment is paramount to improve the effectiveness of the interventions. Personalization can be implemented reactively, by providing real-time physical assistance and feedback, and deliberatively, by planning sessions based on therapeutic goals. Inspired by Kahneman's dual-system theory, this paper proposes a cognitive architecture for a robot-aided rehabilitation platform capable of delivering personalized treatment through deliberative and reactive techniques. The proposed cognitive architecture is described and validated through experimental sessions. Six healthy participants were enrolled in the experiments, simulating a robot-aided rehabilitation session with a TIAGo service robot serving as the physical interface to deliver the planned session. The results highlight that the plans generated according to different clinical objectives elicited distinct physiological responses from the participants, demonstrating the effectiveness of the personalized approach.

Keywords: Robot-aided rehabilitation · Physical rehabilitation · Automated Planning · Task Planning.

1 Introduction

Physical rehabilitation is a crucial component in the recovery process for individuals affected by different disabilities, ranging from neurological conditions to musculoskeletal injuries, aiming to restore lost functionality and improve the quality of life and independence. In 2019, over 2.4 billion people worldwide were living with health conditions that could benefit from rehabilitation, highlighting the immense need for physical treatments [1]. Conventional physical therapy involves a complex interaction between a physiotherapist and the patient, including physically manipulating the patient's body and developing an empathetic relationship [2]. Physiotherapists assess the patient's condition and create tailored rehabilitation exercise plans to push forward motor recovery and promote active participation [3].

In this scenario, rehabilitation robots have started to play an increasingly important role thanks to their capabilities to i) deliver highly intensive and repetitive treatments, ii) implement different care paradigms, iii) integrate feedback systems to gamify exercise and iv) objectively quantify patients' motor performance [4]. Among the opportunities that robotics provides in physical rehabilitation, personalization is one of the features that stands out the most as it enables the delivery of engaging treatments that promote active participation and improve motor recovery [5]. Personalization of robot-aided rehabilitation could be implemented at multiple levels: reactive and deliberative ones. The former generates reactive actions to support the execution of the individual task in real-time according to the patient's needs and state by providing tunable physical assistance or contextualized feedback [6]. The other acts at a higher level to recommend rehabilitation plans consistent with the motor recovery of the individual user [7]. In other words, robot-aided rehabilitation systems can be designed drawing inspiration from Kahneman's dual-system theory [8]. This allows for slow, deliberate planning of each rehabilitation session based on therapeutic goals and fast, real-time feedback and corrections during the execution of individual exercises. This dual approach ensures that each session is optimally tailored to the patient's progress and needs, enhancing the overall effectiveness of the rehabilitation program.

In this context, we propose a cognitive architecture for a robot-aided rehabilitation platform designed to personalize treatment by combining the two reasoning perspectives: deliberately, by generating a physiotherapy exercise plan suitable for the clinical goals of each session, by monitoring task execution and providing contextual feedback in real-time. This work introduces the main functional components of the architecture. A preliminary validation involving six healthy participants in a simulated robot-aided rehabilitation session with the TIAGo robot, demonstrating the feasibility of the approach.

The rest of the paper is structured as follows. Section 2 presents the scientific literature investigating the development of cognitive architectures to provide physical therapy. Section 3 introduces the proposed approach to personalize the rehabilitation treatment, its experimental implementation, and validation with healthy participants. Section 4 shows and discusses the results obtained in the experimental validation. Lastly, Section 5 summarises the main contributions and results and outlines future developments.

2 Related Works

Cognitively sophisticated architectures have been proposed in the literature in the healthcare context to provide robotic systems for rehabilitation with more complex comprehension and adaptation capabilities than traditional robot-mediated rehabilitation systems.

In [9,10], a control architecture for a social robot in rehabilitation is proposed, which handles both physical and cognitive interactions. Monitoring systems analyze users' movements, facial expressions, and spoken sentences to manage different behaviors following the "Stimulus-Response" approach [11]. The robot shifts roles from Demonstrator, the robot explains the motor task, to Observer, it monitors the patient's movements, and to Helper, it physically assists in task execution. This responsive platform effectively engages participants by integrating multimodal monitoring and natural language communication skills. However, it does not customize the treatment plan based on the patient's condition.

Automatic planning techniques have recently begun to contaminate robotic rehabilitation systems as they allow the generation of plans that consider the patient's condition at admission and generate an appropriate schedule of exercises to achieve certain goals within a session [12]. In particular, automatic planning methodologies were used to administer clinical scales, such as the Comprehensive Geriatric Assessment and the Quality of Upper Extremity Skills Test [13]. Both clinical scales require the patient to replicate a series of poses. Once the patient reaches the required pose, the robotic system, which includes the NAO robot and a Kinect camera for upper limb kinematics monitoring, calculates the deviation of the achieved pose from the desired one and automatically compiles the assessment sheet. A similar architecture has been applied to provide physical training to children with neurological disorders such as Cerebral Palsy or Obstetric brachial plexus palsy [14, 15]. This autonomous system plans rehabilitation sessions for children by offering two games: Mirror and Simon [16]. In these games, children mimic upper limb poses demonstrated by the robot. The system employs an automated planning process to generate pose sequences, thereby ensuring the inclusion of various poses within each session, and dynamically adjusting error thresholds for personalized treatment. The poorer the performance, the higher the tolerance. Although this architecture demonstrates its capabilities in generating more engaging and tailored sessions and leading to a notable motor recovery, this planning system has the only objective of planning a session of a certain duration structured as a sequence of warm-up, training, and cool-down phases. Recent research has concentrated on automated planning methodologies for rehabilitation sessions. These methodologies involve the design of choreographies to challenge participants in specific aspects, such as energy expenditure or balance [17]. This system uses the patient's motor condition to generate step sequences targeting the specific clinical objectives. However, this platform lacks a monitoring system to assess the impact of the generated plans on users.

Despite the development of sophisticated robotic rehabilitation systems, existing methodologies frequently fail to i) personalize treatment plans according to the specific circumstances of individual patients; ii) adapt exercise sequences in real-time, and; iii) comprehensively monitor and evaluate the impact of rehabilitation plans on users' progress. Therefore, there is a critical need for a control architecture in robotic rehabilitation that personalizes treatment based on patient-specific conditions and therapist-set clinical objectives. This system should integrate real-time monitoring and feedback to adapt the treatment plan, enhancing the effectiveness of interventions through tailored and responsive care.

3 Materials and Methods

3.1 Proposed Approach

We introduce the REPAIR platform capable of integrating deliberative planning and fast motion classification to support personalized assistance in physical rehabilitation. The architecture is organized according to the Dual Process theory and combines (slow) deliberative and (fast) reactive reasoning capabilities [7,18]. The deliberative layer is in charge of deciding the proper exercise set to administer to the patient and adapting planned and interacting behaviors according to clinicians' feedback and observed state and performance. The reactive layer is in charge of controlling robot actions as well as evaluating the execution of planned activities through the monitoring of patient's movements and physiological state. Figure 1 shows the organization of the functional components and the resulting control flow. It is worth noticing that REPAIR pursues a human-in-the-loop methodology [19] where feedback from clinicians is crucial to tailor reasoning and acting capabilities to specific clinical needs and objectives.

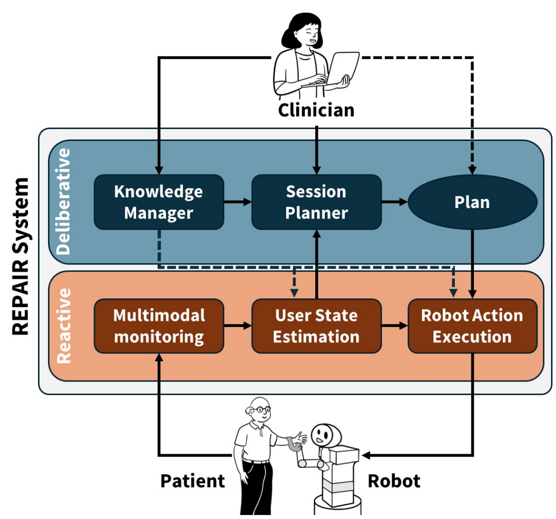

Fig. 1. Architecture of the proposed robot-aided rehabilitation platform.

The overall process begins with the clinician administering a set of clinical scales providing a clear overview of the patient's condition. In this regard, the clinician has the crucial role of endowing the system with the health-related knowledge necessary to make decisions that comply with the clinical practice. Namely, the clinician initializes a *Knowledge Base* encapsulated in the *Knowledge Manager* by providing domain-specific knowledge suitable to contextualize planning decisions according to different clinical objectives. For example, the clinician's input is crucial to model and characterize the effects of physical exercises on health-related conditions/features of patients [17].

Given the specific pathological condition and the patient's level of disability, the clinician specifies a clinical goal to reach in the current session and possibly additional constraints or preferences, e.g. the maximum session duration. Such a goal is fed into the automated planning module, also named *Session Planner*, that leverages the information stored in the *Knowledge Manager*, i.e. the list of exercises that can be used to generate the session along with their characteristics, to compute a proper *Plan*. Automated planning methodologies can be exploited inside the *Session Planner* to solve the planning problem of finding a personalized *Plan* for the patient under examination. In this context, a *Plan* is represented by a list of physical exercises (*stimuli*) that ensures the achievement of the specific clinical goal set by the clinician. Furthermore, the clinician can access the plan generated by the autonomous rehabilitation agent and modify it as appropriate, including the addition, removal, or modification of exercises, to ensure accurate supervision of the session (see dashed line in Fig. 1).

The execution of the rehabilitation session is managed at a low level by the reactive layer. It includes a module for user *Multimodal Monitoring* that collects raw information from several perspectives, ranging from kinematics to physiological measurements, useful to estimate user state. Indeed, the *User State Estimation* module takes as input the data collected from a set of sensors and estimates the complex state of the patient [20]. The user state represents, in an abstract symbolic manner, the quality of the movement performed by the user as well as information regarding the physical or cognitive spheres during the execution of a task. Such an estimation has a twofold effect: it triggers the *Robot Action Execution* in the reactive layer, such as the returning of verbal feedback or providing physical assistance, and it returns the information to the *Session Planner* to check whether re-planning is needed.

Lastly, it is worth noting that the *Knowledge Manager* is interconnected with both the estimation and action modules. This allows it to adapt the reactive modules according to its internal knowledge. The focus of the estimation process may change depending on the context, as well as the action module can generate different behaviors based on the robot's capabilities. For instance, a humanoid robot can physically mirror the tasks, a robot with an anthropomorphic arm can physically support task execution, and a digital system can provide vocal feedback.

3.2 Experimental Evaluation

To validate the proposed methodology, we implemented the approach described in Section 3.1 on a robotic platform to conduct robot-aided rehabilitation sessions. Moreover, a monitoring system was employed to track the user's state throughout the sessions. The following sections provide a detailed account of the materials used in the experimental setup and explain the implementation of each functional block of the architecture to achieve the desired behaviors.

Experimental Setup The service robot TIAGo (PAL Robotics S.L., Spain) was used as the physical robotic system. TIAGo features an anthropomorphic arm with 7 degrees of freedom (DoFs), a liftable torso, and a mobile base on wheels. Additionally, the robot is equipped with a microphone and speakers to manage audio input and output. Its head, which also has pan and tilt degrees of freedom, mounts an Asus Xtion RGB-D camera capable of providing RGB and depth images with a resolution of 640×480 at a frame rate of 30 Hz. Moreover, the participants were asked to wear a GARMIN Vivosmart 4 wristband to collect their Heart Rate (HR) [21]. Figure 2 shows the experimental setup used to evaluate the proposed methodology. All the software components run under the robot operating system (ROS Melodic).

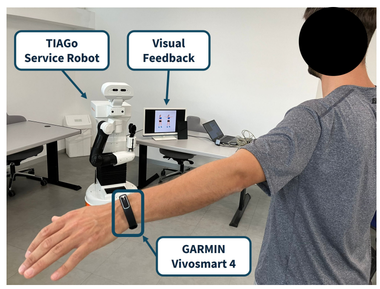

Fig. 2. Experimental setup used to evaluate the proposed methodology.

Knowledge Manager. The *Knowledge Manager* is the central data repository for patient profiles and exercise information. It maintains comprehensive profiles for each patient, including their clinical history, current condition, and progress. Additionally, it houses a detailed list of potential rehabilitation exercises, each described in terms of intensity levels. In the system tested in this paper, 23 exercises were extracted from the "PhysioTherapy eXercises" database [22], see Table 3.2. The selection of these exercises was based on their capacity to elicit a range of intensities, thereby engaging all major muscle groups. This repository is of critical importance for the *Session Planner* module, which accesses the *Knowledge Manager* to generate personalized and effective session plans.

Session Planner. The *Session Planner* is the high-level reasoning component that decides the sequence of physical stimuli suited for a rehabilitation session. The component relies on timeline-based planning [23, 24] and integrates search strategies capable of reasoning on the numeric effects of physical stimuli [17]. Specifically, we have extended the open-source planning framework PLATINUm[1] [25] with a heuristic search suitable to evaluate the clinical qualities of partial plans. Algorithm 1 briefly describes the structure of the search procedure. The

[1] `https://github.com/pstlab/PLATINUm.git`

Table 1. List of the exercises implemented in the Knowledge Base

ID	Exercise	Intensity	ID	Exercise	Intensity
A1	Arm circles	1	A13	Frontal lunges	3
A2	Side stretches	1	A14	Squat	3
A3	Side leg raises	1	A15	Military press	3
A4	Scarecrow arms rotation	1	A16	Side lunges	3
A5	Forward bend stretch	1	A17	Butt kicks	3
A6	Cross-body arm stretch	1	A18	Boxing	4
A7	Standing quad stretch	1	A19	High knees	4
A8	Bicep curl	1	A20	Jumping jacks	4
A9	Arm lateral raise	2	A21	High kick	4
A10	Arm front raise	2	A22	Jump squats	5
A11	Cross-body toe touches	3	A23	Running in place	5
A12	Body crunches	3			

planner receives as input the clinical objective encapsulated into a heuristic function $\mathcal{H}_\pi$ and the number of exercises to be administrated $\mathcal{N}$.

Algorithm 1 Heuristic search procedure of the Session Planner.

Input: $\mathcal{H}_\pi, \mathcal{N}$
Output: $\pi = (FTL, R)$
1: $\Pi' \leftarrow \emptyset$
2: $\pi \leftarrow$ initialize $(\mathcal{SV}, \mathcal{S})$
3: **while** $\neg$meetsRequirements $(\pi, \mathcal{N})$ **do**
4: $\quad \Pi' = \{\pi'_1, \pi'_2, ..., \pi'_m\} \leftarrow$ refine (π)
5: $\quad \pi \leftarrow$ select $(\Pi', \mathcal{H}_\pi)$
6: **end while**
7: **return** π

The termination condition represents a novel aspect of the implemented search procedure. Unlike "classical" planning problems requiring to achieve a certain state or decompose a certain task, the objective here is to synthesize a plan π which considers a sufficient number of stimuli. Plan refinement thus should always consider the possibility of recursively making additional planning decisions (i.e., subgoals) until plan requirements are met (rows 3-5). In the considered problem, the requirement conditions concern the number of stimuli $\mathcal{N}$ specified by the clinician. Alternatively, the clinician can specify the minimum duration of the session instead of the minimum number of exercises.

The planner should search for plans π that achieve a certain (clinical) objective within the specified requirements. In the current work, we consider two extremal clinical objectives: (i) LOW intensity $\mathcal{H}_\pi^{LOW}$; (ii) $HIGH$ intensity $\mathcal{H}_\pi^{HIGH}$. Equation 1 intuitively describes the objective function $f_i(\Pi)$ of the planning problem. It leverages data about exercise intensity in Table 3.2 to evaluate the cumulative intensity of the subset of selected exercises in a given plan π.

$$f_i(\pi) = \sum_{a_i \in \pi} \text{intensity}(a_i) \tag{1}$$

Two search strategies $\mathcal{H}_{\pi}^{LOW}$, $\mathcal{H}_{\pi}^{HIGH}$ have been developed to evaluate sequences of physical exercises $a_i \in \pi$ of a plan π that respectively minimize and maximize the cumulative intensity $f_i(\pi)$. Although we have considered the two simple objective functions mentioned above, the developed approach can support a wider set of more detailed/complex objectives provided by a therapist [17]. It is then worth underlining that our system can be extended to support many different clinical requirements and applications. It is worth noticing that the system has been designed with the explicit intention of assisting healthcare professionals, and not to supplant their role. This ensures that the expertise and judgment of therapists remain at the core of patient care.

Multimodal Monitoring. In this experiment, the multimodal monitoring system collects data related to two dimensions: movement and cardiac activity. The TIAGo robot's built-in camera tracks the user's kinematics during the task execution. Specifically, once a frame is captured by the RGB camera, the Mediapipe pose algorithm is employed to retrieve the user's anatomical landmarks [26]. The three-dimensional joint coordinates in the real world, expressed in the user's origin frame, which is situated between the hips, are collected at a frequency of 30 Hz. In particular, the coordinates of the shoulders, elbows, wrists, hips, knees, and ankles were considered, as the objective was to monitor total body motions. Furthermore, the user's heart rate is collected throughout the session from the wrist-worn device at a frequency of 1 Hz.

User State Estimation. The user state estimation module is responsible for recognising, classifying and counting the user's movements. This software module identifies which of the known activities the user is performing and tracks the duration of each action.

To train the action classification algorithm, data were collected from four healthy participants (26.2 ± 4.1 mean age, 4 males). For each of the 23 exercises, data were gathered over 20 seconds at a sampling rate of 30 Hz, resulting in 600 observation per exercise. Additionally, a further class was included (A0), representing the resting condition, with data collected under identical conditions. Each observation comprised the three-dimensional coordinates of anatomical landmarks of the upper and lower limbs, monitored over 30 frames (equivalent to one second). The supervised model employed in this study to perform action recognition was a Support Vector Machine (SVM) with a radial basis function kernel. The choice of the SVM classifier was driven by its effectiveness in handling high-dimensional data, robustness to overfitting, and its suitability to perform real-time inferences during robot-aided rehabilitation sessions [27].

Robot Action Execution. The robotic tasks considered in this scenario are: i) displaying the currently administrated exercise, and; ii) providing visual and vocal feedback on the task execution. The monitor of the robotic system provides visual feedback, displaying the physical motion to be performed. Additionally, the robot may deliver verbal feedback to guide the user through the exercises. This dual feedback mechanism ensures clear communication and helps maintain user engagement and correct execution of the rehabilitation tasks. The feedback phase allows for the real-time correction and reinforcement of the user's movements. The administration of visual and verbal feedback facilitates the user's comprehension of and ability to modify their task execution, thereby ensuring continuous engagement and immediate guidance [10].

Experimental Protocol In this experiment, 6 healthy right-handed participants (30.8 ± 5.3 mean age, 5 males and 1 female) were enrolled. They provided written consent to participate in the study. To test the capability of the system to deliver different rehabilitation sessions, each participant was asked to perform two sessions, specifically planned with the objectives of *LOW* and *HIGH* intensities and the same requirement in terms of duration, i.e. the session is required to be composed of 10 exercises.

Performance Indicators The offline performance of the implemented action classification model was initially evaluated. Since data from four participants were recorded, a leave-one-subject-out (LOSO) cross-validation approach was employed. This method entails training the model on data from three participants and testing it on the data from the remaining one. This process is repeated for each enrolled participant, thereby ensuring a robust evaluation of the model's performance. Given that the dataset was balanced, accuracy was selected as the metric to assess the classification performance. Moreover, the time needed to train (T_{train}^{SVM}) the model as well as the time to perform inference ($T_{predict}^{SVM}$) were computed.

To quantify the efficacy of the *Session Planner* in generating tailored plans and the impact of the session on the enrolled subjects, the following performance indicators were computed:

- Cumulative Plan Intensity (ΣPI): This indicator quantifies the overall intensity of all exercises incorporated into a rehabilitation session. The indicator is calculated by summing the individual intensities of each exercise and reflects the planner's ability to generate plans with different intensities according to specific input intensity requirements.
- Session Duration (T_{tot}): total time taken to complete a rehabilitation session is defined as the entire period from the beginning to the end of the session, including any intervals or breaks, and is expressed in minutes. The monitoring of T_{tot} enables the assessment of whether plans generated at different intensity levels require different amounts of time.

– Normalized Hear Rate (HR_n): the user's HR mean response, calculated as

$$HR_n = mean\left(\frac{HR - HR_{base}}{HR_{base}}\right) \tag{2}$$

where HR is the actual heart rate during the session and HR_{base} is the baseline heart rate at the beginning of the session. The value of HR_n provides information on the physiological impact of the planned session on the user.

The Mann-Whitney test was employed to ascertain whether there were any statistically significant differences between the computed metrics when *LOW* and *HIGH* intensity plans were administered to the participants. Furthermore, Pearson's linear correlation test was used to compute the relationship between the ΣPI and HR_n. The significance level was set for both statistical tests at a p-value of 0.05.

4 Results and Discussion

Fig. 3A reports the normalized confusion matrix obtained by appending all the predictions inferred on the testing subjects during the LOSO validation. The mean accuracy was found to be $81.80 \pm 4.11\%$. These results demonstrate the model's capacity to maintain accuracy across diverse individuals, which is a crucial attribute for models that handle user characteristic variability while ensuring consistent performance levels. The training time for the SVM model was $T^{SVM}_{train} = 597.25 \pm 50.16$ s, and the inference time per observation was $T^{SVM}_{predict} = 0.07 \pm 0.03$ s. These results suggest that while the SVM model requires a reasonable training period, its quick inference time per observation makes it suitable for applications requiring rapid decision-making and real-time processing of data.

Fig. 3B illustrates the frequency of occurrence of the various exercises in the planes generated at *LOW* and *HIGH* intensities. The two conditions are represented by different colours: green and red for *LOW* and *HIGH*, respectively. The graph illustrates how the rehabilitation plans are adapted according to the required intensity level, employing a combination of exercises to optimize the effectiveness of the session. The *LOW*-intensity plans comprise a more diverse range of exercises, distributed throughout the session in a more balanced manner. In contrast, the *HIGH*-intensity plans focus on a smaller set of exercises, including squats, frontal lunges, and high knees, which are selected by the planner with greater frequency. This distinction reflects the adaptation of the plans to enhance the efficacy of the workout, with a greater emphasis on variety and balance in the low-intensity plans and a more concentrated approach to high-impact exercises in the high-intensity plans.

Fig. 4 shows the performance indicators computed during the experiments carried out in the validation of the proposed planning system. As expected, ΣPI values reflect the clear statistically significant difference between the LOW and HIGH conditions (p-value$< 1 \cdot 10^{-2}$). The HIGH condition requires a greater

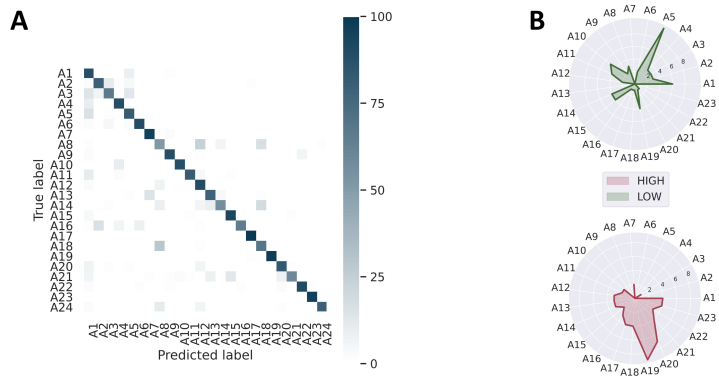

Fig. 3. A. Normalized confusion matrix of the trained SVM model for action classification. B. Occurrences of the exercises in the plans for LOW and $HIGH$ intensity.

ΣPI, consistent with generating more intense sessions and therefore tailoring the session according to the specific clinical goal.

Although the average session duration is slightly longer in the HIGH condition, there is some overlap between the two groups and the distributions are not statistically different (p-value= 0.7). This suggests that HIGH sessions, while more intense, do not always take significantly longer than LOW sessions. Variations in sessions duration may be influenced by several factors, such as individual responses to tasks and variability in exercise performance.

The HR_n metric demonstrates a statistically significant difference between the LOW and HIGH conditions (p-value$< 1 \cdot 10^{-2}$). The mean HR_n values in the HIGH condition are significantly higher, suggesting that the more intense the plan is, the higher the physiological response. This aligns with the expectation that more intensive sessions (HIGH) cause a greater increase in heart rate compared to less intensive sessions (LOW). Negative HR_n values in some LOW sessions indicate minimal or negative physiological response, i.e. a lower HR with respect to the baseline condition, likely due to the relatively low intensity of the exercises. Such a result is also stressed by Pearson's linear correlation between ΣPI and HR_n demonstrating a statistically significant strong linear relationship $\rho = 0.68$ (p-value= 0.01).

5 Conclusions

This work presents a cognitive architecture for a robot-aided rehabilitation platform designed to provide personalized treatment by integrating deliberate exercise planning and real-time reactive feedback. The proposed system was validated with six healthy participants using the TIAGo robot to simulate rehabilitation sessions. Despite the relatively modest cohort size of enrolled subjects, the findings indicate that the planner can effectively produce different sessions according

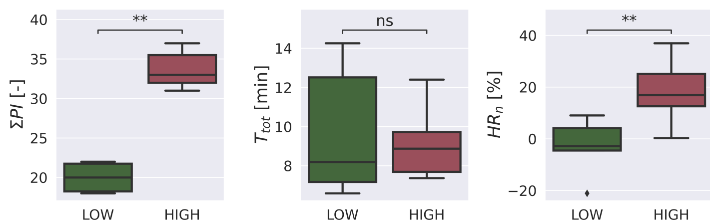

Fig. 4. Performance indicators computed during the experimental validation, categorized by the LOW and HIGH experimental conditions. *ns* indicates no statistically significant difference, and $**$ denotes a significant difference (p-value < 0.01).

to the input intensity level. Different exercises were selected and the ΣPI reflects the intensity. Moreover, the multimodal monitoring system revealed that *HIGH* intensity sessions significantly impacted the participants' HR, reflecting a greater physiological response. Indeed, a statistically significant strong linear correlation was found between the HR_n and ΣPI. The results demonstrate that the planned sessions elicited disparate effects on the body's response.

The preliminary findings demonstrate the platform's capacity to adapt exercise intensity and impact patient response, thereby generating personalized rehabilitation sessions. The long-term objective is to circumvent repetitive exercises and target specific body areas for varied and efficacious rehabilitation. The system can be expanded to support a greater number of therapeutic programs and integrate additional sensors to endow the REPAIR platform with the capability to estimate psychophysiological processes and enhance the *Session Planner* with comprehensive user state estimates for more sophisticated rehabilitation plans. Moreover, from a clinical point of view, extensive experiments will be carried out enrolling the pathological population to assess the applicability of such a system in real clinical practice.

Acknowledgments. This work was supported by the Italian Ministry of Research, under the complementary actions to the NRRP "Fit4MedRob - Fit for Medical Robotics" Grant PNC0000007, (CUP: B53C22006990001).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Alarcos Cieza, Kate Causey, Kaloyan Kamenov, Sarah Wulf Hanson, Somnath Chatterji, and Theo Vos. Global estimates of the need for rehabilitation based on the global burden of disease study 2019: a systematic analysis for the global burden of disease study 2019. *The Lancet*, 396(10267):2006–2017, 2020.

2. Michelle J Kleiner, Elizabeth Anne Kinsella, Maxi Miciak, Gail Teachman, Erin McCabe, and David M Walton. An integrative review of the qualities of a 'good'physiotherapist. *Physiotherapy theory and practice*, 39(1):89–116, 2023.
3. John Naylor, Clare Killingback, and Angela Green. What are the views of musculoskeletal physiotherapists and patients on person-centred practice? a systematic review of qualitative studies. *Disability and rehabilitation*, 45(6):950–961, 2023.
4. Christophe Duret, Anne-Gaëlle Grosmaire, and Hermano Igo Krebs. Robot-assisted therapy in upper extremity hemiparesis: overview of an evidence-based approach. *Frontiers in neurology*, 10:438977, 2019.
5. Helen Rodgers, Helen Bosomworth, Hermano I Krebs, Frederike van Wijck, Denise Howel, Nina Wilson, Lydia Aird, Natasha Alvarado, Sreeman Andole, David L Cohen, et al. Robot assisted training for the upper limb after stroke (ratuls): a multicentre randomised controlled trial. *The Lancet*, 394(10192):51–62, 2019.
6. Christian Tamantini, Francesca Cordella, Clemente Lauretti, Francesco Scotto Di Luzio, Benedetta Campagnola, Laura Cricenti, Marco Bravi, Federica Bressi, Francesco Draicchio, Silvia Sterzi, et al. Tailoring upper-limb robot-aided orthopedic rehabilitation on patients' psychophysiological state. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2023.
7. Alessandro Umbrico, Riccardo De Benedictis, Francesca Fracasso, Amedeo Cesta, Andrea Orlandini, and Gabriella Cortellessa. A mind-inspired architecture for adaptive hri. *International Journal of Social Robotics*, 15(3):371–391, 2023.
8. Jean-François Bonnefon and Iyad Rahwan. Machine thinking, fast and slow. *Trends in Cognitive Sciences*, 24(12):1019–1027, 2020.
9. Lorenzo Cristofori, Claudiu D. Hromei, Francesco Scotto Di Luzio, Christian Tamantini, Francesca Cordella, Danilo Croce, Loredana Zollo, and Roberto Basili. Heal9000: An intelligent rehabilitation robot. In *CEUR Workshop Proceedings*, volume 3060, page 29 – 41, 2021.
10. Christian Tamantini, Francesco Scotto di Luzio, Claudiu Daniel Hromei, Lorenzo Cristofori, Danilo Croce, Marco Cammisa, Arcangela Cristofaro, Maria Vittoria Marabello, Roberto Basili, and Loredana Zollo. Integrating physical and cognitive interaction capabilities in a robot-aided rehabilitation platform. *IEEE Systems Journal*, 2023.
11. Michelle J Johnson, Mayumi Mohan, and Rochelle Mendonca. Therapist-patient interactions in task-oriented stroke therapy can guide robot-patient interactions. *International Journal of Social Robotics*, 14(6):1527–1546, 2022.
12. Raquel Fuentetaja, Angel García-Olaya, Javier García, José Carlos González, and Fernando Fernández. An automated planning model for hri: Use cases on social assistive robotics. *Sensors*, 20(22):6520, 2020.
13. Alejandro Martín, José C Pulido, José C González, Ángel García-Olaya, and Cristina Suárez. A framework for user adaptation and profiling for social robotics in rehabilitation. *Sensors*, 20(17):4792, 2020.
14. José Carlos González, José Carlos Pulido, and Fernando Fernández. A three-layer planning architecture for the autonomous control of rehabilitation therapies based on social robots. *Cognitive Systems Research*, 43:232–249, 2017.
15. Jose Carlos Pulido, Cristina Suarez-Mejias, Jose Carlos Gonzalez, Alvaro Duenas Ruiz, Patricia Ferrand Ferri, Maria Encarnacion Martinez Sahuquillo, Carmen Echevarria Ruiz De Vargas, Pedro Infante-Cossio, Carlos Luis Parra Calderon, and Fernando Fernandez. A socially assistive robotic platform for upper-limb rehabilitation: a longitudinal study with pediatric patients. *IEEE Robotics & Automation Magazine*, 26(2):24–39, 2019.

16. Mısra Turp, José Carlos González, José Carlos Pulido, and Fernando Fernández. Developing a robot-guided interactive simon game for physical and cognitive training. *International Journal of Humanoid Robotics*, 16(01):1950003, 2019.
17. Alessandro Umbrico, Marco Benadduci, Roberta Bevilacqua, Amedeo Cesta, Francesca Fracasso, Elvira Maranesi, Andrea Orlandini, and Gabriella Cortellessa. Combining clinical and spatial constraints into temporal planning to personalize physical rehabilitation. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 33, 2023.
18. Grady Booch, Francesco Fabiano, Lior Horesh, Kiran Kate, Jonathan Lenchner, Nick Linck, Andreas Loreggia, Keerthiram Murgesan, Nicholas Mattei, Francesca Rossi, and Biplav Srivastava. Thinking fast and slow in ai. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17), 2021.
19. Alessandra Sorrentino, Laura Fiorini, Gianmaria Mancioppi, Filippo Cavallo, Alessandro Umbrico, Amedeo Cesta, and Andrea Orlandini. Personalizing care through robotic assistance and clinical supervision. *Frontiers in Robotics and AI*, 9, 2022.
20. Christian Tamantini, Francesca Cordella, Nevio Luigi Tagliamonte, Ilenia Pecoraro, Iolanda Pisotta, Alessandra Bigioni, Federica Tamburella, Matteo Lorusso, Marco Molinari, and Loredana Zollo. A data-driven fuzzy logic method for psychophysiological assessment: An application to exoskeleton-assisted walking. *IEEE Transactions on Medical Robotics and Bionics*, 2024.
21. An-Marie Schyvens, Nina Catharina Van Oost, Jean-Marie Aerts, Federica Masci, Brent Peters, An Neven, Hélène Dirix, Geert Wets, Veerle Ross, Johan Verbraecken, et al. Accuracy of fitbit charge 4, garmin vivosmart 4, and whoop versus polysomnography: Systematic review. *JMIR mHealth and uHealth*, 12(1):e52192, 2024.
22. J Glinsky, L Harvey, C Sherrington, and O Katalinic. www. physiotherapyexercises. com–new exercises and features to help physiotherapists prescribe home exercise programs. *Physiotherapy*, 101:e1381, 2015.
23. N. Muscettola. HSTS: Integrating Planning and Scheduling. In Zweben, M. and Fox, M.S., editor, *Intelligent Scheduling*. Morgan Kauffmann, 1994.
24. Marta Cialdea Mayer, Andrea Orlandini, and Alessandro Umbrico. Planning and execution with flexible timelines: a formal account. *Acta Informatica*, 53(6):649–680, 2016.
25. Alessandro Umbrico, Amedeo Cesta, Marta Cialdea Mayer, and Andrea Orlandini. Platinum: A new framework for planning and acting. In Floriana Esposito, Roberto Basili, Stefano Ferilli, and Francesca A. Lisi, editors, *AI*IA 2017 Advances in Artificial Intelligence*, pages 498–512, Cham, 2017. Springer International Publishing.
26. Sebastian Dill, Andreas Rösch, Maurice Rohr, Gökhan Güney, Luisa De Witte, Elias Schwartz, and Christoph Hoog Antink. Accuracy evaluation of 3d pose estimation with mediapipe pose for physical exercises. In *Current Directions in Biomedical Engineering*, volume 9, pages 563–566. De Gruyter, 2023.
27. Tien-Thanh Nguyen, Dinh-Tan Pham, Hai Vu, and Thi-Lan Le. A robust and efficient method for skeleton-based human action recognition and its application for cross-dataset evaluation. *IET Computer Vision*, 16(8):709–726, 2022.