

Behavioral Cloning of Physiotherapists in Adapting Robot Control Parameter

Rita Molle¹, Christian Tamantini¹, *Member, IEEE*, Clemente Lauretti¹, Davide Sebastiani,
Fabio Santacaterina¹, Marco Bravi, Federica Bressi, Sandra Miccinilli,
and Loredana Zollo¹, *Senior Member, IEEE*

Abstract—Robot-aided rehabilitation effectively supports treatment of upper-limb disorders and enhances outcomes when combined with traditional therapy. Artificial intelligence enables behavioral cloning of physiotherapists' expertise to autonomously modulate robot assistance from real-time multimodal patient data. Therefore, this paper aims to propose and validate a behavioral cloning strategy, namely Physiotherapist-Supervised Parameter Adaptation (PSPA), for online tuning the robot assistance level replicating the physiotherapists' decision-making. The experimental validation was conducted in a clinical setting involving ten post-surgical orthopedic patients who participated in a robot-aided rehabilitation session using the KUKA LWR 4 + robot. The sessions were supervised by physiotherapists

who could adjust the level of robotic assistance as needed, thus labelling the collected patient multimodal data. The validation aimed at *i*) identifying the best-performing input modality, feature set, and classifier, and *ii*) comparing the capability of the approach in tailoring the assistance level with respect to the established performance-based (PB) one. Combining biomechanical and physiological features significantly improved the classification performance across all classifiers, with the highest performance observed for the Multilayer Perceptron on the present dataset. Moreover, using the optimized feature set, the proposed PSPA methodology achieved an even greater alignment with the physiotherapists' decisions with respect to the PB approach ($\Delta F1\text{-score} = 15.40 \pm 30.33\%$, $\rho = 0.56 \pm 0.21$ for PSPA, $\rho = -0.12 \pm 0.43$ for PB).

Index Terms—Supervised learning, behavioral cloning, interaction control, robot-aided rehabilitation, upper-limb, musculoskeletal disorders.

Received 16 September 2025; revised 4 December 2025, 15 January 2026, and 26 January 2026; accepted 26 January 2026. Date of publication 29 January 2026; date of current version 11 February 2026. This work was supported in part by the Italian Ministry of Research, under the complementary actions to the NRRP "Fit4MedRob - Fit for Medical Robotics" under Grant PNC0000007 (CUP: B53C22006990001); in part by the National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.1, Call for tender No. 104 published on 2.2.2022 by the Italian Ministry of University and Research (MUR), funded by the European Union - NextGenerationEU-Project Title 2022P349BK - "PERSEA - Personal Empathic Robot with Sensory-motor and social interaction capabilities for Autis" - CUP C53D23000470008; and in part by Fondazione Compagnia di San Paolo within the project "Bio-inspired stimulation techniques for restoring multimodal somatic sensations in patients with sensorimotor limitations - BIOMIM", funded by the Bando vEiColo - Accompagnamento per la valorizzazione della ricerca. (Rita Molle and Christian Tamantini contributed equally to this work.) (Corresponding author: Rita Molle.)

Rita Molle, Clemente Lauretti, and Davide Sebastiani are with the Research Unit of Advanced Robotics and Human-Centred Technologies, Università Campus Bio-Medico di Roma, 00128 Rome, Italy (e-mail: rita.molle@unicampus.it).

Christian Tamantini is with the Research Unit of Advanced Robotics and Human-Centred Technologies, Università Campus Bio-Medico di Roma, 00128 Rome, Italy, and also with the Institute of Cognitive Sciences and Technologies, National Research Council of Italy, 00196 Rome, Italy.

Fabio Santacaterina and Federica Bressi are with the Research Unit of Rehabilitation, Fondazione Policlinico Universitario Campus Bio-Medico di Roma, 00128 Rome, Italy, and also with the Unit of Physical Medicine and Rehabilitation, Università Campus Bio-Medico di Roma, 00128 Rome, Italy.

Marco Bravi and Sandra Miccinilli are with the Research Unit of Rehabilitation, Fondazione Policlinico Universitario Campus Bio-Medico di Roma, 00128 Rome, Italy.

Loredana Zollo is with the Research Unit of Advanced Robotics and Human-Centred Technologies, Università Campus Bio-Medico di Roma, 00128 Rome, Italy, and also with the Fondazione Policlinico Universitario Campus Bio-Medico di Roma, 00128 Rome, Italy.

Digital Object Identifier 10.1109/TNSRE.2026.3659215

I. INTRODUCTION

ROBOT-AIDED rehabilitation has become a widely adopted approach for supporting motor recovery in patients with neuromuscular or musculoskeletal conditions, offering repetitive and targeted assistance while providing quantitative feedback on patient performance [1], [2], [3], [4], [5]. Despite these advantages, achieving fully personalized rehabilitation remains a significant challenge. A core difficulty lies in continuously adapting the Assistance Level (AL) to match the patient's evolving capabilities, fatigue, and motor control strategies. Improper adaptation can result in suboptimal recovery, over-assistance, or under-challenging exercises, which limits the overall effectiveness of therapy [6].

Physiotherapists play a central role in addressing this challenge. Through real-time observation and clinical expertise, they are able to modulate parameters such as stiffness, force assistance, and exercise intensity to align therapy with each patient's specific needs [7], [8]. This human judgment provides personalized adjustments and expert knowledge that automated systems alone cannot replicate. By interpreting subtle cues, such as compensatory movements or signs of discomfort, therapists ensure that rehabilitation is both safe and effective [9].

While manual control by physiotherapists allows high clinical supervision and patient-centered interventions [10], it also presents intrinsic operational challenges. Real-time adaptation of robot assistance to every movement can be difficult [11].

Simultaneously managing multiple control parameters is cognitively demanding for the therapist, and ensuring continuous and gradual tuning of the *AL* throughout the session may be challenging [12]. Moreover, assistance adjustments may vary across sessions, reflecting differences in patient condition and therapeutic context. Together, these aspects can limit the scalability, continuity, and replicability of manual interventions, particularly in high-intensity or long-term rehabilitation programs. Nevertheless, recent studies have shown that physiotherapists' assistance-level decisions are guided by consistent relationships between observable biomechanical and physiological indicators and the perceived patient needs, suggesting the presence of underlying decision-making patterns that can be learned from data [13].

Building upon these findings, recent robotic rehabilitation systems have started incorporating multimodal data to guide assistance and decision-making strategies. By combining biomechanical information, such as kinematics and dynamics [14], [15], with physiological signals including Galvanic Skin Response (GSR), Heart Rate (HR), and Respiration Rate (RR), therapists can monitor the users' status and robots can more accurately and responsively modulate *ALs* [16], [17]. These multimodal strategies aim to improve the robot ability to estimate patient needs and provide more responsive assistance [18]. However, even with these enhancements, current methods remain limited in capturing the full complexity of patient-specific adaptation and the personalized decision-making typically provided by expert therapists.

For instance, automatic adaptation methods, such as performance-based (PB) impedance control, offer consistency, repeatability, and objectivity [19], [20], [21]. Despite these strengths, they are limited in their ability to provide personalized and clinically expert-based care. Rule-based controllers rely on predefined parameters and often struggle to adapt flexibly to the variability of individual patients [22], [23]. Thus, purely automatic systems may deliver suboptimal assistance that does not fully respond to the patient's moment-to-moment needs.

Recent research has explored therapist-inspired approaches to address this gap. For instance, stiffness modulation based on electromyographic signals has been proposed to estimate patient stiffness and adapt impedance control [24]. Other works have implemented strategies to directly mimic therapist-guided stiffness adjustments in rehabilitation exoskeletons [25]. These studies highlight the potential of transferring clinical expertise into robotic control and often incorporate therapist-guided modulation strategies. However, such approaches typically rely on offline tuning, predefined mappings, or state estimation mechanisms, and do not explicitly model the therapist's real-time, sequential decision-making process within the rehabilitation loop.

Building upon these insights, and to promote a patient-centred approach, it is important to use the insight and expertise of physiotherapists to fine-tune the *AL* and ensure that care is tailored to individual needs [7], [8]. Therefore, based on the current literature and the observation on conventional rehabilitation treatments, it is expected that integrating the physiotherapist's expert judgment with technological tools

through imitation learning, and more specifically behavioral cloning (BC) that transfers expert strategies into robotic assistance [26], will ensure that the robotic therapy is responsive to the patient's needs.

Thus, the objective of this paper is to overcome the limitations of existing automatic methods for adjusting robot *ALs* by proposing a Physiotherapist-Supervised Parameter Adaptation (PSPA) strategy. Framed within imitation learning, and specifically BC, PSPA aims to capture how physiotherapists adapt robot assistance in response to the observed patient state and to transfer these decision-making patterns into the robotic control system [26], [27]. Rather than quantifying inter-therapist agreement or modeling therapist-specific behaviors, the proposed approach focuses on learning transferable mappings between multimodal patient observations and assistance-level adaptations.

In this way, the robot learns to adapt the stiffness control parameter in real time during therapy sessions, ensuring that assistance remains responsive to the patient's needs. By combining biomechanical and physiological measurements collected during treatment, the PSPA strategy enables the development of an autonomous system that delivers tailored rehabilitation by embedding physiotherapists' clinical expertise into the control loop, without relying on predefined rules or expert-designed heuristics. In this sense, the proposed methodology can be framed as a BC approach, as the system learns to imitate expert behavior by mapping the monitored patient state to the corresponding therapist decision [27].

To achieve this goal, the proposed approach relies on collecting in the real clinical setting the decisions made by professionals in modifying the robotic *AL* according to the monitored patients' motor performance and physiological activity. This process generates a labeled dataset, created with input from therapists actively participating in the robot-aided rehabilitation loop, which constitutes the basis for a BC approach, where the model learns to reproduce the therapists' decision-making strategy. Ten orthopedic patients and their therapists participated in the data collection. Several supervised artificial intelligence (AI) models were trained and evaluated to determine their ability to reproduce therapist-driven *AL* modifications. Model predictions were compared with the therapist's actual decisions and with a traditional PB approach to assess the added value of incorporating therapist expertise.

The main contributions of this work can be summarized as follows:

- **Behavioral cloning of expert clinical decision-making in robot-aided rehabilitation:** We propose a novel strategy that formulates assistance-level adaptation as a behavioral cloning problem, enabling the robot to learn clinically relevant decisions directly from expert physiotherapists. The approach leverages a multimodal representation of the patient state, combining biomechanical and physiological measurements, and is validated in a real clinical scenario involving ten orthopedic patients and their physiotherapists.
- **Subject-independent validation against clinically established rule-based methods:** We provide a




Fig. 1. Block scheme of the proposed PSPA approach for adapting robot-aided rehabilitation.

validation of the proposed learning-based approach through a subject-independent evaluation, comparing PSPA predictions against both physiotherapists' ground-truth decisions and a classical PB, rule-driven adaptation strategy adopted in robot-aided rehabilitation. This comparison quantitatively assesses the added value of embedding expert clinical knowledge into robotic rehabilitation platforms.

The paper is organized as follows: Section II describes the proposed approach, the experimental setup, and the experimental validation in the clinical scenario. Section III outlines and discusses the main results, while Section IV provides the conclusions and delineates future works.

II. MATERIALS AND METHODS

A. The Proposed Approach

Fig. 1 illustrates the proposed PSPA framework for robot-aided rehabilitation. The figure provides an overview of the system components, their interactions, and the general workflow followed to adapt the robot assistance in a BC manner. In the specific implementation considered in this study, the proposed strategy is applied using a KUKA LightWeight Robot 4 + robotic manipulator, and the rehabilitation task consists of point-to-point movements in a three-dimensional workspace. Further details about the setup and the protocol will be given in Section II-B and II-C.

The first module encountered in the block diagram is the *trajectory planner*, which generates a Cartesian path $\mathbf{X}$ along with a current desired pose $\mathbf{x}_w(t) = [\mathbf{p}_w(t); \Phi_w(t)]$, where $\mathbf{p}_w(t)$ and $\phi_w(t)$ denote the Cartesian position and orientation, respectively. Both these quantities are fed into the *interaction control* module, in which an inverse dynamic control law is employed to compute the control torques τ_c for robot joints. At the *human-robot interaction* level, kinetic and kinematic data are recorded to analyze the user's and robot movement and interaction. Additionally, wearable physiological sensors placed on the skin, as described in II-B, capture data on heart and respiratory activities, and skin conductance to provide insights into the user's physiological




Fig. 2. Control actions generated by the tunable interaction control. The radial and tangential actions $\mathbf{a}_T$ and $\mathbf{a}_W$ are outlined.




Fig. 3. Experimental setup used to validate the adapting robot-aided rehabilitation approach based on PSPA. The GUI shown to physiotherapists (left) allows entering patient information (ID, sex, age), displays the current AL, and allows modifying it. Additionally, key quantitative parameters such as the mean of heart rate (HR), respiration rate (RR), and interaction force over each repetition are presented. These parameters, together with the physiotherapist's direct observation and knowledge of the patient, guide adjustments of the AL during the session.

responses during the interaction. These measurements are integrated in the *multimodal monitoring* module. Specifically, the raw signals acquired through this system represent the input of the *feature extraction* module, where physiological and biomechanical features are extracted starting from the recorded data. The physiotherapist is provided with a custom Graphical User Interface (GUI) to monitor the therapy by observing the collected parameters coming from the *feature extraction* module. Moreover, the GUI includes three buttons corresponding to different AL, allowing the physiotherapist to select the desired level by pressing the appropriate button (see Fig. 3). Specifically, the PSPA module acts autonomously or with the physiotherapists' intervention, as the physiotherapist has a higher priority than the PSPA module in the proposed approach and can choose when to intervene to modify the AL, allowing the module to operate safely and progressively refine its decisions. The proposed approach is implemented through the acquisition of labeled data that captures the physiotherapist's sequential adjustments during the sessions. These labels represent the expert demonstrations that define the therapist's decision policy, and therefore constitute the teaching phase of a discrete action BC approach: the system learns the mapping between patient states and the therapist's selected discrete action, enabling the reproduction of appropriate AL adjustments. stiffness to match the patient's needs as inferred from physiotherapist-validated labels.

Further details about each functional block are provided in the subsequent explanation.

1) Trajectory Planner: The trajectory planner produces Cartesian trajectories given different initial and final poses ($\mathbf{x}_i$ and $\mathbf{x}_f$, respectively). The trajectory planner calculates the path ($\mathbf{X}$) and the current desired pose ($\mathbf{x}_w(t)$) for each time instant (t) that patients need to follow. It takes into account the initial pose ($\mathbf{x}_i$), the final target pose ($\mathbf{x}_f$), a pause time (T_p) allowing patients to initiate the exercise autonomously without robot intervention, and the time required to complete the task (T_f). The back-wall pose ($\mathbf{x}_w(t)$) remains constant at $\mathbf{x}_i$ for a duration of T_p seconds to wait for patient movement initiation. However, if a patient takes more than T_p seconds to start moving, the desired pose ($\mathbf{x}_w(t)$) moves forward toward the user's limb to complete the assigned task in T_f seconds. Further details about the motion planner used in this approach can be found in [28].

2) Interaction Control: The control strategy is a modified version of the tunable inverse dynamics control law to implement an assist-as-needed approach, as presented in [28], in which the control stiffness is modified solely according to the kinematic performance of the patients. Different from inverse-dynamics formulations that enforce an explicit Cartesian time law, the proposed controller is designed to generate a position-based force field without prescribing desired Cartesian velocity or acceleration profiles. This choice allows the patient to freely explore the trajectory with self-selected timing, while the robot provides spatial guidance and assistance modulated through stiffness parameters selected by the physiotherapist.

The control law is expressed as follows

$$\begin{cases} \tau_c = \mathbf{B}(\mathbf{q})\mathbf{y} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{F}_v\dot{\mathbf{q}} + \mathbf{F}_s \text{sign}(\dot{\mathbf{q}}) + \mathbf{g}(\mathbf{q}) \\ \mathbf{y} = \mathbf{J}_A^\dagger(\mathbf{q})(\ddot{\mathbf{x}}_d + \mathbf{A}_t^{-1}\mathbf{K}_T\mathbf{A}_t\dot{\mathbf{x}}_T + \mathbf{A}_t^{-1}\mathbf{K}_W\mathbf{A}_t\dot{\mathbf{x}}_w + \\ -\dot{\mathbf{J}}_A\dot{\mathbf{q}} + \mathbf{A}_t^{-1}\mathbf{K}_D\mathbf{A}_t\dot{\mathbf{x}}) \end{cases} \quad (1)$$

Here, $\mathbf{B}(\mathbf{q})$ is the robot inertia matrix, $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})$ represents centrifugal and Coriolis effects, $\mathbf{F}_v$ is the viscous friction torque, $\mathbf{F}_s \text{sign}(\dot{\mathbf{q}})$ is the static friction torque, and $\mathbf{g}(\mathbf{q})$ is the gravity contribution. Additionally, $\mathbf{y}$ includes joint space accelerations with stabilizing actions, with the right pseudo-inverse matrix of the robot Jacobian $\mathbf{J}_A^\dagger(\mathbf{q})$, the derivative of the robot Jacobian $\dot{\mathbf{J}}_A$, the adjoint matrix $\mathbf{A}_t$, the stiffness matrices $\mathbf{K}_T$, $\mathbf{K}_W$, the damping matrix $\mathbf{K}_D$, end-effector velocity error $\dot{\mathbf{x}} = \dot{\mathbf{x}}_d - \dot{\mathbf{x}}$ (with $\dot{\mathbf{x}}_d$ and $\dot{\mathbf{x}}$ being the desired and actual velocity, respectively), and end-effector desired acceleration $\ddot{\mathbf{x}}_d$. The pose errors $\tilde{\mathbf{x}}_T$ and $\tilde{\mathbf{x}}_w$ correspond to the actual end-effector pose with respect to the tunnel and back-wall pose. To achieve the tunnel and backwall effects described above without enforcing a specific desired motion dynamics, the desired end-effector velocity and acceleration, $\dot{\mathbf{x}}_d$ and $\ddot{\mathbf{x}}_d$, are set equal to zero.

The pose errors are transformed to a moving frame $[\mathbf{X}_T, \mathbf{Y}_T, \mathbf{Z}_T]$ using the $\mathbf{A}_t$, whose expression is given by

$$\mathbf{A}_t = \begin{bmatrix} \mathbf{R}_T^B & \hat{\mathbf{p}}_T \mathbf{R}_T^B \\ \mathbf{O}_{(3 \times 3)} & \mathbf{R}_T^B \end{bmatrix}^T. \quad (2)$$

The interaction control provides two assistive contributions: $\mathbf{a}_T = \mathbf{A}_t^{-1}\mathbf{K}_T\mathbf{A}_t\dot{\mathbf{x}}_T$ maintains the robot end-effector

actual position ($\mathbf{p}_a(t)$) close to the nearest position on the planned path ($\mathbf{p}_T(t)$), namely the tunnel control action, and $\mathbf{a}_W = \mathbf{A}_t^{-1}\mathbf{K}_W\mathbf{A}_t\dot{\mathbf{x}}_w$ moves patients forward to complete the trajectory in T_f seconds, also called as the back-wall contribution. Stiffness matrices $\mathbf{K}_T = \text{diag}\{0, k_r, k_r, k_{p\phi}, k_{p\phi}, k_{p\phi}\}$ and $\mathbf{K}_W = \text{diag}\{\hat{k}_w, 0, 0, 0, 0, 0\}$ encode the robot behavior for tunnel and back-wall contributions, respectively, where k_r manages the radial stiffness around the tunnel, $k_{p\phi}$ is the orientation control gain and k_w is the back-wall stiffness along the X_T axis. The back-wall stiffness k_w varies based on the relationship between the current end-effector and back-wall positions. Specifically

$$k_w = \begin{cases} k_r, & \langle \tilde{\mathbf{p}}_w, \mathbf{X}_T \rangle \geq 0 \\ 0, & \langle \tilde{\mathbf{p}}_w, \mathbf{X}_T \rangle < 0. \end{cases} \quad (3)$$

Here, $\langle \cdot, \cdot \rangle$ denotes the dot product, $\tilde{\mathbf{p}}_w$ is the position error computed as the difference of the current end-effector position $\mathbf{p}_a$ and the desired one returned by the trajectory planner $\mathbf{p}_w$ (see Fig. 2), and $\mathbf{a}_w$ accelerations are applied when the current end-effector position is behind the back-wall position.

The term $\mathbf{K}_D = \text{diag}\{k_d, k_d, k_d, k_{d\phi}, k_{d\phi}, k_{d\phi}\}$ defines the damping behavior of the impedance controller, which ensures stable non-oscillatory motion of the robot.

Figure 2 reports a graphical representation of the control action generated by the controller.

As evident, the proposed control strategy includes the definition of the parameter k_r which modulates the *AL* the robot is providing to the patient. In the context of tunnel control, k_r specifically regulates movement orthogonal to the desired trajectory, determining how strictly the user is constrained within the movement path. By adjusting k_r , the system controls the degree of freedom around the trajectory, effectively managing both the *AL* and the task difficulty [29]. This selective tuning enables the user to maintain the intended trajectory while progressing toward the assigned target. The core of the proposed approach relies on modifying this control parameter through supervised learning, where the labels are informed by the physiotherapist, as described in detail in Section II-A6.

On the other hand, the parameters T_f , T_p , and k_ϕ were kept constant for all participants to preserve uniformity in the experimental protocol. Notably, the patients began the task before T_p seconds and completed it in less than T_f . Therefore, these parameters only act as lower and upper bounds for task execution. The parameter k_ϕ was included to maintain a rigid end-effector orientation throughout the task, ensuring standardized and ergonomic posture of patients during the session.

3) Patient Multimodal Monitoring: To obtain a comprehensive understanding of the patient's condition and support adaptive interventions, the system captures kinematic, kinetic, and physiological data.

- Kinematic measurements include movement speed and positions, which help assess motor functionality, monitor progress [30], and quantify deviations from the desired trajectory during task execution.
- Kinetic measurements involve the force exerted by the patient on the robot, providing insight into the intensity

and nature of physical interaction, as well as exertion and resistance levels [31].

- Physiological measurements are recorded via wearable sensors presented in Section II-B and include HR, RR, and GSR. These signals provide valuable information on both physical and emotional states, reflecting arousal, cognitive load, and unconscious behavioral responses [32], [33], [34], [35]. For example, changes in GSR are associated with emotional arousal and cognitive effort, while increases in HR and RR often accompany heightened emotional states such as excitement, anxiety, or stress [36].

By integrating these multimodal measurements, the system allows physiotherapists to monitor the patient in real time and adapt the robot-assisted therapy to the patient's individual motor, physical, and psychophysiological needs, supporting more personalized and effective rehabilitation interventions [37], [38].

4) *Feature Extraction*: For each movement, twenty-five features are calculated from the raw signals. Following the framework provided by [18], they can be divided into three groups: biomechanical (9 features), physiological (15 features), and the *AL* itself. The details of the extracted features are listed below:

- Biomechanical features: they describe the forces and the movements applied in the human-robot interaction. The nine features computed are: the mean force, the maximum force, the standard deviation of force, the smoothness computed as the Mean Arrest Period Ratio (*MAPR*) as

$$MAPR = \frac{T_{vel}}{T_{tot}}, \quad (4)$$

where T_{vel} represents the duration during which the patient's movement speed exceeds 10% of the reached peak speed, while T_{tot} is the overall movement duration [13], the mean of the position error (*PE*), where *PE* is computed as

$$PE = \|\mathbf{p}_T(t) - \mathbf{p}_a(t)\|, \quad (5)$$

the maximum *PE*, the standard deviation of the *PE*, the elapsed time in executing the movement. Lastly, the power (*w*) was calculated by multiplying the velocity and force for each movement as

$$w = \sum_{t=1}^T (\langle \mathbf{F}(t), \mathbf{v}(t) \rangle \cdot \Delta t), \quad (6)$$

where $\mathbf{F}$ and $\mathbf{v}$ are the interaction force and the velocity computed at the time instant t , and Δt is the integration time interval.

- Physiological features: they must be computed starting from the collected HR, RR, and GSR. The seven features extracted from the cardiac activity, being NN the time interval between two consecutive R peaks, are: the mean Inter-Beat Interval (IBI), the standard deviation of the IBI, the square root of the mean squared differences of successive NN intervals (RMSSD), the number of interval differences of successive NN greater than 50 ms divided by the total number of NN (pNN50), the

mean HR, the standard deviation of HR, the maximum HR. The three features extracted from the respiratory activity are: the mean of RR, the standard deviation of RR, the maximum RR. Starting from the collected raw GSR, the Skin Conductance Level (SCL) is computed by applying a 4th order low-pass Butterworth filter, with a cutoff frequency of 0.1 Hz. The five features extracted from the SCL are: the mean SCL, the standard deviation of SCL, the maximum of SCL, the mean derivative of SCL and the mean second derivative of SCL.

- Assistance Level (*AL*): the current *AL* provides an essential feature to be fed into the PSPA module.

5) *Graphical User Interface*: A custom GUI enables physiotherapists to monitor therapy progress through quantitative evaluation by tracking parameters collected at each repetition. In agreement with the physiotherapists involved in the study, the interface displays the information considered most meaningful and easy to interpret during rehabilitation. They are the mean HR, RR, and the interaction force over each repetition. The GUI also allows entering patient information (ID, sex, age) and visualizing the current *AL*. Leveraging their knowledge of the patient and the therapeutic relationship developed through direct interaction, the physiotherapist integrates this understanding with the quantitative indicators to update the *AL* every nine movements, increasing, maintaining, or decreasing it as appropriate (see Fig. 3). This combination of personal expertise and real-time data ensures that adjustments are tailored to the patient's needs throughout the session. At the boundary levels (High or Low), some options are disabled: when *AL* is at the minimum value (Low), the "decrease" option is unavailable, and when *AL* is at the maximum value (High), the "increase" option is unavailable.

6) *Physiotherapist-Supervised Parameter Adaptation Module*: The PSPA module represents the core of the proposed approach and is based on the application of a supervised model. The therapist's input has the highest priority and overrides any actions or suggestions proposed by the autonomous PSPA module. This ensures that the therapist's expertise and decisions are paramount in guiding the therapy process, and providing professional supervision to the automated system. The AI module is based on a classification problem, in which there are three classes representing the following actions:

- *IC*: to increase the *AL*,
- *KC*: to keep the *AL* constant,
- *DC*: to decrease the *AL*.

The output of this module is represented by the *AL* to be sent to the robot to vary the stiffness and adapt its behavior to the user's needs. Once the desired *AL* value is predicted, the robot gradually adjusts control stiffness k_r to reach this target smoothly, avoiding sudden force changes that could compromise interaction safety.

To this purpose, four classification models were selected and compared: Linear Discriminant Analysis (LDA) [18], Support Vector Machine (SVM) [39], K-Nearest Neighbour (KNN) [40], and Multi-Layer Perceptron (MLP). In particular, LDA and SVM are among the most effective and simple

solutions for classifying physiological signals [41], [42], while KNN is widely used in processing nonlinear data [43]. A neural-network-based MLP was included to capture complex nonlinear relationships among the input features [44]. The applied MLP consists of two fully connected layers with 128 neurons each and ReLU activation, followed by a softmax output layer for classification into three classes. The network is trained with categorical cross-entropy loss and the Adam optimizer, using early stopping to prevent overfitting.

Beyond the choice of the specific classifiers, the structure of the problem itself, where the model selects one of three discrete actions based on expert-provided examples, naturally aligns with the BC framework. It learns a policy that reproduces expert behavior directly from demonstrations and can be formulated as a supervised learning problem mapping observed states to expert actions [45]. Subsequent works have shown that BC can operate without explicit system dynamics, relying solely on expert observations or actions [46], and theoretical analyses have framed imitation learning within structured prediction and no-regret online learning, inherently accommodating discrete decision spaces [47]. In robotics, BC has been widely adopted for high-level discrete decision-making, as demonstrated by conditional imitation learning approaches [48]. Within this framework, a classifier inferring one among the actions *IC*, *KC*, or *DC* from expert input can be regarded as a BC model, as it aims to replicate the therapist's decision strategy.

B. Experimental Setup

To validate the proposed PSPA approach, the experimental setup was composed of:

- the KUKA LightWeight Robot 4 + implementing the control strategy presented in Section II-A2 through Robot Operating System (ROS) Kinetic middleware on Ubuntu 16.04 LTS, running at 100 Hz. Specifically, the dimensions of the matrices and arrays of Eq. (1) are: 7×1 for q , $\dot{q}$, τ_c , $\mathbf{y}$, $\mathbf{g}(\mathbf{q})$; 6×1 arrays for a_T , a_W , $\tilde{\mathbf{x}}_T$, and $\tilde{\mathbf{x}}_W$; 7×7 matrix for $\mathbf{B}(\mathbf{q})$, $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})$, $\mathbf{F}_v$, and $\mathbf{F}_s$; 6×6 matrix for $\mathbf{A}_t$, $\mathbf{K}_T$, $\mathbf{K}_D$ and $\mathbf{K}_W$; 6×7 matrix for $\mathbf{J}_A^*(\mathbf{q})$. Three different *ALs* were selected with $k_r = \{100, 500, 1000\}$ N/m, corresponding to low, medium, and high, respectively. Selecting specific robot stiffness levels is important for ensuring comparability and reproducibility, thus facilitating easier comparisons of results to evaluate the effectiveness of the rehabilitation treatment [49]. These three levels were chosen because they produce perceptible differences in the assistive force generated by the controller, allowing the physiotherapist to effectively modulate the assistance and provide meaningful supervision during the learning process [28]. The other control parameters were set as $T_F = 10$ s, $T_P = 2$ s, $k_{p\phi} = 800$ Nm/rad, $k_d = 10$ Ns/m, and $k_{d\phi} = 10$ Nms/rad. Furthermore, a software safety measure interrupts the session if the interaction force exceeds 50 N, maintaining the interaction within safe limits for human-robot contact and minimizing the risk of soft tissue injury [50];

- a purposely designed ergonomic flange that is intended to be held by the affected subject's hand and is attached to the robot end-effector;
- a virtual reality game showing the trajectory to be performed and the actual position;
- a wearable physiological monitoring system: the BioHarness 3.0 chest belt, developed by Zephyr™ Technology and working at 1 Hz, is worn against the skin at the height of the sternum to measure the heart and respiration activities, and the GSR is measured with a frequency of 52 Hz by using two electrodes of the Shimmer 3 GSR+ Unit placed on the index and middle fingers of the unaffected hand;
- a custom GUI that allows and helps the therapists to take the action.

A picture of the experimental setup is provided in Fig. 3.

C. Experimental Protocol

The experimental protocol was designed to achieve the aforementioned objectives: first, to identify the most effective input modality and the best-performing classifier to support the rehabilitation process. This involved evaluating various input options and algorithms to determine which combination yielded the highest accuracy and reliability. Second, the study aimed to compare the proposed approach ability to tailor the *AL* against an existing *PB* method. By doing so, the validation assessed whether the new approach could match or surpass the accuracy of the previously clinically established technique in tuning the robot assistance during motor therapy.

The experiments were carried out in a real clinical setting providing robot-aided rehabilitation on ten orthopedic patients (3 males (M), 7 females (F); age 69.80 ± 7.61). The severity of their impairment was quantified by means of the Constant–Murley Score (CMS) and the Disability of the Arm, Shoulder and Hand (DASH) clinical scales. Although the enrolled patients presented different types of shoulder lesions (Table I), these conditions are typically managed through a similar rehabilitation pathway, which focuses on restoring joint mobility, strength, and motor control. For this reason, all participants underwent the same robot-aided therapy protocol, consisting of standardized point-to-point upper-limb movements with the robotic device. This guarantees that, despite the heterogeneity of the lesions, the therapy context and the clinical significance of the collected data remain consistent. The details of the patients' characteristics are expressed in Table I. Each patient, who had no previous experience in the use of robots for upper limb rehabilitation, was enrolled together with their physiotherapist. In total, ten physiotherapists participated in the study, with each patient being supervised by a different therapist during the rehabilitation session.

Each experiment consists in simulating a robot-aided rehabilitation session in which participants are instructed to execute nine cycles of nine 3D point-to-point trajectories. These trajectories initiate from a starting position and extend towards nine distinct targets distributed across three different heights [51]. All the patients were asked to wear the physiological monitoring systems at the beginning of the experiment. At the end of each cycle of nine movements, the therapist was

TABLE I
CHARACTERISTICS OF THE ENROLLED PATIENTS

	Age	Sex	Lesion	Limb	DASH	CMS
P1	64	F	Rotator cuff tendinopathy	L	55	45
P2	77	F	Rotator cuff tendinopathy	R	56	54
P3	59	M	Traumatic anterior dislocation	R	37	53
P4	81	M	Reverse shoulder prosthesis	L	69	33
P5	65	F	Rotator cuff repair	R	67	39
P6	72	F	Frozen shoulder contracture syndrome	L	74	29
P7	66	M	Rotator cuff repair	R	47	45
P8	63	F	Reverse shoulder prosthesis	R	68	29
P9	71	F	Reverse shoulder prosthesis	R	29	51
P10	80	F	Reverse shoulder prosthesis	R	28	59

asked to choose the most appropriate *AL* to tune the robot intervention by performing an action among IC, KC, or DC *AL* [13].

The study was conducted under Ethical Committee approval (Ethical Approval N. 03/19 PAR ComEt CBM) and in accordance with the Declaration of Helsinki. All participants have been adequately informed about the purpose of the study and gave their written informed consent.

The size of the acquired dataset is 1436×25 samples (observations $\times$ features). For the PSPA module computation, data from all ten patients were used within a subject-independent cross-validation framework to train and offline validate the classifiers. Prior to classification, feature normalization was performed. Specifically, features were standardized independently for each subject by applying z-score normalization (zero mean and unit variance), computed using only the samples belonging to the same subject. This procedure prevents inter-subject information from leaking while preserving subject-specific feature distributions. Such subject-wise normalization is particularly important in physiological and biomechanical data, where absolute feature magnitudes can vary substantially across individuals [52]. Moreover, this preprocessing step is essential for distance-based and margin-based classifiers, such as KNN and SVM, which are highly sensitive to feature scale and can otherwise be biased toward features with larger dynamic ranges [53].

Initially, an evaluation was conducted to determine which set of features yielded the best results. Specifically, four modalities were analyzed: using only kinematics and kinetics features (biomechanical), using only physiological features (physiological), using both aforementioned categories (multimodal, i.e., biomechanical, physiological, along with *AL*), and

an additional modality, referred to as optimized. The optimized modality was obtained through a ReliefF feature selection procedure [54]. ReliefF is a filter-based feature selection algorithm, which ranks features according to their relevance to the output variable independently of the classifier employed. In this approach, the ReliefF weight associated with each feature was computed, and only features with positive values were retained. This yielded a compact set of variables identified as informative by the algorithm, reducing redundancy and focusing the model on the most relevant predictors.

Model assessment was conducted on data from ten subjects using a subject-independent Leave-One-Subject-Out (LOSO) protocol. In each iteration, data from one subject were reserved as the test set, while data from the remaining nine subjects constituted the model selection set.

Within each model selection set, an inner stratified k -fold cross-validation (with $k = 5$) was performed on the therapist-provided action labels to compare different supervised classifiers and input configurations, including single-modality, multimodal, and an optimized feature set obtained by applying ReliefF feature selection procedure. Stratification was adopted to account for the uneven distribution of *AL* adjustments, which reflects patient-specific needs. This inner validation step provided balanced performance estimates and was used to identify the classifier-feature configuration yielding the highest F1-score. The choice of $k = 5$ represents a standard compromise between statistical robustness and computational efficiency [55].

Once the best-performing configuration was identified, a final PSPA model was trained on all data from the nine subjects in the model selection set and, subsequently, used to generate predictions for the held-out test individual. This procedure was repeated for each subject, ensuring a fully subject-independent evaluation in which predictions were always obtained on individuals unseen during both model selection and training.

For each test subject, PSPA predictions were then compared against the physiotherapists' ground-truth decisions to quantify both classification accuracy and alignment with the physiotherapists' decisions. In contrast, PB decisions were computed directly on the test subject data and evaluated using the same performance metrics, enabling a consistent comparison between the proposed data-driven behavioral cloning approach and the classical analytical baseline.

Several PB adaptation strategies have been proposed in robot-aided rehabilitation, ranging from early rule-based schemes relying on a single performance metric to more recent approaches that integrate multiple kinematic and dynamic indicators or embed performance evaluation within adaptive or model-based controllers [21], [56]. However, these modern solutions are not yet widely adopted in clinical settings. To ensure transparency, clinical interpretability, and comparability with the literature, the present study adopts the original PB formulation introduced by [19], which remains the most established reference baseline in robot-aided therapy, as evidenced in [1]. In this classical rule-based strategy, PB decisions are derived from relative variations of a single mechanical performance indicator, such as the power.

To evaluate the PB choices, the power (w) for the patients excluded from the training phase was calculated as in Eq. 6.

Subsequently, the changes in power between consecutive blocks of movements (each consisting of nine round trips) were analyzed to assign labels. The power difference for each block was computed as $\Delta w_i = w_i - w_{i-1}$, where w_i and w_{i-1} represent the mean power of the movements within the current and previous blocks, respectively. Based on the methodology outlined in [19], specific rules were established to identify the PB choice for each block (L_i), considering the magnitude of change in power

$$L_i = \begin{cases} DC & \text{if } \Delta w_i > 0.1 \cdot w_{i-1} \\ IC & \text{if } \Delta w_i < -0.1 \cdot w_{i-1} \\ KC & \text{if } -0.1 \cdot w_{i-1} \leq \Delta w_i \leq 0.1 \cdot w_{i-1} \end{cases}. \quad (7)$$

The choice of the 0.1 threshold in Eq. 7 was motivated by the need to adapt the PB decision rule to the characteristics of the power signal w , whose magnitude and physiological variability are substantially larger than those of typical normalized performance metrics. For this reason, an absolute threshold would not be appropriate, and a relative tolerance band was adopted. In particular, variations within $\pm 10\%$ of the previous block value were considered to reflect fluctuations inherent to the force-velocity measurement rather than meaningful performance changes. This proportional interval provides a scale-appropriate slack region and ensures stable PB labels across consecutive blocks, preventing spurious transitions caused by minor oscillations in the power signal. Moreover, this threshold has been set following the choices taken in [19].

All PSPA experiments were executed in a fully deterministic manner by fixing the random seeds across Python, NumPy, and TensorFlow, and by enabling deterministic operations in TensorFlow. Data shuffling, neural network weight initialization, and early stopping were also controlled with fixed seeds to ensure reproducibility of the results.

1) Performance Indicators: The performance of the models (both PB and those tested for PSPA) was evaluated using the F1-score, a statistical measure defined as the harmonic mean of precision and recall. This metric balances both aspects of model performance: precision, which is the ratio of true positive predictions to the total number of positive predictions, and recall, the ratio of true positive predictions to the total number of actual positive instances. The F1-score is calculated using the formula

$$\text{F1-score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}.$$

The F1-score ranges from 0 to 1, with a score of 1 indicating perfect precision and recall, while a score of 0 reflects the worst possible performance.

Given the imbalanced nature of the dataset, where certain classes are underrepresented compared to others, the F1-score provides a balanced measure that considers both precision and recall [57]. By incorporating the model ability to correctly identify positive instances and the accuracy of its positive predictions, the F1-score ensures a robust evaluation across

all classes, thereby mitigating the bias towards the majority class that can occur with metrics like accuracy.

Moreover, the $\Delta \text{F1-score} = \text{F1-score}_{\text{PSPA}} - \text{F1-score}_{\text{PB}}$ was calculated to highlight the differences in the classification performance of the two approaches. In addition, the Pearson correlation coefficient (ρ) was computed between the true block-wise labels and the predictions of both the PSPA and PB models. This analysis was included to explicitly assess the degree of alignment between the autonomous model outputs and the adaptation choices made by the physiotherapist [58]. Rather than measuring point-wise classification accuracy, the correlation coefficient quantifies the linear relationship between the predicted and actual sequences of AL , thus capturing how well the models follow the temporal trends and relative variations observed in the clinician's decisions. Correlation values range from -1 to 1 , where 1 indicates strong positive relationship, 0 indicates no linear relationship, and -1 indicates strong negative relationship.

2) Statistical Analysis: The Shapiro-Wilk test was applied to the collected data to verify the distribution. Due to the non-normal distribution, a non-parametric statistical method, specifically the Wilcoxon rank-sum test with a significance level set to 0.05 , was employed in two different analyses to compare the performance of:

- the different modalities (i.e., biomechanical, physiological, multimodal, optimized) for each classifier;
- the AI algorithms, in order to determine the statistical significance of the differences across their F1 values, specifically considering the best-performing modality for each algorithm;
- the subject-wise mean F1-scores obtained with the PSPA approach versus those obtained with the PB approach;
- the subject-wise Pearson correlation coefficients between predicted and actual labels for PSPA and PB.

Given the presence of multiple comparisons, the Bonferroni correction was employed to control the family-wise error rate. The resulting p-values from the Wilcoxon tests were then used to assess the statistical significance of the observed differences.

III. RESULTS AND DISCUSSIONS

The evolution of different parameters recorded from a patient during three cycles of nine point-to-point movements are shown in Fig. 4. Considering the k_r variation, shown at the third row of the left column, the curves on the left illustrate how adjustments in k_r influence trajectory tracking accuracy and control force response over time. This effect is also visible in the 3D plots on the right, where both the actual and desired trajectories are represented in the cartesian space for each k_r , providing a spatial perspective on trajectory adherence in three dimensions. Moreover, in the graph depicting k_r , it is evident that the adjustment in AL is implemented gradually, ensuring smooth changes in assistance and avoiding abrupt shifts in trajectory or interaction force. In particular, an increase in the stiffness parameter k_r leads to a noticeable reduction in trajectory error, meaning the robot can more accurately follow the desired path. However, as assistance increases, so lightly does the interaction force applied by the robot on the user, enhancing support during movement.




Fig. 4. A representative timeframe of a subject in which the therapist has experienced the three available *ALs*. In particular, three consecutive cycles of nine point-to-point movements are represented. On the left side, time series plots display the evolution of the actual and desired trajectories along the *z*-axis, alongside plots of the error norm ($|err|$) and force norm ($|F|$) as a function of the varying stiffness parameter k_r . On the right side, 3D plots show the spatial paths of both the actual and desired trajectories under different k_r levels.




Fig. 5. Statistical results applying Wilcoxon test with Bonferroni correction on F1-score. The statistical test was applied between the four modalities (i.e. physiological vs biomechanical, physiological vs multimodal, physiological vs optimized, biomechanical vs multimodal, biomechanical vs optimized, multimodal vs optimized) within the same classifier. ** means that the statistical difference between those groups has achieved a p -value < 0.01 .

Building on these recorded parameters and observed interactions, the next analysis evaluates the classification performance obtained. First of all, the feature selection procedure described in Section II-C led to the identification of 11 informative features. Only features with a positive weight after ReliefF application were included in the optimised set. The resulting optimized feature set includes: *AL*, mean SCL, maximum of SCL, maximum HR, mean HR, mean force, maximum *PE*, mean *PE*, maximum force, mean RR, maximum RR. The F1-scores for each classifier, trained on the dataset of all subjects using inner 5-fold cross-validation, across the four different input modalities (i.e., biomechanical, physiological, multimodal, optimized), are reported in Fig. 5. A detailed examination of the offline performance of these classifiers shows that both the multimodal and the optimized feature sets achieve the highest F1-scores across all classifiers, outperforming the biomechanical and physiological modalities. The optimized feature set, consisting of 11 features, achieves performance comparable to or slightly better than the full multimodal set, while using fewer features, which reduces computational cost.

MLP achieves the highest performance with the optimized set ($70.53 \pm 2.85\%$), followed by the multimodal set ($67.71 \pm 3.30\%$), biomechanical ($62.00 \pm 3.77\%$), and physiological features ($50.88 \pm 3.53\%$). For KNN, the F1-score with the optimized feature set is $66.74 \pm 3.01\%$, outperforming the multimodal set ($65.16 \pm 2.71\%$), physiological features ($52.70 \pm 3.57\%$), and biomechanical features ($59.42 \pm 2.85\%$). SVM reaches an F1-score of $59.62 \pm 2.79\%$ with the optimized set, slightly below the multimodal configuration ($60.68 \pm 3.03\%$), and well above physiological ($48.59 \pm 3.72\%$) and biomechanical ($53.84 \pm 4.60\%$) features. LDA achieves an F1-score of $58.58 \pm 3.26\%$ for the optimized set, slightly below the multimodal configuration ($59.49 \pm 4.31\%$), and higher than physiological ($43.21 \pm 2.64\%$) and biomechanical ($56.75 \pm 3.65\%$) features.

Overall, these results indicate that the optimized feature set consistently improves or maintains performance across classifiers, while multimodal features also provide competitive results. Moreover, this finding underscores the advantage of combining biomechanical and physiological information into a comprehensive dataset, which improves the classifiers ability to predict the physiotherapist's decision. By reducing the feature set to the optimized 11 features, similar performance is maintained while lowering computational cost, which is particularly relevant for real-time applications. This is confirmed also by the statistical analysis, as no significant difference was found between multimodal and optimized. Thus, a reduced, optimized feature set can maintain the predictive performance of the full multimodal dataset while decreasing computational complexity.

Additionally, statistical comparisons using paired Wilcoxon tests with Bonferroni correction indicate that, for all classifiers, the optimized feature set significantly outperforms both biomechanical, physiological, and optimized modalities (p -value < 0.01).

Among the evaluated classifiers, the MLP yields the highest F1-scores when trained on the optimized feature set, followed by KNN, SVM, and LDA. The superior performance of the MLP can be attributed to its ability to model complex non-linear relationships among input features through its multi-layer architecture, which is particularly suited to capturing interactions across heterogeneous multimodal signals. KNN and SVM achieve intermediate F1-scores, reflecting their capacity to handle non-linear decision boundaries and high-dimensional feature spaces, respectively, although their performance remains below that of the MLP in this experimental setting. LDA, which relies on a linear separability assumption, exhibits the lowest F1-scores among the four classifiers.

Consequently, the optimized feature set was used to build the PSPA module, and MLP was selected as the classifier due to its robust and superior performance with the optimized features. The MLP-based PSPA module was then trained on the entire dataset using the optimized feature set and evaluated on the held-out subjects of the LOSO cross-validation, comparing its F1-score performance with the PB approach.

Specifically, the decisions made by the PSPA and PB modules are compared to the actual physiotherapists' choices,




Fig. 6. Comparison of choices (increase AL (IC), keep AL constant (KC), decrease AL (DC)) taken by Physiotherapist, Physiotherapist-Supervised Parameter Adaptation (PSPA) strategy, and performance-based (PB) approach for each block of movements and a testing patient. Block 9 is missing because it was the final block of the rehabilitation session, and the physiotherapist did not make a selection as the session was concluding.

revealing that, when averaged across the ten test subjects, the PSPA module outperforms PB in terms of F1-score ($\Delta F1 = 15.40 \pm 30.33\%$) and shows a stronger relationship with the physiotherapists' decisions, as quantified by the Pearson correlation coefficient ($\rho = 0.56 \pm 0.21$), whereas PB exhibits lower and predominantly negative correlations ($\rho = -0.12 \pm 0.43$). Statistical testing using the Wilcoxon signed-rank test indicates that the improvement in correlation is significant (p -value < 0.01), while the F1-score increase shows a positive trend that does not reach significance level (p -value $= 0.08$).

However, the large standard deviation of $\Delta F1$ -score indicates notable inter-subject variability, and PSPA does not outperform PB for all subjects in terms of classification accuracy. Despite this, PSPA shows a significantly and consistently higher correlation with the physiotherapists' decisions, suggesting a stronger alignment with expert decision trends than the PB approach. This result suggests that the proposed approach better captures the temporal trends and decision-making patterns of therapists, providing a more coherent alignment with expert judgment compared to the PB strategy.

Figure 6 illustrates, for a representative test patient, a comparison between the AL decisions made by the physiotherapist, the predictions generated by the proposed PSPA approach, and the decisions computed by the PB method. Considering eight blocks of movements, the comparison reveals that the PSPA module agrees with the physiotherapist's decisions in 5 blocks. On the other hand, the PB module aligns only in 3 blocks. This corresponds to F1-score and correlation of 75.00% and $\rho = 0.79$ (p -value $= 0.02$) for PSPA, respectively, compared to 22.50% and $\rho = -0.54$ (p -value $= 0.16$) for PB.

These results for the example subject illustrate the general trends observed across all participants. The PSPA module is more aligned with the physiotherapist's decision-making, reflecting the adaptive and responsive choices a therapist would make during a rehabilitation session, as it shows a higher relationship with the physiotherapists' decisions, both in terms of F1-score and correlation. The higher ρ of the AI-driven PSPA module with the physiotherapist's choices can be attributed to its ability to dynamically adjust to the complexity of movements through data-driven analysis, effectively integrating a broader range of features. This adaptability enables PSPA to capture intricate patterns within the multimodal data that are otherwise challenging to interpret with simpler approaches. In contrast, the PB approach relies solely on

predefined power-based criteria and thresholds, which limit its flexibility in adapting to different user needs.

The improvement in correlation, which is statistically significant across subjects ($p < 0.01$), suggests that PSPA better mirrors the context-sensitive judgment of physiotherapists. While the F1-score increase shows only a positive trend ($p = 0.08$), the example subject emphasizes the potential of PSPA to produce more accurate and responsive assistance-level adjustments.

A. Limitations

The present findings should be interpreted in light of some limitations. First, the dataset is relatively small, as only ten orthopedic patients were enrolled. Although a subject-independent LOSO validation framework was adopted, the limited sample size limits statistical power and may compromise the stability of the comparison among classifiers. Consequently, the identification of MLP as the best-performing model should be interpreted with caution, as different datasets or larger datasets may yield different classifier rankings. However, the proposed behavioral cloning methodology and its conceptual framework remain unchanged, and applying the same pipeline to different datasets may yield different quantitative results while preserving the validity of the overall approach. Additionally, the study considered only three discrete assistance levels during label collection, which may limit the resolution of patient-specific adaptation. These limitations highlight the need for larger and more diverse datasets, as well as an expanded set of assistance configurations.

IV. CONCLUSION

This paper presented the development of a novel strategy for adapting the AL in robot-aided rehabilitation replicating the decision-making process of a physiotherapist. To achieve the specified aim, experimental sessions were carried out in a real clinical environment. In this session, the therapist is included in the loop to label the patients' data in real-time by setting the appropriate AL , making it possible to develop an autonomous system that relies on measurements of both biomechanical and physiological quantities collected during the therapy session. This process implements a BC approach, where the system learns to reproduce the therapists' decision-making strategy by mapping patient states to the corresponding AL adjustments.

Ten orthopedic patients and their respective physiotherapists were enrolled, and the labeled dataset of all patients was used to train four different supervised models. Statistical analyses identified a best-performing configuration for building the PSPA module on the collected data. Analysis of the results revealed that both the multimodal and optimized feature sets achieved strong classification performance. Among the evaluated classifiers, MLP showed the highest performance on the current dataset; however, this result should be interpreted in relation to the limited sample size and the specific experimental conditions. Importantly, the optimized feature set, comprising a reduced number of features, maintained performance comparable to the full multimodal set while

lowering computational cost, making it an efficient choice for practical implementation.

Moreover, compared to the clinically established PB methodology, PSPA demonstrated an improved ability to follow the physiotherapists' decision trends, as reflected by significantly higher correlation values with expert choices.

Future efforts will focus on expanding the available dataset to enhance the generalization of the algorithms and interpreting the model decisions using explainable AI techniques. Another important direction is the evaluation of the PSPA module as a decision-support tool for less experienced therapists, helping to standardize clinical choices and reduce reliance on individual expertise. Finally, the module will be integrated into adaptive assistance protocols in upcoming clinical studies, enabling an evaluation of how effectively it replicates therapist decisions and contributes to patient outcomes.

ACKNOWLEDGMENT

Rita Molle is the Ph.D. Student enrolled in the National Ph.D. in Artificial Intelligence, XXXVIII Cycle, course on health and life sciences, organized by Università Campus Bio-Medico di Roma.

REFERENCES

- [1] H. Rodgers et al., "Robot assisted training for the upper limb after stroke (RATULS): A multicentre randomised controlled trial," *Lancet*, vol. 394, no. 10192, pp. 51–62, Jul. 2019.
- [2] C. Duret, A.-G. Grosmaire, and H. I. Krebs, "Robot-assisted therapy in upper extremity hemiparesis: Overview of an evidence-based approach," *Frontiers Neurol.*, vol. 10, p. 412, Apr. 2019.
- [3] Z. Chen et al., "Robot-assisted arm training versus therapist-mediated training after stroke: A systematic review and meta-analysis," *J. Healthcare Eng.*, vol. 2020, pp. 1–10, Oct. 2020.
- [4] A. Bertomeu-Motos, A. Blanco, F. J. Badesa, J. A. Barios, L. Zollo, and N. Garcia-Aracil, "Human arm joints reconstruction algorithm in rehabilitation therapies assisted by end-effector robotic devices," *J. NeuroEngineering Rehabil.*, vol. 15, no. 1, pp. 1–11, Dec. 2018.
- [5] R. Molle, C. Tamantini, and L. Zollo, "Artificial intelligence in upper limb robot-aided physical rehabilitation: A systematic review," *ACM Trans. Hum.-Robot Interact.*, vol. 15, no. 2, pp. 1–35, Mar. 2026.
- [6] Y. Zhang et al., "Research on adaptive impedance control technology of upper limb rehabilitation robot based on impedance parameter prediction," *Frontiers Bioeng. Biotechnol.*, vol. 11, Jan. 2024, Art. no. 1332689.
- [7] T. Hoffmann, J. Lewis, and C. G. Maher, "Shared decision making should be an integral part of physiotherapy practice," *Physiotherapy*, vol. 107, pp. 43–49, Jun. 2019.
- [8] M. N. Ozer and T. Kroll, "Patient-centered rehabilitation: Problems and opportunities," *Crit. Rev. Phys. Rehabil. Med.*, vol. 14, nos. 3–4, p. 18, 2002.
- [9] T. He, Y. Chen, L. Wang, and H. Cheng, "An expert-knowledge-based graph convolutional network for Skeleton - based physical rehabilitation exercises assessment," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 32, pp. 1916–1925, 2024.
- [10] C. J. Hasson, J. Manczurowsky, E. C. Collins, and M. Yarossi, "Neurorehabilitation robotics: How much control should therapists have?," *Frontiers Human Neurosci.*, vol. 17, May 2023, Art. no. 1179418.
- [11] D. M. Mahfouz, O. M. Shehata, E. I. Morgan, and F. Arrichiello, "A comprehensive review of control challenges and methods in end-effector upper-limb rehabilitation robots," *Robotics*, vol. 13, no. 12, p. 181, Dec. 2024.
- [12] F. van Dellen, T. Aurich-Schuler, and R. Labruyère, "Within - and between-therapist agreement on personalized parameters for robot-assisted gait therapy: The challenge of adjusting robotic assistance," *J. NeuroEng. Rehabil.*, vol. 20, no. 1, pp. 1–9, Jun. 2023.
- [13] R. Molle et al., "Exploring priority parameters in physiotherapist decision models for tailoring robot-aided rehabilitation," *Int. J. Social Robot.*, vol. 17, no. 11, pp. 1–16, Nov. 2025.
- [14] L. Luo, L. Peng, C. Wang, and Z.-G. Hou, "A greedy assist-as-needed controller for upper limb rehabilitation," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 30, no. 11, pp. 3433–3443, Nov. 2019.
- [15] S. Pareek, H. J. Nisar, and T. Kesavadas, "AR3n: A reinforcement learning-based assist-as-needed controller for robotic rehabilitation," *IEEE Robot. Autom. Mag.*, vol. 31, no. 3, pp. 74–82, Sep. 2024.
- [16] F. J. Badesa et al., "Dynamic adaptive system for robot-assisted motion rehabilitation," *IEEE Syst. J.*, vol. 10, no. 3, pp. 984–991, Sep. 2016.
- [17] R. Molle, C. Tamantini, C. Lauretti, E. M. Romano, and L. Zollo, "An online reinforcement learning method to improve control adaptability in robot-aided rehabilitation," *Eng. Appl. Artif. Intell.*, vol. 161, Dec. 2025, Art. no. 112248.
- [18] D. Novak, M. Mihelj, J. Zihlerl, A. Olensek, and M. Munihi, "Psychophysiological measurements in a biocooperative feedback loop for upper extremity rehabilitation," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 19, no. 4, pp. 400–410, Aug. 2011.
- [19] H. I. Krebs et al., "Rehabilitation robotics: performance-based progressive robot-assisted therapy," *Auto. Robots*, vol. 15, no. 1, pp. 7–20, Jul. 2003.
- [20] R. Colombo, I. Sterpi, A. Mazzone, C. Delconte, and F. Pisano, "Taking a lesson from Patients' recovery strategies to optimize training during robot-aided rehabilitation," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 20, no. 3, pp. 276–285, May 2012.
- [21] L. Pezeshki, H. Sadeghian, A. Mohebbi, M. Keshmiri, and S. Haddadin, "Personalized assistance in robotic rehabilitation: Real-time adaptation via energy-based performance monitoring," *IEEE Trans. Autom. Sci. Eng.*, vol. 22, pp. 13298–13309, 2025.
- [22] E. D. Oña, J. M. Garcia-Haro, A. Jardón, and C. Balaguer, "Robotics in health care: Perspectives of robot-aided interventions in clinical practice for rehabilitation of upper limbs," *Appl. Sci.*, vol. 9, no. 13, p. 2586, Jun. 2019.
- [23] Y. Murakami et al., "New artificial intelligence-integrated electromyography-driven robot hand for upper extremity rehabilitation of patients with stroke: A randomized, controlled trial," *Neurorehabilitation Neural Repair*, vol. 37, no. 5, pp. 298–306, May 2023.
- [24] C. Castellini, A. Arquer, and J. Artigas, "SEMG-based estimation of human stiffness: Towards impedance-controlled rehabilitation," in *Proc. 5th IEEE RAS/EMBS Int. Conf. Biomed. Robot. Biomechanics*, Aug. 2014, pp. 604–609.
- [25] B. Luciani, A. Pedrocchi, F. Braghin, and M. Gandolla, "A therapist-inspired approach to stiffness modulation in rehabilitation exoskeletons," in *Proc. Int. Conf. Rehabil. Robot. (ICORR)*, May 2025, pp. 180–185.
- [26] Q. Wang, Z. He, J. Zou, H. Shi, and K.-S. Hwang, "Behavior cloning and replay of humanoid robot via a depth camera," *Mathematics*, vol. 11, no. 3, p. 678, Jan. 2023.
- [27] I. Bratko, T. Urbančič, and C. Sammut, "Behavioural cloning: Phenomena, results and problems," *IFAC Proc. Volumes*, vol. 28, no. 21, pp. 143–149, Sep. 1995.
- [28] C. Tamantini et al., "Patient-tailored adaptive control for robot-aided orthopaedic rehabilitation," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2022, pp. 5434–5440.
- [29] L. Liu, S. Leonhardt, C. Ngo, and B. J. E. Misgeld, "Impedance-controlled variable stiffness actuator for lower limb robot applications," *IEEE Trans. Autom. Sci. Eng.*, vol. 17, no. 2, pp. 991–1004, Apr. 2020.
- [30] V. D. Tran, P. Dario, and S. Mazzoleni, "Kinematic measures for upper limb robot-assisted therapy following stroke and correlations with clinical outcome measures: A review," *Med. Eng. Phys.*, vol. 53, no. 1, pp. 13–31, Mar. 2018.
- [31] M. Abujelala, A. Lioulemes, P. Sassaman, and F. Makedon, "Robot-aided rehabilitation using force analysis," in *Proc. 8th ACM Int. Conf. Pervasive Technol. Rel. Assistive Environments*, Jul. 2015, pp. 1–2.
- [32] R. Levenson, "The autonomic nervous system and emotion," *Emotion Rev.*, vol. 6, no. 2, pp. 100–112, 2014.
- [33] M. Sharma, S. Kacker, and M. Sharma, "A brief introduction and review on galvanic skin response," *Int. J. Med. Res. Professionals*, vol. 2, no. 6, pp. 13–17, Dec. 2016.
- [34] N. Nourbakhsh, F. Chen, Y. Wang, and R. A. Calvo, "Detecting users' cognitive load by galvanic skin response with affective interference," *ACM Trans. Interact. Intell. Syst.*, vol. 7, no. 3, pp. 1–20, Sep. 2017.
- [35] L. Shu et al., "A review of emotion recognition using physiological signals," *Sensors*, vol. 18, no. 7, p. 2074, Jun. 2018.
- [36] H. A. Shehu, M. Oxner, W. N. Browne, and H. Eisenbarth, "Prediction of moment-by-moment heart rate and skin conductance changes in the context of varying emotional arousal," *Psychophysiology*, vol. 60, no. 9, p. 14303, Sep. 2023.

- [37] C. L. Bethel, K. Salomon, R. R. Murphy, and J. L. Burke, "Survey of psychophysiology measurements applied to human-robot interaction," in *Proc. 16th IEEE Int. Symp. Robot Human Interact. Commun.*, Jan. 2007, pp. 732–737.
- [38] A. P. Hills, N. Mokhtar, and N. M. Byrne, "Assessment of physical activity and energy expenditure: An overview of objective measures," *Frontiers Nutrition*, vol. 1, p. 5, Jun. 2014.
- [39] F. J. Badesa, R. Morales, N. Garcia-Aracil, J. M. Sabater, A. Casals, and L. Zollo, "Auto-adaptive robot-aided therapy using machine learning techniques," *Comput. Methods Programs Biomed.*, vol. 116, no. 2, pp. 123–130, Sep. 2014.
- [40] A. Lanata et al., "A new smart-fabric based body area sensor network for work risk assessment," in *Proc. IEEE Int. Workshop Metrology Ind. 4.0 IoT*, France, Jun. 2020, pp. 187–190.
- [41] J. Jeon and H. Cai, "Classification of construction hazard-related perceptions using: Wearable electroencephalogram and virtual reality," *Autom. Construct.*, vol. 132, Dec. 2021, Art. no. 103975.
- [42] G. Ramos et al., "Fatigue evaluation through machine learning and a global fatigue descriptor," *J. Healthcare Eng.*, vol. 2020, pp. 1–18, Jan. 2020.
- [43] I. Saini, D. Singh, and A. Khosla, "QRS detection using K-nearest neighbor algorithm (KNN) and evaluation on standard ECG databases," *J. Adv. Res.*, vol. 4, no. 4, pp. 331–344, Jul. 2013.
- [44] M. Rosoł, M. Młyńczak, and G. Cybulski, "Granger causality test with nonlinear neural-network-based methods: Python package and simulation study," *Comput. Methods Programs Biomed.*, vol. 216, Apr. 2022, Art. no. 106669.
- [45] M. Bain and C. Sammut, "A framework for behavioural cloning," in *Machine Intelligence 15*. Oxford, U.K.: Oxford Univ. Press, 1995. [Online]. Available: <https://api.semanticscholar.org/CorpusID:10738655mymargin>
- [46] F. Torabi, G. Warnell, and P. Stone, "Behavioral cloning from observation," in *Proc. 27th Int. Joint Conf. Artif. Intell.*, Jul. 2018, pp. 4950–4957. [Online]. Available: <https://api.semanticscholar.org/CorpusID:23206414>
- [47] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proc. 14th Int. Conf. Artif. Intell. Statist.*, 2011, pp. 627–635.
- [48] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2018, pp. 4693–4700.
- [49] L. Zhang, S. Guo, and Q. Sun, "Development and assist-as-needed control of an end-effector upper limb rehabilitation robot," *Appl. Sci.*, vol. 10, no. 19, p. 6684, Sep. 2020.
- [50] N. Kovičič, H. Gattringer, A. Müller, and M. Brandstötter, "Physics guided machine learning approach to safe quasi-static impact situations in human-robot collaboration following the power and force limiting method of the ISO/TS 15066 standard," in *Proc. Int. Design Eng. Tech. Conf. Comput. Inf. Eng. Conf.*, Aug. 2023, p. V010T10A024.
- [51] C. Tamantini, F. Cordella, C. Lauretti, and L. Zollo, "The WGD—A dataset of assembly line working gestures for ergonomic analysis and work-related injuries prevention," *Sensors*, vol. 21, no. 22, p. 7600, Nov. 2021.
- [52] J. Fdez, N. Guttenberg, O. Witkowski, and A. Pasquali, "Cross-subject EEG-based emotion recognition through neural networks with stratified normalization," *Frontiers Neurosci.*, vol. 15, Feb. 2021, Art. no. 626277.
- [53] M. Ahsan, M. Mahmud, P. Saha, K. Gupta, and Z. Siddique, "Effect of data scaling methods on machine learning algorithms and model performance," *Technologies*, vol. 9, no. 3, p. 52, Jul. 2021.
- [54] R. J. Urbanowicz, M. Meeker, W. La Cava, R. S. Olson, and J. H. Moore, "Relief-based feature selection: Introduction and review," *J. Biomed. Informat.*, vol. 85, pp. 189–203, Sep. 2018.
- [55] S. Raschka, "Model evaluation, model selection, and algorithm selection in machine learning," 2018, *arXiv:1811.12808*.
- [56] L. Zhang, S. Guo, and F. Xi, "Performance-based assistance control for robot-mediated upper-limbs rehabilitation," *Mechatronics*, vol. 89, Feb. 2023, Art. no. 102919.
- [57] P. Pramokchon and P. Piamsa-nga, "A feature score for classifying class-imbalanced data," in *Proc. Int. Comput. Sci. Eng. Conf. (ICSEC)*, Jul. 2014, pp. 409–414.
- [58] B. Luciani, L. Roveda, F. Braghin, A. Pedrocchi, and M. Gandolla, "Trajectory learning by therapists' demonstrations for an upper limb rehabilitation exoskeleton," *IEEE Robot. Autom. Lett.*, vol. 8, no. 8, pp. 4561–4568, Aug. 2023.